AI GOVERNANCE SERIES | REGULATORY PERSPECTIVE

AI Governance Is Becoming a Multi-State Operating Model Problem.

The challenge is no longer writing an AI policy. It is operating AI consistently across jurisdictions with different, non-identical requirements.

By Greg Aldrich | Global CIO & Strategic Advisor | June 2026

Most discussions about AI governance focus on principles. Regulators are increasingly focused on implementation.

Connecticut's new omnibus AI law, Senate Bill 5, signed into law on May 27, 2026, is less interesting for its individual provisions than for what it represents in aggregate: the accelerating fragmentation of AI regulation across state lines. WilmerHale published a detailed analysis of SB 5 on June 4 that is worth reading carefully if your organization operates in Connecticut or deploys any of the systems the law covers.

But the strategic implication extends well beyond Connecticut. The framework is not converging. It is fragmenting. Organizations are now confronting a more complicated reality: the governance challenge is shifting from writing AI policy to operating AI consistently across jurisdictions, with requirements that are sufficiently similar to create compliance pressure yet different enough to resist a unified solution.

> ***"The governance challenge is no longer writing an AI policy. It is operating AI consistently across jurisdictions with different requirements, and that is a board-level conversation."***

What Connecticut's Law Actually Does

Connecticut SB 5 is notable for its breadth. Where most state AI laws address a single use case, Connecticut's 39-section statute covers AI companions, frontier model safety, automated employment decision tools, generative AI content provenance, and social media platform regulation, with effective dates staggered from October 2026 through January 2028.

Several provisions warrant specific attention from an enterprise governance perspective:

- * The automated employment-related decision technology framework pertains to technologies that are a "substantial factor in making or materially influencing" employment decisions. This is a different standard than Colorado's "consequential decision" threshold, meaning the same tool deployed in the same way could be in scope under one state's law and out of scope under another's.
- Section 13 specifies that use of an AEDT "shall not be a defense" against employment discrimination claims, and that courts may consider evidence of anti-bias testing when evaluating such claims. This is an incentive to implement proactive bias testing programs now, before the disclosure obligations become operational in October 2027.
- The developer-deployer allocation model requires AEDT developers to provide deployers with all information necessary for compliance or to contractually assume those obligations. This has immediate consequences for AI procurement agreements: organizations buying AI tools need to review vendor contracts for compliance and allocation of responsibilities now, not when the effective dates take effect.
- Provenance requirements for generative AI audio, image, and video content, embedded metadata that allows consumers to assess whether content was created or materially altered by AI, take effect October 1, 2026. Organizations deploying generative AI in content creation workflows have a compressed runway.
- * The AI companion provisions, effective January 2027, include specific minor protections and prohibitions on manipulative techniques that are defined differently from analogous provisions in Oregon's SB 1546. The same product may face different compliance requirements in each state.

WilmerHale's analysis makes a point that deserves to be underlined: "Companies operating in multiple states should not assume that compliance with one state's AI law will satisfy another's." That is not a lawyerly hedge. It is an operational warning that leads directly to a more extensive fragmentation problem.

The Fragmentation Problem

Connecticut is not an outlier. It is a data point in a pattern that has been building for several years and is now accelerating.

State / Framework	Scope	Key Standard	Effective
Colorado ADMT Act	Automated decision-making in consequential decisions	"Consequential decision" standard; consumer rights included	2026

California (CPRA / AB 1008)	Broad data and AI governance; automated decision-making rules	Consumer rights; opt-out of automated decisions	In effect / evolving
Connecticut SB 5	Omnibus: AI companions, frontier models, AEDT, provenance, social media	"Substantial factor" AEDT standard; no consumer rights in AEDT	Oct 2026 – Jan 2028 (phased)
Oregon SB 1546	AI companions	Different companion definition than Connecticut	2025
Illinois AIAA	AI in employment decisions	Bias audit requirements	In effect
EU AI Act	Risk-tiered framework for AI systems	Risk classification: unacceptable / high / limited / minimal	Phased 2024–2026

The operational consequence of this table is not theoretical. An organization deploying AI hiring tools in Colorado, California, and Connecticut simultaneously is operating against three different standards for what constitutes a covered decision, three different disclosure obligation structures, and three different enforcement models, with varying definitions of the same underlying concepts.

This is the multi-state operating model problem. It is not fundamentally a legal problem, though legal counsel is essential. It is a governance architecture problem: how do you build AI systems, workflows, and controls that can satisfy varying requirements across jurisdictions without creating a separate compliance program for each one?

> ***"Key statutory definitions vary across jurisdictions. Close parsing is still needed. Do not assume compliance with one state's AI law satisfies another's."***

The Hidden Cost: Governance Strategy Drift

There is a deeper problem embedded in this regulatory fragmentation that I want to name directly, because it is one that most AI governance conversations are not having yet.

Many organizations built their AI governance programs in 2023 and 2024, when the regulatory landscape was less defined, against assumptions that may already be outdated. They wrote policies calibrated to the EU AI Act's risk-tiered framework. They built disclosure workflows around Colorado's ADMT Act. They designed vendor governance processes around the NIST AI RMF. Each of these was reasonable at the time.

The problem is that the regulatory environment has continued to develop as those governance programs have remained largely static. Connecticut's omnibus approach initiates provenance requirements, companion regulations, and an independent verification pilot program that do not map cleanly to present frameworks. California's continuing evolution of its AI rules adds additional complexity. And the working group established by Connecticut SB 5, tasked with proposing legislation to regulate general-purpose AI models and considering whether to establish a dedicated technology court for AI and data privacy matters, signals that the regulatory surface will continue to expand.

This is governance strategy drift: the gap between the regulatory environment your governance program was designed for and the regulatory environment it is currently operating in. It is the AI governance equivalent of a security posture that has not been updated since the last major review, technically compliant with what was required then, increasingly out of sync with what is required now.

> ***"Governance strategy drift is the gap between the regulatory environment your governance program was designed for and the one it is currently operating in. Most organizations are living in that gap and don't know it."***

The organizations most exposed to governance strategy drift tend to share a recognizable profile: they completed an AI governance initiative, published a policy, stood up a committee, and moved on. Governance became a document rather than a living system. The policy does not update itself when Connecticut passes a new law. The committee does not automatically modify its scope when provenance requirements take effect in October. The vendor contracts do not self-amend to reflect new developer-deployer allocation obligations.

This is why I have persistently argued that governance must be designed as operational infrastructure rather than policy documentation. A policy document ages. An operating system adapts.

What the Multi-State Reality Requires

For organizations operating across multiple jurisdictions, the regulatory fragmentation creates four distinct operational requirements that most current administrative frameworks do not address:

- A jurisdiction-by-jurisdiction AI system mapping. The WilmerHale analysis is explicit: organizations should "conduct a jurisdiction-by-jurisdiction mapping exercise that identifies precisely which AI systems and use cases fall within scope under each applicable law, paying particular attention to differences in defined terms, covered decision types, and exemptions." This is not a one-time legal exercise. It is an ongoing governance function that needs to be owned, resourced, and updated as the regulatory landscape develops.
- Contract review for developer-deployer allocation. Connecticut's developer-deployer model is not unique — similar allocation frameworks appear in Colorado and California. Organizations buying AI tools need to understand who has contractually assumed compliance obligations, whether the developer is providing the information necessary for the deployer to comply, and what indemnification protections exist when the allocation is unclear. Most AI procurement agreements were not drafted with this specificity.
- Proactive bias testing documentation. Section 13's anti-defense provision creates a clear incentive to build and document anti-bias testing programs before disclosure obligations arrive. Courts "may consider evidence of anti-bias testing" when evaluating discrimination claims. That evidence needs to exist and be retrievable. Organizations that treat bias testing as a compliance checkpoint rather than an ongoing operational practice will find themselves at a disadvantage when claims arise.
- Provenance and allocation management. Organizations also need processes to manage provenance requirements for generative AI content and the developer-deployer allocation obligations that affect procurement. These requirements need to be tracked, reviewed, and

embedded in workflows before the effective dates take effect. Without that, the governance program will lag the regulatory environment it is meant to cover.

- * Governance refresh cadence. The organizations that manage multi-state AI compliance effectively will be those that treat their governance programs as living systems with structured review cycles. Not annual reviews triggered by an incident. Quarterly assessments calibrated to the regulatory calendar, effective dates, working group recommendations, enforcement actions, that update the governance posture before requirements arrive rather than after they are violated.

The Board-Level Conversation

I want to be precise about why this is a board-level issue rather than a legal or compliance team issue.

The multi-state regulatory fragmentation creates strategic risk that operates at a different level than individual compliance obligations. Organizations that build AI capabilities on governance architectures that cannot conform to evolving regulatory requirements are making a strategic bet that the regulatory environment will stabilize, converge, or that their current posture will remain sufficient. That bet is increasingly difficult to justify.

The Connecticut SB 5 working group's mandate to consider a dedicated technology court for AI matters is a signal worth taking seriously. If that concept advances, or if other states employ similar approaches, the enforcement landscape for AI governance failures will become significantly more structured and more consequential. Organizations that have treated AI governance as a compliance checkbox will be poorly positioned when that enforcement landscape matures.

The board conversation is not primarily about Connecticut's specific provisions. It is about whether the organization's AI governance architecture is designed to adapt to a regulatory environment that will continue to evolve, and whether the governance program has the operational infrastructure to keep pace.

That is a different question than "are we compliant with Connecticut SB 5?" And it is the more important one.

What Good Looks Like

Organizations ahead of this curve tend to share a few common practices:

- They have inventoried their AI systems not just against risk categories but against jurisdictional scope, maintaining a running map of which systems are subject to which state frameworks.
- They have designated ownership for regulatory monitoring, someone whose job includes tracking state AI legislation, assessing its pertinence, and triggering governance updates before effective dates take effect.
- * They have reviewed AI vendor contracts specifically for developer-deployer allocation, information-sharing obligations, and the assignment of compliance responsibilities.
- They treat bias testing as an ongoing operational function with documentation, not a pre-launch checklist item.

- They have built governance review cycles into their operating calendar, calibrated to the regulatory timeline rather than to internal project schedules.

None of this requires a new platform or a new team. It requires a governance architecture that was designed to be a living system rather than a completed project.

Final Thought

Connecticut's SB 5 is a significant law that calls for careful attention from any organization operating within its scope. WilmerHale's analysis is the right starting point for comprehending its specific requirements.

But the more important conversation is the one Connecticut's law reinforces: the era of building a single AI governance program and expecting it to hold is over. The regulatory context is fragmenting faster than most governance frameworks can adapt to, and the organizations that recognize this as an operating-model problem, rather than a legal problem to be managed on a jurisdiction-by-jurisdiction basis, will be significantly better positioned than those that continue to treat governance as a document-based exercise.

The question is not whether your organization is compliant with Connecticut today. It is whether your governance architecture is designed to remain compliant with wherever the regulatory environment is twelve months from now.

> ***"The governance challenge is no longer writing an AI policy. It is operating AI consistently across jurisdictions with different requirements and building a governance system that adapts as those requirements evolve."***

Related reading: *WilmerHale: What Companies Should Know About Connecticut's New Omnibus AI Law (June 4, 2026)*

AI Governance Series: *The AI Governance Tightrope* | *The New Control Layer* | *The Operating Model for AI Governance* | *Three Frameworks* | *The Specification Owner* | *OutcomeOps: The Organizational Intelligence Layer* | *Your Employees Aren't Ignoring AI Policy* | *Agentic Orchestration is a Compelling Idea* | *AI Governance Is Becoming an AI Security Problem*

About the Author

Greg Aldrich is a Global CIO and Strategic Advisor with 30+ years of experience helping boards, executive teams, and C-suite stakeholders navigate IT strategy, AI governance, and digital transformation. He serves as a fractional CIO and Senior Strategy Advisor at Blue Tree Technology Group and Field CIO/CTO and Senior Transition Architect at SDS Consulting and currently serves as CIO at Andrew Wommack Ministries & Charis Bible College, where he chairs both the AI Committee and IT Steering Committee. He has advised organizations across financial services, gaming, healthcare, higher education, logistics, and nonprofit sectors.

Connect: linkedin.com/in/galdrich | Aldrich.ai