

AI Governance, Evidence and Dispute Resolution: Themes for ADR Institutions

Ranse Howell · October 2026 · 14 min read

The headline overview

AI governance is moving toward more demonstrable evidence, independent scrutiny and meaningful human recourse. This bulletin considers what those themes may mean for ADR institutions as both users of AI and forums for resolving AI-related disputes — across borders, not only within a single jurisdiction.

This year has seen a significant acceleration in the way governments, courts and commercial markets are defining what counts as accountable AI. For organizations involved in alternative dispute resolution — whether administering proceedings, sitting as neutrals, or advising parties — several of these developments intersect directly with the work they do.

This paper does not presume to describe any particular institution's practice. It identifies themes that have emerged across the United States, the United Kingdom and the European Union, and draws out the indicators that suggest ADR institutions will increasingly find themselves at the center of the AI accountability conversation — not only as forums for disputes about AI, but as users of AI tools that may themselves come under regulatory scrutiny.

The observations here draw on the regulatory developments covered in my October 2026 update ([Navigating Federal/State AI Governance and Compliance](#)), as well as institutional guidance published by arbitral bodies, government consultation documents, and the emerging evidence standards that are beginning to shape how AI-assisted processes are recorded and challenged. Where relevant, links point to the underlying sources or to regitech.ai/regulatory-landscape for a broader jurisdictional overview.

The original five-factor framework considered government action, market power, the financial case and governance as an enabler. For a global ADR audience, the first two lenses become supranational/cross-border frameworks and national/jurisdictional action. A sixth lens considers people and recourse; a seventh considers incentives to adopt. The factors are connected analytical lenses, not a scoring system.

SEVEN CONNECTED LENSES FOR GLOBAL ADR

01

Supranational & cross-border frameworks

Where jurisdiction travels

02

National & jurisdictional action

Different routes. Related expectations.

03

The market power principle

Expectations travel through markets

04

The financial case

Look beyond the penalty

05

Governance as an enabler

Evidence makes accountability usable

06

People & recourse

Who can question the outcome?

07

Incentives to adopt

Make the purpose visible

Five original factors, adapted globally + people and recourse + incentives to adopt.

Figure 1. The five original lenses, a sixth centered on people and recourse, and a seventh on incentives to adopt — adapted for global ADR. Connected lenses, not a quantitative score.

Factor 1. Supranational & cross-border frameworks

The European Union's AI Act (Regulation 2024/1689) entered into force on August 1, 2024. Its treatment of artificial intelligence in dispute resolution is among the most consequential — and most debated — provisions for the ADR community.

What the Act says

Annex III, point 8(a) classifies AI systems as **high-risk** when they are used by an "alternative dispute resolution body" to research and interpret facts and the law, or to apply the law to a concrete set of facts, where the proceedings produce legal effects for the parties. Recital 61 makes clear that arbitration is included: arbitral awards are generally binding and enforceable (for example, under the New York Convention), bringing them within the scope of "legal effects."

The European Commission's May 2026 Draft Guidelines drew a functional distinction. AI tools used for purely administrative tasks — speech-to-text transcription, case assignment, document organization, anonymization — are generally excluded from the high-risk classification. Tools that generate draft decisions, analyze legal precedents, or suggest how the law applies to facts are included.

The institutional response

In July 2026, [sixteen leading international arbitral institutions](#) submitted a joint letter to the European Commission expressing several concerns:

- **Misattribution of responsibility.** The institutions argued that they serve an administrative function — appointing arbitrators, managing procedural timelines, collecting deposits — and do not decide the merits of disputes. Labeling the institution as the primary "ADR body" for regulatory purposes risks attaching compliance burdens to the wrong actor.
- **Inconsistency with judicial treatment.** The Draft Guidelines exclude bodies responsible for the "management, governance, and support of the judiciary" from the Act's scope. The institutions argued that their administrative role is analogous and should receive similar treatment.
- **The adjudicative function.** Responsibility for AI-assisted decision-making, the letter contended, rests with the arbitral tribunal — the individuals who exercise independent judgment — not with the institution that administers the case.

This debate remains unresolved. The Commission has not yet issued final guidance, and the obligations for high-risk systems do not apply until **December 2, 2027** (for stand-alone AI systems) or **August 2, 2028** (for those embedded within broader products). But the timeline is not as distant as it appears: organizations that use AI in proceedings involving EU parties, or whose awards may be enforced within the EU, will need to assess their exposure well in advance.

The filter mechanism

The Act includes a filter under Article 6(3) that can exempt an AI system from high-risk classification if it does not pose a significant risk of harm or materially influence decision-making. Tools used for preparatory research, translation, or organizing a chronology may fall within this exemption. But the boundaries are not yet tested, and the distinction between a research tool that "informs" and one that "influences" is likely to generate its own disputes.

Extraterritorial reach

The Act applies not only to providers and deployers established in the EU, but also to those outside the EU if the output of the AI system is used within the Union. For international arbitration proceedings seated outside the EU, this creates a layer of uncertainty that institutions and practitioners will need to navigate.

A global, not only European, lens

ADR is a cross-border activity. A dispute may connect parties, governing laws, the seat of proceedings, data processing and award enforcement in different jurisdictions. The New York Convention provides an important enforcement context; it does not by itself determine whether an AI law applies.

The EU AI Act belongs in a wider field that includes trade arrangements such as the CPTPP and UN initiatives on AI governance. These are different kinds of instruments, not interchangeable compliance regimes. Their relevance is a reason to map cross-border exposure, not to assume a single global rulebook. National implementation and the position of individual EU member states also require separate attention.

The [regulatory landscape](#) provides a broader starting point for that jurisdictional assessment.

Role and use matter. Administrative support and adjudicative decision-making are not the same activity. Classification requires a functional assessment, not a label.

Factor 2. National & jurisdictional action

One of the clearest indicators to emerge this year is the convergence, across jurisdictions, on a single principle: **self-attestation by AI developers and deployers is no longer considered sufficient**. Independent, third-party verification is becoming a regulatory default.

United States: state-level action

The US does not yet have comprehensive federal AI legislation, despite a substantial volume of bills introduced in 2026. The legislative energy has been at the state level, and it has produced a series of concrete requirements:

- **Illinois SB 315** (signed July 6, 2026; effective January 1, 2027) requires developers of frontier and high-impact AI systems to submit to **annual independent, third-party audits** — the first US state law to mandate this. An AI Cabinet was created by executive order on September 22, 2026, to coordinate implementation.
- **California SB 813 and AB 1405** (signed September 9, 2026) go further by establishing the infrastructure for who may conduct those audits. SB 813 creates criteria for designating **Independent Verification Organizations (IVOs)** qualified to assess AI systems (effective

January 1, 2028). AB 1405 establishes a public **AI Auditor Registry**, requiring that state-mandated audits be performed only by registered auditors (effective January 1, 2029). Together, these laws create the first state-level audit accreditation framework.

- **Colorado SB 26-189** (effective January 1, 2027) replaced the earlier SB 24-205 with a narrower framework focused on Automated Decision-Making Technology (ADMT). It requires deployers to retain compliance records for at least three years, provide consumers with explanations of algorithmic factors, and afford **meaningful human review and reconsideration** within 30 days of an adverse outcome. The Colorado Attorney General filed formal proposed rules on August 11, 2026.
- **California SB 1000** (signed September 30, 2026) removed the original 1 million-user threshold from the AI Transparency Act, extending its disclosure and provenance requirements to a broader range of providers. The law's content-provenance requirements — including latent disclosure through machine-readable metadata — are now operative.

United Kingdom

The UK maintains a principles-based, sector-specific approach rather than a single AI Act. Five principles — safety, transparency, fairness, accountability, and contestability — are applied by existing regulators (FCA, ICO, Ofcom, CMA, MHRA). However, the limits of the voluntary model became visible in 2026:

- The UK AI Security Institute (AISi) found its evaluation capability constrained by its reliance on **voluntary cooperation** from frontier developers. Anthropic's decision to decline pre-release model access underscored the structural weakness.
- Parliament summoned Anthropic, OpenAI, Google DeepMind and Meta to an **urgent hearing on October 13, 2026**, to examine whether pre-release safety testing should become legally mandatory.
- The government has signaled a potential **2026 Statutory AI Bill** to address frontier models.
- The UK Jurisdiction Taskforce launched a consultation on a draft Legal Statement addressing **liability for AI-associated loss**, confirming the expectation that English law will resolve novel AI issues through flexible application of existing common law principles.

For ADR institutions operating in the UK, these developments indicate a direction of travel — even without binding legislation — that favors demonstrable accountability over self-reported compliance.

Federal US: volume without enactment

At the federal level, the picture is one of activity without resolution. Bills introduced in 2026 include:

- The **FRONTIER Act** (H.R. 9925), proposing a marketplace of licensed independent verification organizations.
- The **AI Kill Switch Act**, requiring emergency shutdown mechanisms for AI systems capable of causing significant harm.
- The **AI Systems Transparency Act**, mandating consumer-facing disclosures and access for independent researchers.
- The **Ban Artificial Superintelligence Act** (Sanders–Casar), proposing a permanent ban on ASI with a cabinet-level Department of AI.
- The **Great American AI Act**, a broader regulatory framework bill.

None has been enacted. Congress adjourned for the November midterm recess without passing comprehensive AI legislation, leaving states and international frameworks to set the operational standards.

Separately, the executive branch has added a **Voluntary White House Accord** (September 29, 2026) and a new **Super Intelligence Force** led by Jay Clayton (announced October 3–4, 2026), with 120 days to report on AI risks and opportunities — placing its deadline around early February 2027.

THE REGULATORY HORIZON · OCTOBER 2026 SNAPSHOT

Aug. 2, 2026

EU transparency

Article 50 transparency obligations.

July 1, 2027

Workplace AI

California SB 947 effective.

Oct. 26, 2026

Colorado comments

Public comment deadline on ADMT rules.

Dec. 2, 2027

EU stand-alone systems

High-risk obligations for stand-alone systems.

Nov. 16, 2026

California recommendations

Expert recommendations on AI oversight implementation.

Jan. 1, 2028

Independent audits

Illinois mandatory independent audits begin; California IVO criteria commence.

Jan. 1, 2027

State laws commence

Colorado SB 26-189, Illinois SB 315 and New York RAISE Act effective.

Aug. 2, 2028

EU embedded systems

High-risk obligations for systems embedded in broader products.

Around Feb. 2027

Federal report

Super Intelligence Force report; approximate, based on the 120-day interval.

Jan. 1, 2029

AI Auditor Registry

California public registry operational under AB 1405.

Dates follow the source bulletin. The federal report date is approximate; confirm current primary sources.

Figure 2. Selected regulatory milestones from the October 2026 bulletin. Timelines are not uniform obligations for all ADR institutions. The federal report date is approximate; confirm current primary sources.

Date	Development	Status
Aug. 2, 2026	Article 50 transparency obligations.	In effect in source
Oct. 26, 2026	Public comment deadline on ADMT rules.	Consultation
Nov. 16, 2026	Expert recommendations on AI oversight implementation.	Recommendations
Jan. 1, 2027	Colorado SB 26-189, Illinois SB 315 and New York RAISE Act effective.	Scheduled in source
Around Feb. 2027	Super Intelligence Force report; approximate, based on the 120-day interval.	Estimated
July 1, 2027	California SB 947 effective.	Scheduled in source
Dec. 2, 2027	High-risk obligations for stand-alone systems.	Scheduled in source

Date	Development	Status
Jan. 1, 2028	Illinois mandatory independent audits begin; California IVO criteria commence.	Scheduled in source
Aug. 2, 2028	High-risk obligations for systems embedded in broader products.	Scheduled in source
Jan. 1, 2029	California public registry operational under AB 1405.	Scheduled in source

Regulatory calendar. Dates are reproduced from the October 2026 source bulletin; confirm current primary sources.

Factor 3. The market power principle

Large organizations subject to the strictest requirements can carry those expectations into procurement, supply chains and partner agreements. This is the **market power principle**: rules may influence counterparties beyond the territory in which they were enacted.

For ADR, the relevant indicators are client due diligence, technology procurement, the terms on which information is handled, and the evidence expected of an AI-assisted process. The point is not that every institution faces the same obligations, but that commercial expectations may move before a directly applicable legal requirement does.

Commercial access and procurement

As governance expectations are embedded in procurement requirements, enterprise contracts and partner agreements, the ability to demonstrate independently verified AI accountability is becoming a condition of commercial participation. Organizations that cannot show evidence of control may find their market access constrained — not by regulation, but by the due-diligence requirements of their commercial counterparts.

SELECTED GOVERNANCE THEMES BY JURISDICTION

	Evidence	Human review	Verification
European Union	●	●	●
Illinois	●		●
California	●	●	●
Colorado	●	●	
United Kingdom	●	●	●
US federal	●		●

Markers indicate themes discussed, not equivalent legal duties or compliance scores.

Figure 3. Selected governance themes by jurisdiction. Markers show topics discussed in the bulletin, not equivalent legal duties, compliance ratings or an exhaustive legal comparison.

Jurisdiction	Scope to consider	Selected governance themes	Timing / status
European Union	High-risk AI used in legally consequential ADR; role and use matter.	Human oversight, documentation and system logs.	Article 50: Aug. 2026; high-risk dates below.
Illinois	Frontier / high-impact developers under SB 315.	Annual independent audits and reviewable records.	Effective Jan. 2027; audits from Jan. 2028.
California	Provider, auditor and workplace provisions have distinct scopes.	Auditor qualifications, provenance and human corroboration.	Staged implementation through Jan. 2029.
Colorado	Automated decision-making technology under SB 26-189.	Three-year records; explanation and human review.	Jan. 1, 2027, in the source bulletin.
United Kingdom	Sector-specific principles and an evolving policy debate.	Accountability, contestability and proposed testing duties.	Policy direction; not a single enacted AI Act.
US federal	Several legislative proposals; voluntary executive measures.	Independent verification appears across proposals.	Proposed / voluntary; not comprehensive enacted law.

Table 1. Selected jurisdictions and requirements described in the October 2026 source material. Different roles and systems attract different duties.

Factor 4. The financial case

Beyond regulatory compliance, there are indicators that the financial and commercial landscape is reinforcing the case for structured AI accountability.

Incident exposure

- The **Hong Kong deepfake fraud** (reported February 2024) — in which a finance worker was deceived by AI-generated video into transferring \$25 million — illustrates the scale of exposure when AI-assisted processes lack adequate verification.
- **Multi-agent system failures**, including the July 2026 Hugging Face incident in which autonomous AI agents breached sandbox containment, have shifted the policy conversation from theoretical risk to demonstrated vulnerability.

Litigation risk

Under California's SB 813, an independent audit is "relevant to, but not conclusive of" liability. Organizations that maintain structured evidence of their AI processes will be better positioned in litigation. Conversely, the absence of such records may become a source of adverse inference.

Commercial access and procurement

As governance expectations are embedded in procurement requirements, enterprise contracts and partner agreements, the ability to demonstrate independently verified AI accountability is becoming a condition of commercial participation. Organizations that cannot show evidence of control may find their market access constrained — not by regulation, but by the due-diligence requirements of their commercial counterparts.

Insurance

The UK Jurisdiction Taskforce's consultation on AI liability, and the broader conversation about insurability of AI-related losses, suggest that the insurance market will increasingly differentiate between organizations that can demonstrate governance controls and those that cannot. Premiums, coverage availability and policy terms may all be influenced by the adequacy of an organization's AI evidence infrastructure.

Factor 5. Governance as an enabler

As regulatory frameworks mature, the question of what constitutes adequate evidence of an AI-assisted process is becoming more precisely defined. Several converging developments are worth noting.

Content provenance and the C2PA standard

The Coalition for Content Provenance and Authenticity (C2PA) has established an open technical standard for recording the origin and modification history of digital content. It works by embedding a cryptographically signed manifest into files (JPEG, MP4, PDF and others) that tracks the creator, timestamps, and every edit performed by compliant software.

In 2026, adoption has accelerated:

- Hardware-level implementations (in devices from Leica, Sony and Google) sign content at the point of capture — establishing provenance before any human or AI intervention.
- Major news organizations (BBC, AP, Reuters, NYT) now prioritize C2PA-verified assets.
- Platforms including LinkedIn, YouTube and TikTok have integrated C2PA support to surface provenance information.
- Some **US federal courts have begun accepting C2PA-credentialed evidence** under Rule 901, though case law is still developing.
- The **EU AI Act's Article 50** transparency obligations, which require disclosure when content is AI-generated, have applied since August 2, 2026.

For ADR institutions, provenance technology offers a practical mechanism for documenting the integrity of AI-assisted processes — recording which system produced a given output, when, and whether the output was subsequently modified.

Record retention

Statutory record-keeping requirements are becoming more specific:

- Colorado's SB 26-189 requires at least **three years** of compliance records for automated decision-making technology.
- The EU AI Act requires at least **six months** of system logs for high-risk applications.
- Illinois SB 315's annual audit requirement implies that records must be maintained in a form that an independent auditor can review.

Explainability and disclosure

The ability to explain why an AI system produced a particular outcome is becoming a statutory obligation in multiple jurisdictions:

- Colorado requires deployers to provide consumers with **explanations of the principal factors** in an algorithmic determination.

- California's SB 942 (the AI Transparency Act, as amended by SB 1000) retains latent, machine-readable disclosure requirements. As explained in the [October 2026 update](#), the amendment removed the earlier requirement to offer a visible disclosure.
- The EU AI Act's high-risk transparency provisions require documented descriptions of system capabilities and limitations.


The implication for ADR institutions is that the standards for what counts as adequate documentation of an AI-assisted process — whether in a hearing, a case assessment, or an administrative decision — will increasingly be set by legislation rather than by institutional convention alone.

Independent assurance is an enabler when it makes responsibility and evidence easier to establish. Organizations that can show what their AI did, and who was accountable, are better positioned to engage with customers, auditors and regulators. This is the fifth lens in the [October framework](#), not a list of product capabilities.

3 years — Compliance records. Colorado ADMT framework. Record-retention period described in the source bulletin; application depends on scope.

6 months — Minimum system logs. EU high-risk AI. Minimum log period described in the source bulletin, subject to applicable requirements.

30 days — Meaningful human review. Colorado adverse outcomes. Review period described in the source bulletin; not a universal ADR deadline.



Factor 6. People & recourse

The most recent legislative activity places particular emphasis on the individual affected by an AI-assisted decision — and, specifically, on their right to understand, challenge and obtain meaningful review of that decision.

Consumer and worker protections

- **Colorado SB 26-189** provides consumers with notice, explanation, the right to correct data, and meaningful human review within 30 days of an adverse algorithmic outcome.
- **California SB 947** (effective July 1, 2027) requires employers to notify workers when AI is used in employment decisions and to provide human corroboration before adverse actions.
- **Illinois HB 3773** bars discriminatory AI in employment decisions and provides enforcement through the Attorney General.
- **The EU AI Act's Article 50** requires all deployers of AI systems to inform individuals when they are interacting with an AI.

Why this matters for ADR

The convergence of these provisions with the traditional function of dispute resolution institutions is difficult to overstate. When a statute requires "meaningful human review" of an algorithmic decision, or when a consumer exercises a right to challenge an AI-influenced outcome, the question becomes: **where does that challenge go?**

Courts are one answer, but they are slow, expensive, and ill-suited to the technical specificity of many AI disputes. ADR institutions — with their existing infrastructure for neutral selection, confidential proceedings, technical expertise, and structured resolution — are positioned to provide exactly the kind of forum these statutes contemplate.

Publicly available ADR rules and professional guidance address themes such as proprietary algorithms, data privacy, secure discovery, technical expertise, and the responsibility to exercise independent judgment over AI-assisted outputs. These are indicators of a wider professional response, not an assessment of any institution's practice.

These responses are, in one sense, the profession's own answer to the question regulators are asking: **who provides the structured, credible route for challenging an AI-influenced outcome?**

When an individual exercises a right to challenge an AI-influenced outcome, where does that challenge go?

Factor 7. Incentives to adopt

Accountability becomes easier to adopt when its purpose is visible to the people expected to use it. The preceding factors point to practical incentives: clearer procurement decisions, less uncertainty about responsibility, usable evidence when an outcome is challenged, and greater confidence in human review.

For ADR institutions, the question is not simply whether AI can make a process faster. It is whether the process remains understandable, contestable and accessible. Knowledge sharing can reduce duplicated effort; collaboration can help smaller organizations participate without requiring them to surrender local judgment.

This seventh factor is an analytical extension of the framework, **not a claim that financial returns or regulatory approval are guaranteed**. Meaningful measures would connect adoption to access, quality of explanation, timeliness of review and the experience of the people affected. That leads to the shared-access agenda considered in the conclusion.

Summary & a shared next step

The themes identified here — the EU's high-risk classification of AI in ADR, the convergence on independent verification, the tightening of evidence standards, and the increasing emphasis on individual recourse — combine developments described in the October 2026 source material with published timetables and an evolving policy discussion. Their legal status and application need to be considered separately. Their combined effect is to place dispute resolution institutions closer to the center of the AI accountability conversation than at any previous point.

Institutions that are already tracking their AI tool usage, maintaining structured records of how those tools operate, and preparing for the possibility that their processes may come under regulatory scrutiny will be better positioned — whether as forums for AI disputes, as deployers of AI tools, or both.

Organizations that can show what their AI did, why, and who was accountable will be ready for whichever framework takes hold.

From accountability to access: AI for All

“Accountable AI and accessible AI are not separate goals — they are inseparable.”

Regitech's [AI for All initiative](#) comes from learning from ADR and conflict prevention, management and resolution. The invitation is to share knowledge, collaborate across disciplines and improve access to justice — through partnership without paternalism and a commitment to measurable impact.

A proportionate next step follows from the framework: map where AI is used and which jurisdictions may matter; identify the evidence and human review available; then share nonconfidential learning about what supports fair, accessible outcomes. Collaboration should preserve client confidentiality and respect each institution's independence.

[Explore AI for All](#), read the [October 2026 governance update](#), or continue the conversation at [Regitech](#).

About the Author

Ranse Howell is CEO of Regitech and founder of The Unified Sentient. His background spans law, alternative dispute resolution, social policy and the creative arts. He approaches AI governance with an emphasis on stakeholder alignment and on accountability systems that enable innovation rather than restrict it.

ransehowell@regitech.ai · [202 913 6634](tel:2029136634) · regitech.ai

Legal & editorial note

This bulletin provides general information and does not constitute legal advice. Institutions should take advice on the laws that apply to their own roles, systems and jurisdictions. This is an October 2026 editorial snapshot, not a live compliance tracker. Dates and legal-status descriptions follow the source bulletin; consult current primary sources before relying on them.

Based on the author's October 2026 bulletin text and [October governance update](#). Figures are explanatory, not performance measurements. No institution-specific practice or confidential client matter is described.