Before You Issue an AI Governance RFP: What to Build Before You Buy

A Governance Architecture Primer for Regulated Enterprises

Audience Chief Compliance Officers · VP Risk & Audit · General Counsel · COO · Chief Innovation Officers · AI Governance Leaders

Verticals Life Sciences / CROs · Regulated Infrastructure · Federal / Public Sector

Topic AI governance architecture as an organizational design problem, not a procurement one

Corevident

corevident.com · adeola@corevident.com

An organization gets the budget approved. The mandate is clear. An AI governance program is required. The next step, almost universally, is a search for vendors.

That search quickly becomes disorienting. Every vendor claims to deliver AI governance: one monitors ChatGPT and Copilot usage, another maps AI risk into an existing SOC 2 framework, a third produces a policy document library. They are all, technically, doing something related to AI. None of them is doing the same thing.

The buyer has no framework for evaluating the differences because the foundational question has not been asked yet: what does governance actually need to accomplish here, and what organizational design makes that possible? That question is not a procurement question. It is an architecture question. And it has to come first.

The Procurement Trap

Organizations with AI governance mandates are routing budget toward tools. The tools are real and the controls they provide are legitimate, but controls operate downstream of governance decisions that have not yet been made.

A monitoring tool tells you when an AI system behaves outside a defined boundary. But someone has to define the boundary. A risk register tracks AI-related risks. But someone has to determine which risks are in scope, who owns them, and what constitutes an acceptable threshold. An access control framework restricts which models are available to which users. But someone has to establish the decision rights that make those restrictions meaningful rather than arbitrary.

> *Buying controls without governance architecture is retrofitting oversight onto decisions that were never deliberately designed. When an audit surfaces, a regulatory inquiry arrives, or an autonomous system takes an action no one can account for, the absence of the underlying architecture becomes visible, and consequential.*

What the AI Governance Market Actually Sells

The AI governance market is not one category. It is at least four distinct activities sold under the same label. Each is legitimate. None of them starts with the foundational question that determines whether governance actually works.

Market Category	What It Addresses	What It Does Not Define
Security and Access Control	Monitors AI usage, blocks unauthorized models, flags data exfiltration and shadow AI adoption	• Who authorized the boundary • Who is accountable if an AI-driven decision produces a regulated outcome
Policy Documentation	Produces acceptable use policies, explainability reports, impact assessments, incident response narratives	• Operational decision authority is not established • Policies derived without governance architecture, document decisions that may not have been made
GRC Extension	Maps AI risk into ISO 27001, SOC 2, or existing model risk frameworks	• Escalation logic for autonomous systems is undefined • Inherited frameworks were not designed for agents that act across organizational boundaries
Data Governance	Catalogs data assets, defines ownership, lineage, classification, and access policy	• Decisional accountability is not addressed • A mature data governance program does not answer who authorized an agent to act, what the escalation threshold is, or who is named in the audit record

KEY INSIGHT

Enterprise AI lifecycle platforms have recognized this gap. Sophisticated toolkits now exist spanning model documentation, bias and drift monitoring, compliance mapping, and audit trail automation. These are well-engineered products. What they share, without exception, is a starting assumption: that a governance architecture already exists. The lifecycle framework tells you how to monitor a boundary and document a model's lineage. It does not tell you who drew the boundary, under what authority, or what organizational design makes the documentation defensible. That foundational work is left for teams to connect the dots on their own, and most organizations do not have a defined path to do so.

None of these categories answers the foundational question: who owns the decision when AI acts?

The Layer No Vendor Sells: The Decision Authority Layer

There is a foundational layer of governance work that has no vendor, no off-the-shelf certification, and no procurement path. Corevident calls it Layer 0. It is organizational design work: the structural decisions that make every downstream activity coherent.

Layer 0 asks and answers: Who has authority to approve an AI system for a given use case, and under what conditions? When an autonomous system takes an action, who is accountable for the outcome, and is that documented in a form that will survive regulatory review? At what threshold does a decision require human review before execution? What constitutes an adequate evidence trail for a regulated industry audit?

These are not tool configuration questions. They are governance design decisions. Until they are made explicitly, the tools, policies, and controls built on top of them are built on assumptions that may be invisible until tested.

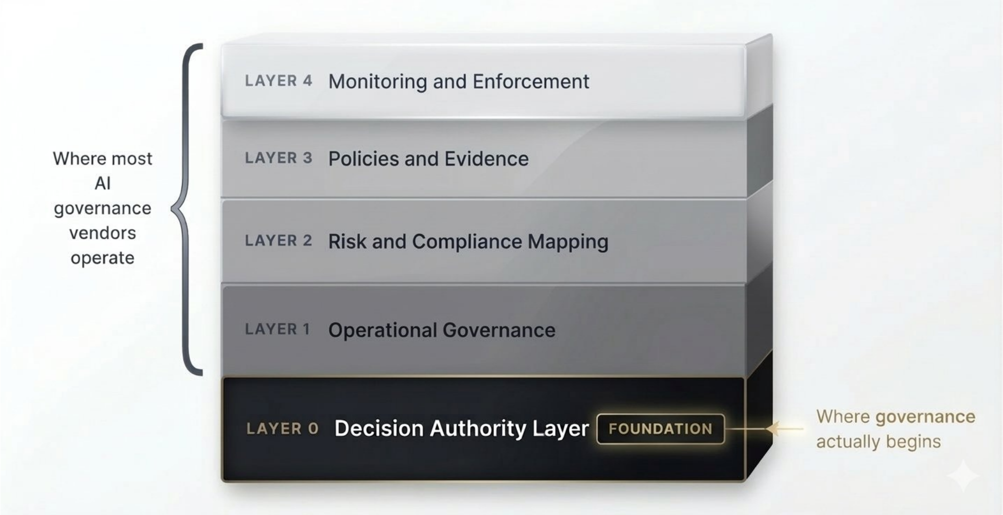
LAYER 4 Monitoring and Enforcement
LAYER 3 Policies and Evidence
LAYER 2 Risk and Compliance Mapping
LAYER 1 Operational Governance
LAYER 0 Decision Authority Layer
FOUNDATION
Where most AI governance vendors operate
Where governance actually begins

> *Layer 0 is the architecture that gives every subsequent governance decision its authority and its auditability. It is not a checklist or a policy template. It has to be built first.*

Why Regulated Industries Cannot Skip This Step

General enterprise contexts surface governance gaps slowly: inconsistent AI adoption, audit friction, accountability gaps. In regulated industries, the same gaps surface as findings.

FDA inspections do not accept system logs as audit trails. The distinction between a log and a moment-of-decision governance record is a finding, not a recommendation. SR 11-7 model risk management requirements assume human accountability structures that agentic systems disrupt by design. Board-level AI oversight obligations require documented decision rights that most current governance programs have not produced. Clinical trial environments require validated governance frameworks before AI systems touch production workflows, and validation requires documentation of architectural decisions that may never have been made explicitly.

The agentic AI risk landscape compounds this exposure further. Standards bodies have begun cataloging risks that emerge specifically from autonomous operation: unsupervised autonomy where agents act without defined human review thresholds; tool choice hallucination where an agent selects the wrong tool in a multi-step chain; cascading errors in multi-agent systems where no human checkpoint exists between steps; and accountability gaps where no defined owner exists for an autonomous decision. In regulated industries, each of these is a potential audit finding, not a theoretical scenario.

The absence of governance architecture does not stay theoretical. It surfaces in inspections, in regulatory inquiries, and in deployment timelines that collapse when the certification burden arrives without adequate preparation.

> *In regulated industries, each of these is a potential audit finding, not a theoretical scenario.*

The Question That Changes the Frame

Before evaluating any AI governance vendor, before issuing an RFP, scheduling a demo, or comparing feature sets, the right question is this:

> *Do we have a governance architecture that defines how AI decisions are authorized, overseen, and evidenced in ways that would hold up under regulatory scrutiny?*

If yes, vendor evaluation is appropriate. The tools, policies, and GRC extensions become meaningful because they operate inside a defined architecture. If no, or unclear, the vendor evaluation is premature. Selecting tools before the architecture exists produces controls that enforce undefined boundaries and compliance postures that cannot be defended when tested.

The governance architecture work is not a consulting engagement that competes with vendor selection. It is the precondition that makes vendor selection rational.

AI Governance RFP Readiness Scorecard

Score each question 0 (not defined) · 1 (partially defined) · 2 (clearly defined and documented)

Question	Score
Do we know who owns the decision when an AI system acts?	0 1 2
Have we defined escalation thresholds for AI-driven decisions?	0 1 2
Can we reconstruct the authority structure behind any AI outcome?	0 1 2
Do we know which AI use cases require human review before execution?	0 1 2
Do we have a defined evidence trail that would survive regulatory scrutiny?	0 1 2

Scoring: 0–4 = Governance architecture not yet defined · 5–7 = Partial governance structure · 8–10 = Governance architecture in place

If your score is below 8, governance architecture likely needs to be addressed before issuing an AI governance RFP.

Corevident LLC

Governance architecture and operating model design for regulated enterprises navigating agentic AI.

Life Sciences / CROs · Regulated Infrastructure · Federal / Public Sector

corevident.com · adeola@corevident.com

Corevident