


The background of the entire page is a photograph of a grand, ornate legislative chamber, likely the U.S. House of Representatives. The room features high ceilings with intricate moldings and paintings, large columns, and tiered seating for members. The lighting is soft, and the overall atmosphere is formal and historic.

A CONFLUENCE PSG WHITE PAPER

Getting AI Regulation Right

Why a neutral, structured process, not an ordinary legislative process, is the most reliable path to AI policy that works



CONFLUENCE
POLICY & STRATEGY
GROUP

AUGUST 2026

Opening

A co-executive director of a disability-rights organization sat across the table from a lobbyist representing one of the world's largest technology companies. Two years of legislative fighting had put them on opposite sides of nearly everything. Neither had been asked, yet, to propose anything. The only question on the table that day was what each of them was actually trying to protect, and why it mattered. The advocate's concern was technical and specific: under existing law, meeting the burden of proof required to pursue a claim against every party responsible for an automated decision — a developer, a deployer, sometimes both — was often functionally impossible for the person harmed by it. By the time she finished explaining it, the lobbyist's position had visibly shifted. He said he had not thought about the problem that way, and that he believed the room could solve for it.



That shift did not happen because either person in the room was unusually persuasive. It happened because of the sequence: nobody at the table had been allowed to negotiate yet, only to explain.

That shift did not happen because either person in the room was unusually persuasive. It happened because of the sequence: nobody at the table had been allowed to negotiate yet, only to explain. Months later, that same lobbyist stood in front of the bill's legislative sponsors alongside the disability-rights advocate, the ACLU, and the head of the state's chamber of commerce, a coalition that would not have shared a table two years earlier, and made a case none of them would have made alone: they did not love every piece of the resulting bill, but they did not want it reopened, because every piece had been negotiated against every other piece, and pulling one thread risked the whole agreement.

Confluence PSG designs and runs rooms like that one for state and local government leaders across the country.

About this paper. This paper argues that states need a deliberate, structured process to govern how AI policy gets made, not what that policy should say. Confluence PSG takes no position on the substance of AI regulation. Its expertise is in the process that produces workable law when the underlying questions are too contested, too technical, or moving too quickly for ordinary legislative process to resolve well. This paper is current as of August 2026; given the pace at which state AI legislation is moving, some of the specific developments described below may have changed by the time of publication, a fact this paper treats as evidence for its own argument rather than a flaw to apologize for.

That argument rests on two related claims. Where parties are dug in and distrustful of one another, a neutral structured process can break gridlock that the ordinary legislative process has repeatedly failed to break; Colorado’s AI law is the clearest current example, and the section below sets out the mechanism in enough detail to be usable by a legislator, a governor’s office, or an attorney general’s office considering the same approach. More often, and more relevant

to most of the domains this paper examines, the same process resolves a quieter problem: converting broad agreement in principle into workable statutory language. Many AI policy fights are not standoffs between people who fundamentally disagree. They are rooms full of people who agree, within the first few sessions, that consumers deserve protection, that compliance obligations should be clear enough for businesses to actually meet, that the right parties should be accountable when something goes wrong. The difficulty surfaces once the definitions get written: what does it mean for an AI system to “materially influence” a decision it is only one input into, and should liability be broad and general or anchored specifically to violations of law that already exists. Colorado’s Working Group reached agreement on principle within its first few sessions and spent nine months on the definitions, a pattern worth naming directly: triggers, not principles, are where policy actually gets made or broken. That is not gridlock. It is the substantive work a five-month legislative session, run by people who hold other full-time responsibilities, is poorly positioned to do well.

Every domain examined in this paper is already being decided. Not eventually, not pending further study — now, this session, in statehouses across the country, whether or not any given legislature

Two Related Claims

01

A structured process can break genuine gridlock between distrustful parties.

02

It converts broad agreement in principle into workable statutory language.

has fully reckoned with what it's deciding. A legislator voting on a hiring-discrimination bill this year is setting the terms for how a resident finds out, or doesn't, why an algorithm denied them a job. A governor's staff drafting chatbot guidance this month is deciding whether a child in crisis is talking to something bound by any standard of care at all. These are not hypothetical future questions. They are live votes, live executive orders, live agency guidance, arriving faster than most legislatures have built the capacity to handle them well.




A bill introduced fully formed forces every affected party into a defensive posture before anyone has had the chance to actually understand what the others are trying to protect.

The standard tool for resolving contested policy — a bill drafted by a small number of people, introduced, debated in public for a matter of weeks, amended, and voted on — works when the underlying problem is stable and well understood by the people voting on it. AI regulation is neither. The technical expertise required to write a workable definition of “materially influences” or “frontier model” sits almost entirely outside the legislature, in companies, in advocacy organizations, in practitioners who see the technology's effects firsthand and have no formal seat at the drafting table. A bill introduced fully formed forces every affected party into a defensive posture before anyone has had the chance to actually understand what the others are trying to protect. That is not a flaw in any individual legislator's judgment. It is a structural mismatch between the shape of the problem and the shape of the tool being used to solve it, and it is why so many of the domains that follow show real, often avoidable problems accumulating state by state: internal contradictions, provisions that turn out to be technically unworkable once implementation begins, conflict that a better process could have caught while it was still cheap to resolve. Not every difference between states is evidence of a defective process; different values, markets, and institutions legitimately produce different answers, and this paper takes no position on which answer any state should reach. What a structured process changes is not the guarantee of a uniform outcome, but the reliability of the reasoning behind whatever outcome a given state lands on, and that outcome's odds of surviving contact with implementation.

There is a different path, and it is not theoretical. Colorado ran it. The section below describes exactly how, in enough detail to be usable, not just credible. What follows after that is a survey of eight areas where states are currently attempting to answer AI policy questions using the standard tool, at real cost, and a closing argument, built on everything the survey shows, for why the alternative Colorado already proved is the more reliable path available to any state facing the same problem now.


Federal Context

Over the past year, the federal government has made several attempts to settle who has authority to regulate AI. Each has stalled, and the pattern of stalling is itself relevant to any state weighing whether to wait for federal resolution before acting.



In December 2025, President Trump signed Executive Order 14365, establishing an AI Litigation Task Force within the Department of Justice and directing Commerce, the FCC, and the FTC to challenge state AI laws considered unduly burdensome. The order explicitly exempted three categories from that effort: child safety, data center infrastructure, and a state’s own procurement and use of AI. That carve-out matters for how this paper is organized, since several of the domains examined later are not part of the federal preemption fight at all.

In April 2026, xAI filed suit against Colorado challenging SB24–205, the state’s original AI law, on constitutional grounds. The Department of Justice moved to intervene on xAI’s side shortly afterward, the first instance of the federal government seeking directly to invalidate a state AI statute in court. That litigation remains unresolved, but the federal government’s posture is not: it has taken an active role in attempting to unwind a specific state’s approach to AI regulation.



Congress has twice attempted to preempt state AI law outright and failed both times. In July 2025, the Senate voted 99 to 1 to remove a ten-year moratorium on state AI regulation from the reconciliation bill. A second attempt, attached to the FY26 National Defense Authorization Act,

also failed to advance. On June 4, 2026, Representatives Jay Obernolte and Lori Trahan released a bipartisan discussion draft, the Great American AI Act, narrower than prior attempts in scope: a three-year preemption limited to laws governing model development rather than model use, paired with new federal transparency and audit requirements for the largest developers. It drew rejection within hours from labor organizations, consumer advocates, and a formal House Democratic commission. Three federal attempts. Three stalls. Twelve months.



A related fight over youth online safety has moved further through Congress than AI-specific preemption efforts but remains unresolved in a way that illustrates the same underlying pattern. The House passed the Kids Internet and Digital Safety Act in a bipartisan 267 to 117 vote on June 29, 2026, after months of negotiation that removed an earlier draft's broader state-preemption language in favor of a narrower standard preserving state laws that offer greater protection than the federal floor. The bill still faces uncertain prospects in the Senate, where the original bipartisan sponsors of the underlying Kids Online Safety Act have said the House version is effectively dead over its removal of a "duty of care" standard they consider essential. Even a heavily negotiated, twice-revised bipartisan effort has not resolved, after more than four years of attempts, a question a number of states have already been legislating on individually.

None of the foregoing is offered as an assessment of whether federal preemption would be sound policy. It is offered as a record of recent attempts and outcomes. Three federal efforts to resolve AI regulation authority, across two separate issues, have stalled within the past twelve months. A future attempt may succeed. States that structure their AI policy around the assumption that Washington will resolve the question are making a bet against recent history.

The Colorado Model

In May 2024, Colorado enacted SB24-205, the first comprehensive AI consumer protection law in the country. The law addressed the use of AI in consequential decisions, hiring, lending, insurance, healthcare, built around a duty-of-care standard: developers and deployers of high-risk AI systems were required to take reasonable steps to prevent algorithmic discrimination, conduct impact assessments, and report harms to the Attorney General.

Governor Polis signed the bill with stated reservations and requested that the legislature revisit it before it took effect. The 2025 regular session did not produce a revision. A special session convened specifically to address the law produced only a delay of its effective date. By the fall of 2025, Colorado had a law that satisfied few of its original stakeholders: a business community warning of competitive disadvantage, and consumer advocates unwilling to concede protections already won. Continued delay was no longer a viable strategy.

Governor Polis convened the Colorado AI Policy Working Group that fall: fourteen members representing constituencies that had spent two years in direct conflict with one another, consumer advocates, civil rights organizations, AI developers and deployers, business associations, and representatives of the healthcare and insurance sectors. The Working Group met nine times over five months. Midway through that process, the xAI lawsuit described earlier in this paper landed, adding fresh, external urgency to work that was already underway rather than prompting it.

14

Stakeholders

9

Sessions

5


months

The process rested on a deliberate structural choice at its outset: no draft was brought to the first meeting. Before any party proposed a solution, each was asked to explain what it was trying to protect and why, a sequence that produced the exchange described in the opening of this paper and shaped much of what followed. The group also maintained a formal distinction throughout its work between “ideas for consideration” and formal “proposals.” That distinction allowed participants to raise and explore options without the other side treating each option as an opening position to be countered, keeping the group in a mode of shared exploration longer than most negotiations allow, and surfacing each side’s underlying reasoning before anyone had staked out a position they would need to defend rather than reconsider.

“

A closed room is not the same as an unaccountable one.

The negotiated trades that followed were specific rather than symbolic. On liability, the group considered both joint-and-several liability, allowing a harmed consumer to pursue any party in the chain, and full insulation of either developers or deployers depending on their position in that chain. The



group settled instead on comparative fault: both developers and deployers could be held liable, but only in proportion to their actual contribution to the harm. To prevent that compromise from being negotiated away in practice, the law also bars parties from contractually shifting liability regardless of fault, closing off the leverage either side could otherwise exploit: a large developer conditioning licensing on a deployer's agreement to absorb liability the developer had actually caused, or a large deployer using its purchasing power to push liability back onto a smaller developer that couldn't afford to refuse.

On disclosure, the trade ran in the opposite direction: businesses accepted disclosure obligations for AI use they had previously preferred to keep unstated, in exchange for a tiered standard defining what triggers that obligation, rather than a single rule applying equal weight to a hiring decision and a routine scheduling function.

One design choice merits direct explanation rather than passing mention: the Working Group's sessions were closed to the public. Legislators were kept informed of the group's progress throughout but did not sit at the table. This was a deliberate decision, made after two prior public, legislator-led attempts to resolve the law's problems had failed. Part of the reason: participants in an open session must manage two audiences simultaneously, the people across the table and the constituents watching from outside it. A closed room is not the same as an unaccountable one. The Working Group held no legislative authority; it could only produce a recommendation, and the group required broad, tested support before finalizing any recommendation at all. Everything that followed was conducted in public: legislators debated the resulting bill, amended what they chose to amend, and voted on the record.

This was a deliberate choice suited to the problem at hand, not a fixed preference independent of it. Where the objective is broad public legitimacy for a decision affecting an entire community, an open process is typically the better tool, and Confluence has run engagements built that way. The relevant question in each case is not which format appears more transparent in the abstract, but which format is likely to produce authentic exchange given who is at the table and what they are being asked to say.

The Working Group's process also depended on participants who were never in the room. Advocates on multiple sides of the issue sustained public pressure throughout the group's work, in the press and in legislators' offices, and that pressure was instrumental both in prompting the Governor to convene the group and in keeping the legislature engaged through a third attempt at resolution.

Their role required holding a public position, which is a different function from negotiating language privately, and the process needed both functions operating simultaneously rather than the same people attempting both at once.

34–1

Senate vote

57–6

House vote

13

Days to law

SB26-189 was introduced on May 1, 2026, passed the Senate 34 to 1, passed the House 57 to 6, and was signed into law thirteen days after introduction. Outcomes of that speed and margin are uncommon in legislative practice generally and rare on the most contested issues specifically, where a fully drafted bill more often arrives as a surprise to most of the parties affected by it, with each subsequent amendment fight a public negotiation over what those parties are willing to lose.



The resulting statute is better understood as a modernization of existing consumer protection and anti-discrimination law than as a new category of AI-specific regulation.

The resulting statute is better understood as a modernization of existing consumer protection and anti-discrimination law than as a new category of AI-specific regulation. Its core commitment, that a person is entitled to know when an automated decision has been made about them and to have a means of challenging it if it is wrong, extends a right that already existed when the decision-maker was human. Senator Rodriguez, the bill's sponsor, has been candid that the final statute does not fully satisfy any single stakeholder in the room that built it. What it produced instead was a framework negotiated by the parties who would have to operate under it, each of whom gave up something specific to reach it.

The business community's own account of the process is worth noting directly. Brittany Morris Saunders of the Colorado Technology Association and Loren Furman of the Colorado Chamber of Commerce, writing publicly after the bill's passage, credited the structured, facilitated process, one that ensured every voice in the room was actually heard, as essential to reaching an agreement on an issue of that complexity. A business-association endorsement of a consumer-protective outcome is not a common occurrence, and it did not happen without the process described above producing it.

Colorado has continued to legislate on AI since the Working Group concluded. On May 29, 2026, the state signed a second AI-related statute, the Chatbot Safety Act, restricting AI's role in clinical mental health settings. That law is discussed in full in the domain addressing mental health and minor safety below.

None of the above is a claim that Colorado is alone in taking a thoughtful, strategic approach to AI or other policy. It is a demonstration of what is achievable, in both senses described at the opening of this paper, resolving genuine gridlock and doing the slower work of converting agreement into workable statutory language, when a state builds a deliberate process rather than relying on a five-month session and a floor vote to produce one. Each domain examined in the remainder of this paper describes a state, or several, still positioned where Colorado was in 2024, or still working through the definitional problem the Working Group spent nine months resolving, without a structured process built to do it.



What a Structured Process Requires, and Where It Doesn't Work

The Colorado story above illustrates several distinct design choices: no draft brought to the first meeting, interests explained before positions were taken, ideas for consideration kept separate from formal proposals, closed deliberation paired with public accountability, and a requirement of broad, tested support before anything was finalized. Naming those choices individually is useful. Naming what they add up to is more useful still.

A structured process of this kind requires five things to be in place before it starts. It needs a sponsor with real authority to convene it, someone who can bring reluctant parties to the table and whose backing signals the process is worth participating in seriously. It needs a defined mandate, a clear statement of what the group is and is not being asked to decide, so participants are not negotiating the scope of their own negotiation. It needs a deliberately constructed table, built to include the parties who are actually affected and the expertise actually required, not simply the loudest or best-resourced voices already in the room. It needs protected space, time and structure enough for evidence and understanding to precede proposals, rather than skipping straight to positions. And it needs a mechanism for testing whether an agreement actually holds, the role Colorado's supermajority requirement served, before that agreement becomes public and effectively irreversible.



Sometimes the most honest thing a good process can do is clarify exactly where and why people disagree, without resolving it.

Where any of those five is missing, the process is likely to underperform, and worth naming plainly rather than assuming this tool solves every problem it's pointed at. A convener without real authority produces a process stakeholders can safely ignore. A mandate left vague invites the negotiation to consume itself arguing about its own scope. A table built too narrowly reproduces the same power imbalances an open floor debate would have, just with better catering. And even a well-built process can fail for reasons outside its own design: if the political will to act is simply absent, if there is no realistic path from recommendation to enacted law, or if the underlying issue is not yet ripe, positions

have not softened enough anywhere in the system for a negotiated agreement to hold, a structured process will surface that clearly rather than manufacture consensus that is not really there. That is a legitimate outcome, not a failure of the method. Sometimes the most honest thing a good process can do is clarify exactly where and why people disagree, without resolving it.



A structured process also does not, by itself, neutralize the real differences in resources, technical capacity, and political access that participants bring into the room. Colorado's Working Group addressed this partly through design, insisting on interests before positions gave less-resourced participants room to be heard on substance before the negotiating started, and partly through deliberate table-building, ensuring disability rights advocates and consumer groups sat alongside developers and deployers rather than commenting on their proposals afterward. But a facilitator has real, ongoing responsibility here: treating lived experience as its own form of expertise, guarding against surprise proposals sprung on parties without time to respond, and keeping a place at the table, or at minimum a clear channel into it, for constituencies too diffuse or under-resourced to organize their own seat.

Reading the Domains That Follow

The eight domains that follow are not uniform in the kind of difficulty they present, and recognizing which kind is at play in each one is useful for understanding what a solution would actually require. Three patterns recur throughout.

In some domains, the difficulty is a genuine collision of values, not a misunderstanding waiting to be cleared up. Reasonable people disagree, on real grounds, about how much liability a developer should bear for harm a deployer caused, or where the line sits between a helpful AI tool and the unlicensed practice of a profession. No amount of technical clarity resolves a disagreement like this. It requires a process built to find a compromise every party can actually live with, not a process built to declare a winner.

In others, the difficulty is technical and definitional rather than a values conflict at all. The parties agree, often quickly, on the underlying principle. What remains is the harder, slower work of writing language precise enough to govern real behavior without either failing to protect anyone or making compliance impossible to achieve. This is the triggers-not-principles problem named above, and it recurs across more of the domains that follow than any other single pattern.

In still others, the difficulty is not disagreement or complexity at all. It is absence. No meaningful process has occurred, and the relevant decision is being made by default, through adoption that has simply outpaced any governing framework. This pattern shows up most starkly in how state governments are using AI on their own behalf, but it is present, to some degree, in nearly every domain that follows.

The same underlying tool, a structured process that brings the right expertise and the right competing interests into a room before positions harden into legislation, addresses all three patterns, though the specific work required differs meaningfully depending on which is actually at play. Each domain below is labeled with the pattern or patterns it represents, so the connection back to this framework stays explicit rather than left for the reader to infer.

DOMAIN MAP GETTING AI REGULATION RIGHT

How states are navigating the challenges of artificial intelligence

PART ONE

How We're Going to Allow AI to Be Used

- | | | |
|----|------------------------------|-------------------------------------|
| 01 | Consequential Decisions | Employment, lending, pricing |
| 02 | Mental Health & Minor Safety | Chatbots, therapy, minor protection |
| 03 | Frontier AI Safety | Model transparency laws |
| 04 | AI in Public Safety | Wildfire detection, policing |

PART TWO

How AI's Existence is Changing the Way we Live

- | | | |
|----|------------------|---------------------------|
| 01 | Data Centers | Siting, energy, water use |
| 02 | Workforce Impact | Career pathways, job loss |

Connective Threads – Issues that do not fit neatly into either category

- | | | |
|----|-------------------------|--------------------------------|
| 01 | Government's Own AI Use | Adoption without governance |
| 02 | AI & Synthetic Media | Deepfakes, likeness, elections |

PART ONE:

How We're Going to Allow AI to Be Used

The domains that follow share a common shape. In each, a policy question has emerged that is too technically dense, too fast-moving, or too dependent on specialized expertise for an ordinary legislative session to resolve well on its own. States are attempting to answer these questions anyway, one at a time, largely without reference to what other states are learning in parallel. What follows is not a comprehensive account of AI regulation nationally. It is a description of where the genuine, unresolved tension currently sits in each area, and why that tension is characteristic of the kind of problem a structured process is built to address.

Consequential Decisions

Primarily a definitional-complexity problem, what triggers a disclosure or liability obligation, compounded by a genuine values collision over how liability should be allocated between developers and deployers.

The question at the center of this domain predates AI by decades: when an automated system contributes to a decision about someone's employment, credit, housing, or insurance, what obligations does the party using that system owe to the person affected by it. Colorado's SB26-189, discussed above, is one state's answer. It is not the only one being written.

California has moved on two fronts simultaneously. The Civil Rights Council's employment regulations, effective October 2025, require employers using automated decision systems in hiring to retain records of the system's criteria and to conduct anti-bias testing. Separately, the California Privacy Protection Agency's regulations governing automated decision-making technology impose significant-decision obligations, including a consumer's right to opt out and to access an explanation, with phased compliance dates running through April 2027. The two regimes overlap in places and diverge in others, a detail that matters less for its own sake than for what it illustrates: even within a single state, AI-related consequential-decision rules are accumulating faster than they are being reconciled with one another.

Illinois has taken a narrower, disclosure-based approach specific to hiring, requiring employers to notify candidates when AI is used to analyze video interviews and to obtain consent before doing so. Connecticut's automated-employment-decision law addresses similar territory with its own definitional structure, and the differences between Connecticut's and Illinois's approaches, though each addresses roughly the same underlying concern, mean that an employer operating in both states is complying with two distinct sets of obligations for functionally the same hiring practice. Texas has gone a different direction structurally, folding AI hiring obligations into a broader technology-governance statute, HB149, sometimes referred to as TRAIGA, rather than legislating hiring practices as a standalone AI issue. Three states, three structurally different approaches to the same underlying question, developed independently and roughly simultaneously.

1.

Illinois

Disclosure and consent

- Employers must notify candidates when AI analyzes video interviews.
- Candidate consent is required.

2.

Connecticut

Separate automated-employment rules

- Uses its own definitions and compliance structure.
- Employers operating across states face different obligations for similar practices.

3.

Texas

Broad technology-governance model

- AI hiring obligations are incorporated into HB149, also referred to as TRAIGA.
- Hiring is regulated as part of a broader technology statute rather than through a standalone employment law.

Where no legislative process exists at all, the courts have begun to fill the gap. *Mobley v. Workday*, still pending in federal court as of this writing, tests whether a vendor providing AI-driven hiring tools can be treated as an “agent” of the employer for purposes of anti-discrimination liability, a question with no clear answer under existing law because existing law was not written with this arrangement in mind. A federal court resolving that question case by case, without the benefit of a legislative record built through hearings and stakeholder input, is not a substitute for a deliberate policy process. It is what happens in the absence of one.



Algorithmic pricing. A related but distinct pattern has emerged in how AI is used to set prices rather than to screen applicants, and it is worth treating as an extension of the same underlying question rather than a separate topic: an opaque system using data the consumer cannot see to make a decision that affects them, in this case what they pay rather than whether they are approved. Maryland has enacted restrictions on certain forms of AI-driven dynamic pricing. New York now requires disclosure when a price shown to a consumer has been set using their personal data. California’s AB325 amended the state’s antitrust law, the Cartwright Act, to address the use of shared pricing algorithms among competitors, a mechanism that raises coordination concerns distinct from the individual consumer-protection question the other two states are addressing. As with the domain’s employment and lending examples, the pattern is a genuine and growing area of consumer concern arriving in front of state legislatures faster than a consistent framework has developed to address it.

Mental Health, Companion AI, and Minor Safety

Primarily a values collision, whether AI should occupy a clinical or therapeutic role at all, compounded by definitional complexity in disclosure and scope-of-practice standards.

No domain in this paper is more actively contested, or more clearly bipartisan in its underlying urgency, than this one. Nearly one hundred state bills addressing AI chatbots specifically were introduced across thirty-four states in the 2026 legislative sessions alone, and the range of approaches reflects genuine, unresolved disagreement about the right model rather than simple delay.

Nearly
100 AI chatbot bills
introduced across

34 states
during the 2026 legislative sessions



The core question is one of scope of practice: whether, and under what conditions, an AI system may occupy a role historically reserved for a licensed professional. Illinois and Nevada have answered it

by prohibiting AI from providing behavioral health or therapeutic services outright. Utah has taken a materially different approach, permitting the practice under a disclosure-and-guardrails framework rather than a prohibition. Tennessee sits between the two, barring AI systems from representing themselves as licensed mental health professionals and attaching real civil penalties to violations, without a blanket prohibition on AI's involvement in the space. These are not variations on a theme so much as three distinct judgments about where the line between helpful tool and unlicensed practice should sit, made by three legislatures working independently.

Where Should the Line Be Drawn?



Prohibition

Illinois + Nevada

- AI may not provide behavioral-health or therapeutic services.



Guardrails

Utah

- AI involvement is permitted.
- Disclosure and safeguards are required.



Professional-identity restriction

Tennessee

- AI cannot represent itself as a licensed mental-health professional.
- Civil penalties apply.
- AI involvement is not completely prohibited.

Is AI a prohibited provider, a regulated tool, or something in between?

Colorado's own second AI statute belongs in this domain in full. The Chatbot Safety Act, signed May 29, 2026, restricts AI to supplementary and administrative functions within clinical mental health settings and prohibits AI systems from implying a level of licensure or confidentiality protection, including HIPAA-equivalent confidentiality, that they do not actually possess. Considered alongside SB26-189, it reflects a legislature that has continued to legislate carefully on AI well after its highest-profile work group concluded, evidence that the process described earlier in this paper produced durable institutional capacity, not a single successful outcome.

Layered on top of the general scope-of-practice question is a sharper, higher-stakes version of it specific to minors. Companion chatbots, AI systems designed for ongoing, relationship-like interaction rather than a single task, have been described by researchers tracking this legislative area as the newest and fastest-moving frontier of the broader children's online safety debate, itself the subject of nearly three hundred state bills in 2026. Documented cases involving AI chatbots and self-harm

among minors have driven bipartisan legislative urgency in a way few AI-related issues have, including testimony from American Psychological Association researchers documenting the specific risks these systems pose to adolescent users. Age-verification requirements are now in place in roughly half of all states, addressing online platforms generally rather than AI chatbots specifically, but increasingly intersecting with the same underlying technology.

This state-level activity is now running alongside a parallel federal fight, one that shows the same underlying congressional gridlock described earlier in this paper, if not the identical preemption dynamic. The House passed the Kids Internet and Digital Safety Act on June 29, 2026, in a bipartisan 267 to 117 vote, after negotiators narrowed the bill's state-preemption language to preserve stronger state protections rather than override them; the bill's prospects in the Senate remain uncertain, with the original Kids Online Safety Act's own bipartisan sponsors calling the House version dead over a dropped "duty of care" standard. At the same time, five separate federal bills addressing AI chatbots specifically, the CHAT Act, the GUARD Act, the SAFE BOTs Act, the Youth AI Privacy Act, and the CHATBOT Act, have been introduced without any advancing, a congressional gridlock pattern playing out on this specific issue that closely tracks the broader federal preemption stalemate described earlier.

This is a domain in which nearly every stakeholder, consumer advocates, state legislators across the political spectrum, and much of the AI industry itself, agrees that some form of regulation is necessary. The disagreement is almost entirely about what that regulation should require, precisely the condition under which a structured, multi-stakeholder process produces the most value relative to ordinary legislative process. Thirty-four states are currently running that experiment separately and simultaneously, with limited coordination among them.



Frontier AI Safety

Primarily a definitional-complexity problem, what counts as a frontier model and what a transparency obligation should require, with less of a values collision than most domains in this paper, a fact reflected in how closely California and New York's approaches already align.

A distinct question sits apart from how AI is used in any particular decision: whether the organizations building the most capable AI systems have any obligation to disclose what they are building, and to whom, before those systems are deployed. This domain addresses that question directly, and it is newer and narrower than the others in this paper, involving a small number of large developers rather than the broad universe of businesses deploying AI in hiring, lending, or customer service.

Three State Laws Are Becoming a De Facto National Standard

40%

California, New York, and Illinois together represent roughly 40 percent of the domestic AI market.

California

Transparency and incident reporting

- No pre-release approval requirement
- Developers disclose safety practices
- Significant incidents must be reported

New York

Transparency plus threshold-based audit

- Similar “trust but verify” structure
- Requires a single audit after a developer crosses the statutory threshold

Illinois

Annual independent audits

- First state requiring annual third-party safety audits
- Applies to large frontier developers

Three state frameworks now influence companies operating across much of the national AI market.

California's Transparency in Frontier AI Act, SB53, and New York's RAISE Act both adopt a similar structural approach: rather than requiring pre-approval before a frontier model can be released, both laws impose transparency and incident-reporting obligations on developers, a “trust but verify” model under which the largest AI companies, including OpenAI, Anthropic, Google DeepMind, Meta, and Microsoft, must disclose safety practices and report significant incidents after the fact. The degree of substantive convergence between California's and New York's approaches is notable, since most of the domains in this paper show states developing meaningfully different frameworks independently; here, two of the largest state markets have arrived at structurally similar answers.

Illinois has moved from consideration to enactment since this pattern first emerged. Governor Pritzker signed the Artificial Intelligence Safety Measures Act, SB 315, on July 6, 2026, making Illinois the third

state with a frontier safety law and going further than California or New York in one respect: it is the first state to require large frontier developers to undergo annual independent third-party audits of their safety practices, rather than the single audit New York's law requires once a developer crosses the law's size threshold. Illinois's law defines a large frontier developer by an annual revenue threshold of \$500 million, requires disclosure of "catastrophic risk," defined by reference to incidents causing death or serious injury to more than fifty people or more than a million dollars in property damage, and gives exclusive enforcement authority to the state Attorney General rather than creating a private right of action. Lawmakers backing the bill estimated that California, New York, and Illinois together account for roughly forty percent of the domestic AI market, meaning three state laws are now functioning as a de facto national standard in the continued absence of federal action. This domain sits squarely within the federal preemption fight described earlier in this paper: the frontier AI safety obligations imposed by all three states are precisely the category of law the December 2025 executive order and the subsequent congressional preemption attempts have targeted most directly, since these obligations apply to model development rather than to any specific use case exempted from those efforts. That fight gained a concrete new data point on July 23, 2026, when Obernolte and Trahan formally introduced the FRONTIER Act, the frontier-governance piece of their broader discussion draft, as standalone legislation aimed at replacing California, New York, and Illinois's three separate frameworks with one uniform national standard, the clearest sign yet that Washington views these state laws as the target of preemption, not a gap left for states to fill.



AI in Public Safety

Spans two different patterns depending on the application: a governance vacuum in wildfire detection, where adoption has outpaced any formal framework at all, and a genuine values collision in criminal justice, where due process and public safety are in real tension.

Government's use of AI in public safety spans a genuine range, from applications that have generated broad, largely uncontested support to applications generating some of the sharpest civil-liberties concerns in AI policy today. Treating that full range within a single domain is deliberate: no state currently has one consistent framework governing government's use of AI in this space, and the contrast between its least and most contested applications makes clear why that gap matters.

6

states using AI-assisted
wildfire detection

126

cameras operated across
Colorado by Xcel Energy

1,240+

cameras in California's
ALERTCalifornia network

Wildfire detection. AI-assisted wildfire detection has been adopted across Arizona, California, Colorado, Oregon, Washington, and Texas, typically through camera networks that use AI image analysis to identify smoke or fire conditions well before a human observer or traditional detection method would. Xcel Energy operates 126 such cameras across Colorado. California's ALERTCalifornia network includes more than 1,240 cameras statewide. In Arizona, an AI-monitored camera identified what became known as the Diamond Fire before it had grown large enough to be reported through conventional means. Colorado has codified support for this approach through SB23-032, establishing a statutory Wildfire Detection Technology Pilot Program. This application of AI in public safety has generated little organized opposition. In most of the states that have adopted it, it reflects AI deployment proceeding through agency and utility adoption rather than contested legislation, precisely because the underlying question, whether earlier fire detection saves lives and property, does not admit much genuine disagreement.

Criminal justice. The same government function, deploying AI to support public safety, generates substantially more contested outcomes once the application shifts from detecting fire to identifying people. Law enforcement and corrections agencies have adopted AI tools that have produced genuine investigative value: AI-assisted DNA analysis has resolved cold cases previously considered unsolvable, and AI-driven evidence-management systems have benefited public defenders' offices as much as prosecutors' offices in managing case loads that have grown faster than staffing.

Set against those benefits is a documented pattern of harm concentrated in facial recognition technology specifically. Multiple wrongful arrests nationally have been traced to facial recognition misidentification, with the burden of those errors falling disproportionately on Black individuals. Colorado has its own case within this pattern: a woman and her four children were held at gunpoint after an automated license plate reader misidentified her vehicle as stolen, a documented instance of the same underlying failure mode, an automated system's error translating directly into a public safety response with real consequences for people who had done nothing wrong.

A constitutional dimension runs alongside the practical one. Wisconsin's *State v. Loomis* raised, without fully resolving, a confrontation-clause question with implications well beyond that case: whether a criminal defendant has a meaningful ability to challenge a proprietary, legally protected algorithm's role in a sentencing recommendation made about them. The Department of Justice's own December 2024 report and the National Governors Association have both identified facial recognition, predictive policing, and algorithmic risk assessment in bail and sentencing decisions as areas squarely suited to state legislative action, a rare instance of the federal government and a bipartisan state-level membership organization converging on the same assessment of where state authority is both appropriate and currently underused.

One existing governance model offers a useful, if imperfect, precedent for how a state might approach this domain. Sentencing commissions — standing, pluralistic bodies with dedicated research staff, of the kind operating in North Carolina and Minnesota — have earned durability across changes in political leadership precisely because they were built as expert, multi-stakeholder institutions rather than one-time legislative fixes. That durability, developed for a pre-AI purpose, is a meaningful data point for any state considering how to govern AI's growing role in the same institutions those commissions already oversee.

Government's deployment of AI in public safety currently spans from an application that has proceeded with minimal controversy and clear benefit to an application generating wrongful arrests and unresolved constitutional questions, within the same broad category of "AI used by government to keep people safe." No state has yet built a framework that addresses that full range coherently.



*Triggers, not principles,
are where policy actually
gets made or broken.*



PART TWO:

How AI's Existence Is Changing the Way We Live


The domains above address how governments will permit AI to be used within decisions people already understood governments to regulate: hiring, lending, healthcare, public safety. The two domains that follow are different in kind. Neither is primarily a question of permission. Both concern infrastructure and economic conditions that AI has altered fast enough to outpace the institutions built to manage change of this scale.

Data Centers

Primarily a values collision, a community's jobs and economic development weighed against ratepayer costs and environmental impact, with no technical or definitional fix available to resolve it.

Unlike every domain addressed so far, this one is not fundamentally about regulating a technology. It is a land-use, ratepayer, and environmental-review fight, the kind states have managed for decades, that AI has forced onto legislative agendas at a pace those existing institutions were not built to absorb. Framed correctly, the connection to AI policy is real but indirect: the demand driving this fight is AI's growth, while the substance of the fight itself is closer to zoning and utility regulation than to anything else in this paper.

The scale of legislative activity reflects that pace mismatch directly. More than three hundred data center-related bills were introduced across thirty state legislatures in the first six weeks of the 2026 session alone, and twenty-seven states are currently advancing legislation addressing how utilities structure rates for large electricity customers, a category data centers dominate.



New York offers the clearest current illustration of the underlying tension. The state legislature passed the Responsible Data Center Development Act, a one-year moratorium on new data center construction, in the final hours of its session on June 4, 2026, by a Senate vote of 44 to 16 and an Assembly vote of 102 to 39. Rather than act on that bill, Governor Hochul issued Executive Order 62 on July 14, 2026, imposing her own one-year moratorium on new data center permits above fifty megawatts while state agencies study the industry's environmental and infrastructure impact, making New York the first state with a statewide data center permitting moratorium, achieved through executive action.

rather than the legislature’s bill. That bill itself remains undelivered and unresolved; the Assembly Speaker has said he still expects Governor Hochul to sign it, but her decision to act unilaterally first has led most observers to expect an eventual veto. She retains until the end of the year to act on it either way. Maine’s governor faced a comparable bill in April and vetoed it, not because she opposed the underlying concern but because the moratorium, as drafted, would have blocked a five-hundred-fifty-million-dollar project already promising jobs to a former mill town. She simultaneously signed separate legislation stripping data centers of state tax incentives and announced a study council to examine the issue further, a response that reads less as resolution than as an acknowledgment that the legislature and the governor could not agree on how to weigh one community’s jobs against the state’s broader ratepayer and environmental concerns.

Data-Center Growth Creates a Genuine Values Collision Jobs and economic development VS Ratepayer costs and environmental impact	
New York One-year construction moratorium	Maine Moratorium vetoed
● Responsible Data Center Development Act ● Passed June 4, 2026 ● Senate vote: 44–16 ● Assembly vote: 102–39 ● Would temporarily pause new data-center construction	● The proposed moratorium could have blocked a \$550 million project ● The project promised jobs to a former mill town ● The governor separately: ■ Removed state tax incentives for data centers ■ Announced a study council

The pattern is not confined to state legislatures. More than one hundred local moratoriums have already been adopted at the county or municipal level nationally, and roughly a dozen states currently have statewide moratorium bills pending. The fight has begun moving toward the federal level as well; a moratorium bill introduced by Senators Sanders and Ocasio-Cortez is unlikely to advance but signals the scale of concern reaching Congress regardless of its prospects.

No other domain in this paper illustrates as starkly the mismatch between how fast AI’s physical footprint is expanding and how slowly the institutions responsible for zoning, utility rate design, and environmental review are able to respond. A single project’s jobs weighed against a region’s water and electricity costs is precisely the kind of question a five-month legislative session, structured around up-or-down votes on fully drafted bills, is poorly equipped to resolve well.

Workforce Impact

Primarily a governance vacuum. No state has yet built a process to answer the underlying question at all, let alone resolved a disagreement about the answer.

This domain follows a single continuous arc rather than several unrelated concerns: the gap already visible in what recent graduates can do relative to what employers now need, followed by the sharper, more immediate question of what happens as AI begins eliminating positions outright rather than merely changing how they are performed.



Most of what states have done regarding AI in K-12 classrooms has proceeded through agency guidance and local district policy rather than legislation, and that pattern is largely working as intended. Ohio's requirement that every school district adopt a written AI policy by mid-2026 is a useful example of guidance-driven governance producing a workable outcome without a full stakeholder process, a reminder that not every domain in this paper requires the same tool.

The more consequential and less publicly examined gap sits one level up, in the transition from education to employment. Entry-level job postings have declined significantly since 2023, with particularly steep drops in technology and data-related roles, and a substantial share of recent graduates report being underemployed relative to their degrees. No state has yet built a cross-sector process that brings employers directly into a conversation about what AI competency should actually mean across specific career pathways, as opposed to leaving individual institutions and individual employers to define the term independently and inconsistently. That gap,

between what credentialing programs are teaching and what hiring managers are actually screening for, is a policy vacuum in the fullest sense: not a fight between opposing positions, but an absence of any structured effort to close it at all.

The AI Skills Gap Is Becoming a Policy Vacuum

Education and credentialing

Programs teach skills and award qualifications.



Employer expectations change

Employers increasingly screen for AI competency but define it independently.



Entry-level access declines

Entry-level job postings have fallen since 2023, especially in technology and data-related fields.



Graduates face underemployment

Many recent graduates are working below the level expected from their qualifications.



No coordinated process exists

No state has built a cross-sector process connecting employers, education providers, and career pathways.

The same underlying disruption is beginning to generate a second, more acute concern: job elimination rather than job transformation. The federal Great American AI Act discussion draft, described earlier in this paper, includes a provision requiring sixty days' notice before AI-driven mass layoffs, evidence that this concern has reached the federal legislative conversation directly rather than remaining a matter of state-level anxiety alone. This domain sits furthest from anything resembling regulation and closest to what might fairly be called a reckoning, which is precisely why an employer-anchored, structured process is more useful here than either inaction or a purely defensive legislative response.

Connective Threads

The two domains above sort cleanly into the categories this paper has used to organize AI policy: permitting AI's use, and responding to AI's broader effects. The two that follow do not sort as cleanly into either, and are presented separately for that reason.

Government's Own AI Use

The clearest governance vacuum in this paper: not a values collision or a definitional dispute, but the near-total absence of any process addressing the question yet.

Unlike most of the domains addressed so far, this one is not primarily a story of legislative gridlock. It is a story of adoption outpacing governance so quietly that, in most states, no real public debate has occurred at all. A recent survey found that 55.7 percent of government organizations now use AI in some capacity, while only 42.9 percent have adopted any formal policy governing that use, a governance gap of real consequence: the absence of a framework, not a contested disagreement about what the framework should say.

Adoption Is Outpacing Governance

55.7%

of government organizations use AI

42.9%

have a formal policy governing its use

California's Executive Order N-5-26, signed by Governor Newsom in March 2026, establishes new AI vendor certification and procurement standards for state agencies. One structural feature is worth noting specifically: the order builds a mechanism allowing California to route around federal supply-chain-risk designations the state disagrees with, an illustration of state government asserting control over its own AI adoption even as the broader preemption fight over private-sector AI regulation continues elsewhere. A related accountability concern has drawn less legislative attention than it likely deserves: reporting has documented instances of AI-generated government communications losing clear attribution, leaving residents with no reliable way to determine who, or what, is actually responsible for information a government agency has provided them.

The pattern of ad hoc, inconsistent internal policy is not unique to any single state. One Colorado agency's experience is illustrative without being unusual: permission requirements for employee AI use shifted shape repeatedly over a short period, from a blanket prohibition, to a requirement that any AI tool be approved before installation, to an inconsistent, use-case-by-use-case patchwork that varied by division. In one instance, a specific division prohibited the use of AI meeting-transcription tools on video calls under a policy that, when asked, no one in that division could actually produce in

writing. This is not a description of one state being unusually disorganized. It is a description of what happens by default anywhere adoption is allowed to outpace policy, which appears to be most places.

This domain also occupies a distinct position relative to the federal preemption fight described earlier in this paper: a state's own procurement and use of AI is one of the few categories explicitly exempted from that fight altogether. The tension here runs through governors' executive authority and internal agency policy rather than through the state-versus-Washington dynamic driving nearly every other domain in this paper, a structural difference worth naming plainly rather than assuming.

Of every domain examined in this paper, this may be the starkest illustration of its central argument, not because the stakes are highest, but because the absence of process is nearly total. This is not a hard question that states have answered differently. It is a question most states have not yet formally asked.



This is not a hard question that states have answered differently.
It is a question most states have not yet formally asked.

AI and Synthetic Media

A useful counterexample precisely because it fits none of the three patterns strongly, the values at stake are symmetric across parties and the technical definitions involved are comparatively simple, which is a plausible explanation for why ordinary legislative process handled it without difficulty.

This domain sits apart from the rest of this paper for a different reason: the evidence suggests it has largely been resolved through ordinary legislative process, making it a useful point of contrast rather than another example of the pattern this paper otherwise documents.

Thirty states now have laws addressing AI-generated deepfakes in the specific context of elections, and this legislation has drawn bipartisan support in every state where it has passed, according to tracking maintained by Public Citizen. The reason is not difficult to identify: the harm is narrow and

concrete, a fabricated video or audio clip designed to deceive voters about a candidate's actual words or actions, and both major parties have symmetric reasons to want it addressed, since any candidate, of either party, could be the next target. Where the underlying disagreement is this contained, ordinary legislative process appears to function adequately without the kind of structured, multi-stakeholder engagement this paper otherwise advocates for.

One caveat is worth stating plainly: passage has not uniformly meant durability. Some of these laws, including California's, have been struck down or narrowed by courts on First Amendment grounds, a reminder that swift, low-conflict passage does not necessarily produce settled law.

"Swift, low-conflict passage does not necessarily produce settled law."




A related but distinct harm belongs in this domain as well, and connects to a much older fight than AI itself. The unauthorized use of a real person's likeness, voice, or creative style in AI-generated content, whether a synthetic political ad or an AI-generated illustration mimicking a specific working artist's style, raises the same underlying concern as an election deepfake: a person's identity or work used without consent, aimed in one case at a candidate and in the other at a performer or illustrator whose livelihood depends on that style being recognizably theirs. This is not a new fight. It is a continuation of tensions that first surfaced under the Digital Millennium Copyright Act in the platform-and-creator disputes of an earlier technological era, now reshaped by a new technology into a familiar argument about who controls the value of someone's own identity and work.

What These Eight Domains Have in Common

Eight domains, eight different subjects, and one recurring shape. In each, states are attempting to resolve a genuine values collision, a technically dense definitional problem, or a governance vacuum, using a tool, the ordinary legislative session, built for none of the three particularly well. The result, visible across nearly every domain examined above, is not failure exactly. It is fragmentation: fifty states, or some meaningful subset of them, each producing a partial, independently developed answer to a question none of them is actually answering alone, since AI systems, and the companies that build and deploy them, operate across state lines regardless of where any single legislature draws its boundaries. Not all of that fragmentation is a defect. Different states have different values, markets, and institutions, and a well-run process in two different states can still produce two different answers; this paper is not arguing for uniformity, and takes no position on which of those answers is correct. What a structured process changes is the reliability of the reasoning behind each state's answer, not the guarantee that the answers will match.

A structured, neutral process addresses this fragmentation in three specific ways that ordinary legislative process does not.



It creates space to establish reasoning before anyone is asked to take a position. The exchange described at the opening of this paper, a disability-rights advocate and a technology lobbyist arriving at genuine understanding rather than a negotiated truce, was possible only because neither was asked to propose anything until each had explained, in real detail, what they were actually trying to protect. Ordinary legislative process compresses that sequence, or skips it entirely: a bill arrives drafted, and every party's first act is to defend or attack a position rather than to understand one.

It gives the people who actually hold the relevant expertise a seat at the table before the language is written, not a chance to comment on it afterward. Nearly every domain in this paper turns on technical distinctions, what triggers a disclosure obligation, what counts as a frontier model, what constitutes the practice of therapy, that the people best equipped to get right are rarely the same people with a formal vote. A structured process built around convening those people directly produces language

that reflects how the technology and the harm actually work. A process that relies on lobbying and public comment after a bill is already drafted produces language that reflects whoever showed up loudest, soonest.

And it tests whether an agreement actually holds before it becomes law, rather than after. Colorado's AI Policy Working Group required broad, tested support before finalizing any recommendation, using the deliberate distinction between ideas for consideration and formal proposals described earlier to surface disagreement while it was still cheap to resolve. The alternative, discovering that a compromise cannot survive contact with implementation, tends to happen after passage, through amendment fights, litigation, or simple non-enforcement, all of which cost more, in time and public trust, than resolving the same disagreement before a vote.

The evidence for this is not theoretical, though it is still early. SB26-189 passed the Colorado Senate 34 to 1 and the House 57 to 6, and in the months since has not faced the kind of immediate reopening fight that typically follows contested legislation of comparable scope. That is a strong indicator of durability, not yet proof of it: the law does not take effect until January 2027, and its real test is implementation, not passage. That distinction is itself the actual measure of success this paper is arguing for, not simply that a law passes, but that it is built to survive the implementation still to come, because the people who will have to live under it built it together, and tested it against each other's objections, before it ever reached a floor vote.



Every domain examined in this paper presents a state, or several, currently deciding these questions using a tool poorly matched to the problem in front of it. None of them has to. The alternative is not a proposal. It has already been run, in public, on the record, by a state facing exactly this problem eighteen months ago.

A structured process is not the right tool for every policy question, and claiming otherwise would undercut the argument this paper has just made. It is the right tool when a state recognizes its own situation in some combination of the following:

- **A bill has already failed once or twice on an issue that has not gone away.**
- **The parties involved disagree less about principle than about definitions that will determine who a law actually reaches.**
- **Authority over the problem is split across multiple agencies or branches with no single sponsor able to resolve it alone.**
- **Positions are likely to harden the moment a bill is introduced, before anyone has had the chance to understand the other side's actual concern; or, as with several of the domains above.**
- **Adoption of the technology itself is already outpacing any governance process at all.**

The remaining question for any governor, legislator, or agency head who recognizes their own situation in that list is not whether a better path exists. It is whether they are willing to use it before the next session's bill, drafted the old way, adds one more state to the count of contested, unresolved answers this paper has just spent thirty pages documenting.

The Confluence PSG Advantage

Confluence PSG was founded in Colorado. Its principals split their time between Colorado and North Carolina, and the firm works with state and local government leaders nationally. The firm designs and facilitates structured, multi-stakeholder processes for governments and organizations facing decisions too contested, too technical, or too consequential for a standard approach to resolve well. Confluence does not take positions on the policy questions it is retained to help resolve. Its work is the process, not the outcome, and it accepts engagements on that basis: any team member who has publicly advocated for a specific result on an issue does not facilitate a process addressing it. Confluence PSG is led by practitioners who have been in the rooms they now facilitate. The work is not handed off or managed from a distance. Every engagement is designed, led, and delivered against the same standard the firm was built around, and by people who are directly accountable for what it produces.” And then continue with sentence that starts The Colorado AI Policy Working Group described throughout this paper was one engagement built on that discipline. It was not the only one.



Its work is the process, not the outcome.

Recent Projects & Team Experience

Confluence’s practice spans civil rights, education, criminal justice, transportation, public safety, housing, healthcare, and workforce and economic development, across state, regional, and local government. A selection of engagements relevant to the questions this paper addresses:

Colorado AI Policy Working Group. Described in full above. Nine sessions, fourteen members, unanimous support, SB26-189 signed into law thirteen days after introduction.

RTD Accountability Committee. A committee of fourteen voting members plus one non-voting member, created by SB25-161 to resolve years of gridlock over transit governance in one of only three fully elected transit board structures in the country. On the central governance question, a proposed nine-member board structure fell two votes short of the required supermajority in its first vote, 6 to 5, with three members absent. The following week, after further deliberation grounded in shared evidence rather than restated positions, the same proposal passed 11 to 2. The committee’s thirty-

one recommendations, including the restructured board, were codified into law in SB26-150, and the legislature’s own declaration credited the committee’s process and its neutral facilitator directly.

Governor-Elect Polis Transition. At the invitation of Governor-Elect Polis and his Chief of Staff, Confluence President Berrick Abramson managed the full scope of Boldly Forward, Colorado’s official gubernatorial transition committee, including transition operations, the hiring process for twenty state agency directors, and development of reforms and opportunities reports for each executive agency. The transition was built to be unusually open: committee membership combined appointed civic leaders with members chosen through a public application process, and town halls and an online portal gave residents across the state a direct channel into the process before the new administration took office.

Governor and Legislative ARPA Listening Tour. When \$3.8 billion in federal American Rescue Plan Act funds became available to Colorado in 2021, Governor Polis, Treasurer Young, and legislative leadership from both parties needed to understand what residents across the state’s diverse communities and regions actually needed before priorities were set. Abramson designed and led a statewide listening tour across seven geographic regions and ten additional audience-specific sessions. The input gathered directly shaped the state’s prioritization of the funds, and Abramson went on to facilitate legislative discussions of the findings, support discussions between the legislature and the executive branch, and contribute to drafting the state’s first report to the U.S. Treasury.

Task Force on the Rights of Coloradans with Disabilities. Created by the Colorado legislature to examine government accessibility, housing, outdoor access, and a full rewrite of the state’s disability civil rights statutes. Confluence facilitated the full task force and its four simultaneous subcommittees, integrating testimony from dozens of stakeholders with direct lived experience. The task force published fifty-eight recommendations in its January 2025 final report.

Colorado Educator Safety Task Force. A seventeen-member task force examining aggressive and violent behavior toward school staff, supported by four subcommittees. Confluence designed and administered a survey of more than 1,100 educators to generate data the state’s existing records had missed; more than half reported having been physically injured by a student, a finding that reframed what state action the task force ultimately recommended.

1,100+

educators surveyed

More than half

reported having been physically injured by a student

Working Group on Transforming Criminal and Juvenile Justice. Convened by executive order after legislation to reauthorize Colorado’s twenty-year criminal justice commission was postponed indefinitely, sunseting the body with no replacement. Confluence facilitated the seventeen-member group, which included prosecutors, defense attorneys, and reform advocates who had been in direct opposition for years, bringing in more than fifty outside voices before the group turned to its own recommendations. The group reached consensus on the mission, structure, and membership of two new permanent commissions, a genuine outcome among historic adversaries after a two-decade stalemate. The legislature has not yet acted to implement those recommendations, a separate question that remains outside the process’s control.

Wildfire Resiliency Code Board. Created by SB23-166 to develop Colorado’s first model code governing construction in the wildland-urban interface. Confluence supported the twenty-one-member board through code development and a series of meetings held across the state, ensuring communities facing different wildfire conditions had a voice in the resulting code.

CBI Forensic Services Committee. Convened by the Governor after the discovery of errors in forensic evidence handling at the Colorado Bureau of Investigation. Confluence facilitated a committee spanning forensic scientists, prosecutors, defense attorneys, and victim advocates to recommend process improvements and oversight structures for the agency’s forensic labs.

Colorado LTSS Work Group. An eighteen-month, multi-stakeholder process examining Colorado’s Long-Term Services and Supports system, convened by the Colorado Cross-Disability Coalition and funded by the Colorado Health Foundation. A steering committee of fifteen to twenty members, supported by up to four expert subcommittees, is developing recommendations for the Governor, the Joint Budget Committee, and legislative health committees, with final recommendations due in September 2027.

Clean Transit Enterprise, Strategic Planning and Stakeholder Engagement. Confluence designed the stakeholder engagement structure for CTE’s formula grant rollout following SB25-230, then was retained in 2026 to lead development of the enterprise’s first five-year strategic plan. A full-day board retreat surfaced the central strategic tension facing the enterprise, whether it should operate primarily as a funding partner or a more accountability-focused entity, a question the retreat named and clarified rather than resolved for the board.

Puerto Rico Career Pathways. After Hurricanes Irma and Maria devastated Puerto Rico’s school system, the Governor and Secretary of Education faced a question beyond rebuilding facilities: how to redesign career pathways to produce real economic opportunity for students. Abramson designed

and led a multi-phase initiative establishing first-of-their-kind partnerships between schools and employers. In the first year, more than 1,400 students were matched with employers for work-based learning, with more than 800 guaranteed a job at a living wage upon completion, work recognized by the Clinton Global Initiative with an Ideas Into Action award.

City of Manitou Springs, Budget Community Engagement. Facing a budget shortfall of \$1.3 to \$1.6 million, the city retained Confluence to give residents a genuine opportunity to weigh in before the City Council acted. A community survey drew more than 500 responses, and two in-person sessions moved participants from shared understanding of the budget's actual constraints to specific, consistent input: protect public safety, treat parks and recreation as central to quality of life, and prefer sales tax, tourism fees, or parking revenue over property tax increases. The conclusion that the city could not cut its way out of the shortfall came from the data the city presented, not from Confluence, and participants across both sessions arrived at it independently.

City of Longmont, City Council Strategic Planning Retreat. A city manager who had previously participated in a Confluence-facilitated multi-stakeholder process brought the same methodology to his own council, seeking a facilitator specifically so he could participate in the retreat rather than run it. A pre-retreat survey asked council members not about policy positions but about what they hoped to achieve and what they had learned from their constituents, surfacing common ground before the group ever sat down together. The retreat produced a substantive shift on a long-standing, low-grade point of council disagreement, the future of the municipal airport, moving the conversation from stalemate to a shared, specific vision.

The sequence, not the subject matter, is what determines whether a room moves.

The methodology that produced unanimous agreement on AI regulation among parties in two years of direct conflict is the same methodology that moved a city council past a decade of stalled conversation about an airport. The sequence, not the subject matter, is what determines whether a room moves.

Author



Berrick Abramson is President and Co-Founder of Confluence PSG. Over more than twenty-five years, he has served as process architect and executive facilitator for some of the most consequential collaborative governance processes in state government, including the engagements described throughout this paper. His practice spans civil rights, education, criminal justice, transportation, public safety, housing, healthcare, and workforce development.

Before founding Confluence, he worked in institutional sales and trading and municipal underwriting, founded a venture consulting firm involved in more than a billion dollars in transactions, and led policy and government affairs for a four-hundred-person national nonprofit. He has served on the Colorado Commission on Higher Education since 2019, appointed by Governor Polis and reappointed in 2023.



CONFLUENCE
POLICY & STRATEGY
GROUP