

REPORT

The State of AI in U.S. Arbitration 2026




Table of Contents

Foreword	3
Executive Summary	5
Methodology	7
Featured Voices	9
The Trust Spectrum in Legal AI	10
Concerns About AI: Accuracy, Confidentiality, and the Limits of Trust	15
The Judgement Line	19
Measuring the ROI of Legal AI	23
Workforce and Career Implications	26
The Institutional Access Divide	30
AI Adoption by Organizational Profile	32
The Maturity Model	35
Accelerating AI Adoption	38
Looking Ahead	40
Conclusion	42

Foreword

Alternative dispute resolution stands at an inflection point. Over decades, the practice of arbitration has adapted to successive waves of technological advances. However, the emergence of artificial intelligence (AI) and large language models (LLMs) heralds a transformation that may prove orders of magnitude more profound.

Unlike passive digital storage or video transmission tools, AI platforms actively process, analyze, synthesize, and generate legal and factual content. LLMs can extract key dates from large volumes of unstructured documents, draft complex chronologies, perform in-depth legal research, and summarize party submissions in far less time than it would take humans to do the same work.

The speed at which AI works speaks directly to time and cost efficiency in arbitral proceedings. Generative AI offers a way to significantly reduce document review costs, streamline tribunal preparation, and enhance procedural velocity overall.

Yet, the integration of AI into dispute resolution raises fundamental questions regarding procedural integrity, data privacy, confidentiality, evidentiary reliability, and the nature of the adjudicative function. Arbitration relies on party autonomy, procedural fairness, and the tribunal's judgment and decision-making. The deployment of AI systems into this ecosystem may create, depending on how it is used, tension between efficiency and norms that emerged in a pre-AI world.

Arbitral institutions, including the American Arbitration Association® (AAA®), are already formulating practice notes, procedural orders, and guidelines to address AI disclosure and usage. However, formal regulation by legislatures, regulators, and court systems is still developing, leaving practitioners and institutions to operate in an environment of uncertainty. Individual law firms have adopted divergent internal policies ranging from outright bans on commercial LLMs to active, enterprise-wide deployments. Meanwhile, independent arbitrators often lack corporate IT infrastructure, creating a stark digital divide as compared to those arbitrators associated with large firms.

To establish a rigorous, empirical foundation for this evolving landscape, the AAA and Jus Mundi initiated a comprehensive joint study. This report presents our findings, evaluating how AI is used, perceived, and regulated across the dispute resolution community in the U.S.

Thank you to every respondent who took the time to complete this survey, and to those who went further and sat down with us for an in-depth interview. Your participation shaped this report.



Annie Lespérance
Head of Americas,
Jus Mundi




Linda Beyea
Group Vice President,
Innovation and
Thought Leadership,
American Arbitration
Association



Greg Kochansky
Vice President,
Thought Leadership
and Applied AI,
American Arbitration
Association



Jean-Rémi de Maistre
Co-Founder & CEO,
Jus Mundi



Executive Summary

AI has permeated so many facets of the arbitral process so quickly that the field has not yet adopted consistent approaches on how to use it. Tools that draft, research, and analyze in minutes are already in practitioners' hands, but there is still room to debate questions that matter: where AI belongs, how far to trust it, and who is at risk of being left behind.

Derived from a survey of 557 respondents, as well as qualitative interviews with leading arbitrators and law firm partners, this report presents a baseline for AI adoption and governance in U.S. arbitration.

Six findings stand out:


1. **The Usage-Trust Feedback Loop:** Trust is workflow-specific and driven by exposure to tools rather than tool architecture. Daily users report an average trust score of 3.42 out of 5.00, compared to 0.83 among non-users¹.
2. **The Judgment Line:** Most respondents draw a clear boundary between task execution and arbitral decision-making. While 51% of respondents deploy AI for document review and 46% for legal research, 58% of respondents believe award-drafting should not be “primarily handled by AI” and 49% take the same view on evidentiary hearings. The findings do not, however, differentiate based on matter complexity or amount in controversy, yielding an avenue for future research.
3. **Shifting Human Capital:** Over 50% of respondents expect AI to absorb document review and cite-checking work, freeing humans to focus more on strategic judgment, ethical oversight, and oral advocacy. Concerns about AI displacing human practitioners remains low (1.70).
4. **The Infrastructure & Access Divide:** Adoption is influenced by access to sufficient resources, organizational support, and infrastructure, or lack thereof. While 83% of respondents at global law firms operate under clear AI policies, only 15% of independent arbitrators do. Independent arbitrators are three times as likely to cite cost and four times as likely to cite training as primary adoption barriers.

1. Non-percentage metrics in this report reflect Likert-scale scores, all of which are out of a possible 5.00 total.

Independent arbitrators face steep barriers to adoption: less training, less support, and fewer resources. It may be that boutiques leverage operational agility to deploy domain-specific tools faster, without corporate procurement friction.

5. **Boutique Agility vs. Large Firm Scale:** Boutique law firms lead the field in realized efficiency, with 66% reporting significant time savings compared to 50% at global law firms. It may be that boutiques leverage operational agility to deploy domain-specific tools faster, without corporate procurement friction.
6. **Industry Priorities:** The dispute resolution community prioritizes authoritative guidelines (30%) and domain-specific arbitration tools (24%) over cost reduction (5%) or commercial ROI metrics (3%).

Together these findings describe a field at an early and uneven moment: confident that AI will not supplant human judgment but divided in terms of their familiarity and comfort level with existing AI tools. In this report, we will explore where the field stands today, and what it will take to move forward.




Methodology

To set a baseline measure for integration of AI into arbitration, the AAA and Jus Mundi conducted a joint research study. That study employed a mixed-methods approach that combined quantitative survey data gathered between April and July, 2026 with in-depth qualitative interviews conducted with leading arbitrators and law firm partners.

Empirical Scope & Survey Demographics

The quantitative survey gathered responses from 557 respondents from the field of arbitration, during the fielding period. The survey was distributed and promoted specifically within the U.S. arbitration community.

Respondent Organization Type (n = 557)



Respondents span the full range of arbitration practice: arbitrators, whether independent or affiliated with a law firm; law firm counsel from global, mid-size, and boutique firms; in-house counsel; and consultants. The majority of respondents have 15 or more years of experience in arbitration. Because independent arbitrators make up the largest single segment of respondents, findings throughout this report are broken out by organization type wherever the data shows a meaningful difference, so readers can see where their own segment of the field stands.

Throughout the report, we also classify respondents into four adoption archetypes: Leaders, Adopters, Experimenters, and Skeptics, based on actual reported behavior.

Because we are still early on the AI adoption curve, the distinctions between tiers, particularly between Adopters and Leaders, remain modest for now, resting mainly on structure and scale rather than on usage or realized value. These lines are likely to sharpen as AI use matures:

1. **Leaders:** Individuals combining consistent, deliberate use with the formal systems and scale that Adopters lack, structured training programs, formal ROI tracking, firm-wide governance, and defined plans to expand AI investment. The structure that is partial and informal for Adopters is deliberate and practice-wide for Leaders.
2. **Adopters:** Individuals using AI consistently, supported by limited governance and usage guidelines, and basic risk controls. They have not yet built the formal training, ROI measurement, or expansion plans that define Leaders.
3. **Experimenters:** Individuals who have begun using AI but only occasionally and without structure, typically general-purpose tools used in an ad hoc manner, without formal organizational guidelines or data protection policies in place. They are testing out the technology rather than relying on it.
4. **Skeptics:** Individuals with little to no experience using AI. Most have never used the tools and have no plans to. They typically operate without institutional policy, access, or support, and report the lowest trust and benefit of any group.

This framework is introduced in full later in the report and used as a benchmarking lens throughout.

Quantitative Instrument & Analytical Modeling

The quantitative survey comprised 18 questions across five analytical domains:

1. Usage and adoption
2. Sentiment and trust
3. Impact, measurement, and ROI
4. Workforce and career implications
5. Risk, governance, and usage policies

We analyzed quantitative responses using cross-tabulation modeling to identify significant correlations between respondent demographics (such as firm size, role, and usage frequency) and primary attitudinal outputs (including trust metrics, risk perception, and task rejection rates).

A copy of the survey questions can be [found here](#).

Featured Voices

To complement the statistical dataset, the research team conducted structured qualitative interviews with independent arbitrators and partners in arbitration practice groups at law firms.

While quantitative data establishes industry trends, qualitative interviews provide context on adjudicative philosophy, thoughts on non-delegable duties, and procedural dynamics. The insights gathered from these interviews are integrated throughout the report to provide practical framing to the empirical findings.

The following respondents participated in interview sessions:



Jack Levin
Arbitrator/Mediator,
Levin ADR



Jo Colbert Stanley
Arbitrator/Mediator,
Stanley Legal Services LLC



John Hanish
Partner,
Maynard Nexsen PC



Sarah Reynolds
Partner,
Kaplan & Grady



Steven Gregory
Partner,
Gregory Law Firm PC


The Trust Spectrum in Legal AI

Trust is a prerequisite for technological adoption in arbitration. However, the survey data reveals a conundrum: distrust is the initial barrier to adopting AI, but AI usage is precisely what engenders trust in the technology. Initial skepticism prevents some individuals from acquiring the direct operational exposure needed to form a nuanced, evidence-based assessment of legal AI tools.

Usage as the Primary Predictor of Trust

Respondents' trust scales in proportion to operational usage. Daily users report roughly four times the trust score of non-users (3.42 vs. 0.83).

Trust Rises With How Often Someone Actually Uses AI




Daily users report roughly four times the trust of respondents who have never used AI. Crucially, this trust gap does not appear to be driven by software architecture or platform choice. Instead, whether respondents deploy general-purpose LLMs, legal-specific platforms, or purpose-built arbitration engines, average trust scores converge within a range of 2.20 to 2.93 once the individuals regularly use AI.

The clearest divide in the data is between those who use AI and those who don't. That distinction reframes what the survey's average trust score actually tells us. According to survey data across all 557 respondents, the mean score was 2.18, low enough to look like broad skepticism. But that figure is pulled down by the roughly one in five respondents (21%) who have never used AI and don't plan to. As illustrated in the chart above, this group rated their trust at just 0.83 on average. Their low scores reflect distance from the technology, not a considered judgment of it, and they weigh down the overall average enough to obscure how the practitioners actually using AI feel about it.

Institutional Environment and Policy Impact

While individual hesitation may account for initial non-usage, institutional infrastructure can also influence whether individuals feel empowered to bridge the trust gap.

Where Trust Runs Highest and Lowest Across Usage, Practice Area, Firm Type, and Policy



Respondents operating within global law firms (2.83) or organizations with clear, encouraging AI policies (2.98) exhibit higher trust metrics. Institutional backing, through enterprise software licensing, clear risk guidelines, and IT support, appears to provide the psychological and structural safety needed for arbitration professionals to test and integrate these tools.

Conversely, where institutions discourage (1.49) or explicitly prohibit (0.50) AI deployment, trust seems to collapse. Domestic labor and employment practices, the second-largest practice area represented in the survey, and independent arbitrators report similarly low trust levels. This disparity points less to the nature of the legal work than to differences in resource allocation across sectors of the legal market.


The Generational Curve

When broken down by years of practice, the divide sharpens into a distinct generational shift:

- **Junior Respondents (<3 years experience):** Over 50% use AI on a daily or weekly basis, with fewer than 5% ruling out its future adoption.
- **Senior Respondents (15+ years experience):** Regular usage drops below 30%, while nearly 24% state they will never adopt AI tools in their practice.

This generational divergence illustrates that hesitation can be rooted in established habits, heightened risk sensitivity, and comfort with technology, reinforcing the initial barrier to AI adoption.

Usage Habits by Years of Experience



Contextual Trust: Task, Transparency, and Human Verification

A single trust score does not convey how practitioners evaluate AI across specific legal workflows. As Jack Levin, Arbitrator and Mediator at Levin ADR, noted when evaluating the survey's baseline trust score:

"It's trust for what? The question is: what level of trust do you need in a tool before you can rely on it? If I'm asking AI to alphabetize something, I don't need to trust it, because the output is so easy to check. Similarly, cite and quote-checking cases requires a low level of trust. But understanding and applying the reasoning of the case law delegates quite a lot to the tool."

Levin rates AI 5.00 out of 5.00 for legal research, verification, and similar time-intensive tasks, but near zero for tasks touching core legal judgment. Trust in AI is not evenly distributed across ADR tasks; it varies based on two crucial structural factors:

Algorithmic Transparency

Trust depends heavily on whether a platform reveals its underlying reasoning. Sarah Reynolds, Partner at Kaplan & Grady, highlights the distinction between opaque and transparent systems:

“Not all tools are created equally and are designed for different purposes. Some tools are opaque: they provide an output without showing how they reached it, and should be treated with caution. Others show their step-by-step reasoning and sources, allowing you to verify the basis for their output. Those tools warrant greater confidence.”

A tool that shows each step lets you follow its logic and catch where it went wrong, which is why an opaque tool earns less trust than a transparent one doing the same job. But this strengthens trust only when two conditions are met: the user knows the reasoning is available to inspect and takes the time to review it.

Human Verification and Accountability

Transparency is only effective if practitioners perform rigorous verification of the AI’s output. For this reason, Jo Colbert Stanley, an independent arbitrator, shifts the ultimate responsibility from the software to the user:

“Trusting AI comes down to whether people trust themselves to do the extra verification, to take the time to actually read it and double-check everything.”

The record bears this out. A [legal database maintained by Damien Charlotin](#), a research fellow at HEC Paris, has catalogued more than 1,400 court and tribunal filings containing AI-hallucinated citations. Those filings are a direct measure of missing human oversight, not of the technology’s limits. Any tool can produce an error; what allowed these errors to reach a tribunal was not the software but the absence of a human check between the output and the filing. The lesson is not that AI cannot be trusted, but that its use demands verification and that where verification lapses, as Stanley mentions, the failure is a human one.

Ultimately, trust in AI is not about the tool. It’s dependent on the complexity of the task, the transparency of the software, and the rigor of human verification.

Building Trust in AI: Strategic Takeaways

The survey's findings on trust in AI indicate that initial distrust seems to act as a barrier that prevents engagement, while continued non-usage perpetuates that very skepticism.

Breaking this cycle requires more than individual willingness to experiment. It requires AI governance, a term used here to describe a framework through which a firm, institution, or individual practitioner adopts AI. Such a framework establishes which tools are approved for use, what information may be shared with them, how outputs are verified, and where accountability for the result ultimately rests. The trust data presented in this section illustrates why that structure matters. Respondents appear to engage when their environment makes engagement safe, and to retreat when it does not. This work is already well underway at the institutional level. The AAA-ICDR and other arbitration institutions have published standards for the responsible use of AI and adopted rules governing AI's role in arbitral proceedings. Bar associations and courts have issued guidance of their own.

For individuals and firms seeking to build that structure, the survey data analyzed in this section points to three structural realities:

- **Exposure Demystifies Technology:** Trust is built through hands-on interaction within protected environments. Organizations that offer clear guidelines and enterprise-grade software enable practitioners to move past speculative anxiety and develop realistic assessments of AI capabilities through usage.
- **Trust is Workflow-Specific:** For respondents in this survey, confidence varies significantly depending on whether a tool is used for low-risk administrative organization or high-stakes substantive analysis.
- **The Human-in-the-Loop:** Tool transparency and step-by-step auditability are necessary, but only have an impact if humans actually take advantage of those features.


Concerns About AI: Accuracy, Confidentiality, and the Limits of Trust

The previous section established that trust in AI grows with use, and that much of the respondents' apparent skepticism reflects distance from the technology. But distance is not the whole story. Even among active users, caution persists. When asked to identify their primary concerns about AI in arbitration, respondents pointed to a specific and consistent set of risks that could also be hindering their use of AI.

A Field Concerned with Reliability Above All

When asked about their primary concerns, respondents named 2.5 concerns on average, and 71% flagged three or more. But the concerns cluster clearly around a small number of issues.

Accuracy and Confidentiality Dominate the Field's Concerns




AI inaccuracy, the risk of “hallucinations,” stands far above every other concern, cited by 78% of respondents. Confidentiality and data security follows at 48%, then ethical concerns (41%), potential bias (31%), and lack of transparency (28%). The concerns that dominate public commentary about AI and the professions, such as cost and job displacement, sit at the very bottom, each named by fewer than 10% of respondents.

When it comes to AI in arbitration, respondents are worried about whether the technology can be relied upon to produce accurate, confidential, and defensible work. These are professional risks with direct consequences; a confidentiality breach, for example, can compromise a proceeding and a client relationship. The concerns respondents named are the ones that bear directly on their professional obligations.

Concerns Become More Concrete with Experience

The types of concern respondents report shift as they gain exposure to AI.

Primary Concerns Among AI Users and Non-Users



Among respondents who have never used AI, ethical concerns rank highest (56%). “Ethical concerns” likely captures a range of worries depending on the respondent, from disclosure and appropriate use to broader questions about the use of AI in a process built on human judgment. Among regular users, the figure falls to 35%.

This likely reflects a combination of direct experience and the growing body of guidance from institutions and bar associations that has given individuals clearer answers to these questions. Whatever the cause, the concern is less acute among those closest to the technology.

Confidentiality moves in the opposite direction. Cited by just 31% of non-users but 54% of regular users, it becomes a leading concern once practitioners are working with these tools regularly and confronting the concrete challenge of protecting privileged information in systems they may not fully control.

Transparency follows its own path: concern about “black box” reasoning is high among non-users (41%) and lower among regular users (22%). As experienced practitioners we interviewed note, the concern is real and operational, it bears directly on whether output can be verified, but it is one that active users can learn to manage by choosing tools that show their reasoning and treating opaque ones with caution. Rather than vanishing, the concern possibly converts into a criterion for tool selection.

One concern holds steady regardless of experience. Fear of inaccuracy and hallucinations is cited by roughly three-quarters of both regular users and non-users, a consistency that sets it apart from every other concern in the survey and suggests it reflects a persistent feature of the technology. However, that persistence is not necessarily a sign that the risk is unmanaged. Among experienced users, the concern about inaccuracy is real but routinely handled through verification.

John Hanish, Partner at Maynard Nexsen PC, described how that works today:

“Before anyone relies on it, they go back and check, make sure the cases are really cases, that everything the tool gave you actually exists and isn’t just made up.”

The more revealing concern, in his view, is not the technology’s current tendency to err but what happens as it improves:


“Right now we check everything because we don’t fully trust it. Attorney judgment, confidentiality, and accountability remain non-negotiable. The risk is that as it gets more accurate, people stop checking. You can imagine someone a few years from now saying, ‘we used to check everything back when we worried about hallucinations, but it’s fine now, the computer’s never wrong.’ That’s exactly when things slip through.”

This reframes the report’s most-cited concern in a useful way. The danger Hanish describes is that growing reliability could erode the verification habit that currently keeps its errors in check. It is a concern that guidance and firm policy can speak to directly, by making verification a fixed professional expectation rather than a discretionary step that fades as trust grows.

Concerns are Largely Shared Across Organization Types

Broken down by organization type, the field’s concerns are remarkably consistent. As chart 2.3 shows, fear of inaccuracy is near-universal regardless of firm size, cited by 77% to 84% of every segment.

Primary Concerns by Firm Type



A few differences are worth noting; confidentiality concern is highest among boutique firms (59%). Global firms (50%) and independent arbitrators (47%) sit close to the field average, with mid-size firms slightly lower (44%). Independent arbitrators stand out on two counts, citing both potential bias (34%) and ethical concerns (42%) more often. This aligns with the arbitrator’s distinct position: as the neutral responsible for a fair and impartial outcome, bias in an AI tool is not merely a quality issue but a threat to the integrity of the role itself.

What stands out, however, is the consensus. Across firm sizes and roles, the field worries about the same things in roughly the same order.

Concern as a Blueprint for Guidance

The significance of these findings becomes clear when read alongside what the profession says it wants. As later sections detail, the single most requested accelerant of responsible AI adoption is continued guidance. The data from this survey around respondent concerns gives a clue as to what that guidance should address.

In this report, respondents have clearly identified their risks: accuracy first, confidentiality second, as well as ethical and transparency questions, which weigh most heavily on those least familiar with the technology. Each of these can be addressed through standards and best practices. Guidance on verification protocols speaks to the fear of inaccuracy. Guidance on data handling and tool selection speaks to confidentiality. Guidance on disclosure and acceptable use speaks to the ethical questions that keep hesitant users on the sidelines.

The Judgment Line


The previous sections showed that trust in AI grows with use, and that the respondents' caution centers on a specific set of risks. That caution finds its sharpest expression in a boundary the data unearthed: a threshold separating administrative and analytical preparation from AI having a primary, as opposed to supporting, role in adjudicative decision-making.

This demarcation, referred to throughout this report as the “Judgment Line”, is consistent across all respondent cohorts. It does not change with years of legal experience, firm size, or individual frequency of AI usage. To understand where the arbitration community is prepared to use AI primarily, one must first examine where it currently is not.

Categorizing Workflows: Preparation vs. Adjudication

To map the boundaries of AI integration, respondents were asked two complementary questions: the phases of arbitration where they currently deploy or would deploy AI, and the phases they believe should not be primarily handled by AI.

Where AI Is Welcome, and Where the Profession Draws a Line



Mapped side-by-side, two categories of dispute resolution workflows emerge:

Pre-Adjudicative Preparation

Document review (51% usage, 9% rejection), legal research, case management, and translation exhibit broad acceptance and rapid integration. These tasks share three characteristics:

- They involve sorting, summarizing, and retrieving unstructured data.
- They operate as inputs into a broader legal analysis rather than final outputs.
- Their outputs can be independently verified against source records by human practitioners.

The Adjudicative Core

Other tasks reveal skepticism about AI taking on a primary role in the work. On award drafting, 58% of respondents state that AI should not primarily handle the task, more than double the 26% who report they do or would use AI for this purpose. Evidentiary hearings display a somewhat wider gulf, with 49% rejecting primary AI involvement compared to 14% operational usage. Drafting submissions and post-award enforcement follow the same pattern, with respondents still seeing human judgment as the primary driving force.

The divide isn't about difficulty or time saved. It's about whether the task ends in a decision. Arbitration's authority has always rested on a human tribunal weighing conflicting evidence, judging witness credibility, and taking personal responsibility for the award.


Respondents are willing to let AI prepare the inputs to that decision. They are currently less willing to let it take the primary role in making the decision.

Rather than easing with familiarity, this boundary tends to hold among those most experienced with AI.

The Infrastructure Paradox: Familiarity Sharpens the Line

A common assumption in legal technology is that institutional sophistication and heavy tool usage will inevitably lead well-resourced law firms to hand off increasingly complex tasks to AI. The empirical data runs counter to this premise.

The Better-Resourced the Firm, the Firmer the Line on Awards



Seventy-five percent of respondents at global law firms maintain that award-drafting should not be primarily handled by AI. In contrast, independent arbitrators, who represent the least-resourced segment in terms of dedicated IT infrastructure, prove slightly more open, at 55%, to AI driving that process. The global law firms that use AI most heavily elsewhere, such as for document review, hold the tightest line on AI being the primary awards drafter.

The data from this survey suggests familiarity with AI does not erode the Judgment Line; it sharpens it. Resistance among advanced users may not stem from inexperience or fear, but from a clear understanding of the capabilities and limitations of AI as well as the non-delegable nature of decision-making.

Qualitative Perspectives: Accountability, Mandate, and Personal Judgment

The qualitative interview findings align with the quantitative data. Levin articulates a boundary regarding award drafting:

“The award is the one thing that’s mine. It’s my judgment, my reasoning, my name on it. I’ll use AI to check a citation or organize my notes, but the decision has to be mine.”

Steven Gregory, Arbitrator at Gregory Law Firm PC, connects the boundary to procedural accountability and legal standing:


“The problem isn’t whether AI could draft a competent award. It’s that when a party challenges that award, I’m the one who has to stand behind every word of it. I can’t stand behind reasoning I didn’t do myself.”

Neither objection rests on the premise that generative AI produces poor outputs. Both accept that language models can draft coherent text. The line holds because of human legal responsibility for the final award. Levin and Gregory do not believe that a tribunal’s responsibility for the ultimate decision in a case should be primarily delegated to a tool.

Strategic Implications of the “Judgment Line”

These survey results are a map for anyone building AI into arbitration practice. In chart 3.1, the tasks listed at the top of the chart (document review, research, case management) are not just where AI is accepted, they are where adoption meets the least resistance and delivers value fastest.

The tasks listed at the bottom of chart 3.1 are a different matter, and the lesson isn’t that AI has no role there. Award review, where 28% say AI should not be the primary driver, is far less controversial than award drafting, where 58% say AI should not be the primary driver.



Respondents are more open to AI checking a draft award they wrote, but just over half are currently less comfortable with AI being responsible for writing the award. The opening here isn't in replacing judgment, it's in supporting it: catching an inconsistency, flagging a missed argument, surfacing a citation the drafter overlooked. Essentially, a tool that helps an arbitrator pressure-test their own reasoning. And whenever AI is used, human oversight and verification should be a critical component of every step.


Measuring the ROI of Legal AI

While mapping the “Judgment Line” defines where respondents currently are comfortable with AI usage, evaluating the actual value realized from adoption reveals a split between significant efficiency gains and persistent operational barriers.

Measuring Realized Value: The Efficiency Paradox

When survey respondents were asked to identify the specific tangible benefits gained from using AI, the dataset highlighted a sharp divide between time savings and systemic adoption friction.

The Biggest Benefit Is Time, but a Third See No Benefit Yet




While 43% of respondents cite time savings as the primary benefit, 37% report realizing no significant benefit at all, splitting the dispute resolution community between those deriving immediate value and those who haven't yet seen any measurable return.

Furthermore, far more value is seen in raw operational speed than in improved quality (e.g., output accuracy at 16%, deeper analytical insights at 16%). Most respondents view AI as a utility to reduce hours spent on labor-intensive tasks rather than a tool that fundamentally transforms legal analysis.

Structural Disparities Across Practice Types

The 37% of respondents reporting no measurable benefit are not randomly distributed across the field; their distribution reflects differences in access to institutional infrastructure, software licensing, and organizational support.

Boutique Firms Get the Most Out of AI; Solo Arbitrators the Least



Structural Disparities and Infrastructure Support

Boutique practices report realizing the greatest value of any organizational category by a decisive margin: two-thirds cite significant time savings, while only 15% report no measurable benefit. Global full-service and mid-size firms occupy a middle tier, with approximately half observing tangible time savings. Independent arbitrators are the outlier in the opposite direction, and the only group where “no benefit” (44%) actually outranks time savings (37%).


As demonstrated across the survey findings, respondents deriving the greatest value from AI are those engaging with it most often. Daily active users are overwhelmingly supported by institutional law firm environments that provide software licensing, explicit permission to experiment, and peer-learning networks. Conversely, independent arbitrators typically acquire software independently, operating without dedicated IT departments to conduct due diligence or peer groups to refine usage protocols. The nature of the work adds to this: an arbitrator’s role sits closest to the Judgment Line the field is least willing to cross.

Boutique Firms Outpace Others in AI Efficiency

The performance of boutique firms challenges any assumption that global law firms with expansive financial resources inevitably extract the greatest value from emerging legal technology.

Boutique practices lead across nearly every usage and efficiency benchmark in the study: they record the highest rate of significant time savings (66%), the lowest proportion of non-users across all organizational categories (3.4%, outperforming even global law firms), and the highest concentration of daily active users.

This operational edge is structural. Boutique practices often possess sufficient capital to license enterprise-grade legal AI platforms and foster shared operational habits, yet may remain unencumbered by the multi-layered committees, pilot programs, and protracted IT procurement cycles characteristic of larger institutions. They may be able to evaluate a specialized tool, establish usage protocols, and deploy it more quickly. This combination, capital resources less inhibited by administrative friction, can be a competitive advantage for small but sophisticated firms.




Workforce and Career Implications

Any conversation about AI inevitably confronts a central question: the long-term impact on human capital. Arbitration is no exception, and the survey put this question directly to individuals across the field.

Reassessing Role Replacement: The Confidence Baseline

Most Practitioners Aren't Worried AI Will Replace Them



When asked to evaluate the degree to which AI might partially or fully replace aspects of their professional role, the respondents expressed low levels of anxiety.

Half of all respondents rated their level of concern at either 0 or 1.00 out of 5.00, while a modest 13% reported high concern (4.00 or 5.00). The overall average score of 1.70 represents one of the lowest scores recorded across the survey.


This confidence baseline partly reflects the demographic composition of the respondent pool. The professional standing of experienced arbitrators and lawyers rests on adjudicative judgment, strategic authority, and institutional reputation. However, the low baseline in the survey data may not merely be an artifact of seniority but may also reflect attitudes regarding the current limits of AI and the role of human judgment in dispute resolution.

The Functional Bifurcation: Task Absorption vs. Skill Appreciation


Low displacement anxiety does not imply that respondents anticipate an unchanged working environment. On the contrary, respondents expect a substantial portion of dispute workflows to be absorbed by AI over the coming years, divided across two categories:

AI Absorbs the Task Work; Judgment Becomes More Valuable

Where AI is expected to absorb the work



Skills expected to matter more, not less



The respondents expect AI to absorb labor-intensive, high-volume tasks whose value historically derived from manual effort and scale: document review (51%), proofreading and cite-checking (47%), timeline creation (43%), and legal research (40%). Conversely, skills anchored in normative evaluation, ethical stewardship, and strategic advocacy are expected to appreciate in value: strategic judgment (52%), ethics and governance (42%), substantive legal expertise (39%), and the ability to develop facts and weigh evidence (30%).

AI is stripping away process-oriented execution and placing higher premiums on analytical judgment. The tasks suitable for delegation become standard AI inputs, while the attributes that define professional expertise become the primary basis for service valuation.

Hanish observes this transition in the day-to-day management of junior associates:

“Instead of an associate writing the first draft from a blank page, AI handles that. They still spend the hours, but at the back end, refining. That should help associates spend more of their billable time making the work product better, not producing boilerplate.”

This shift moves junior legal work up the value chain, transitioning the associate’s primary function from initial drafting to critically evaluating and refining AI-assisted outputs.


Similarly, Gregory emphasizes that, as procedural drafting becomes automated, junior practitioners must cultivate core advocacy skills earlier in their career trajectory:

“Associates need to learn how to present a case to a panel. You’ve only got one shot in an arbitration hearing. Preparation, demeanor, the way you approach the room, those are the skills young lawyers, and older ones too, need to have. AI can’t replicate them.”

Taken together, these perspectives point towards a fundamental shift: the core competencies of senior practice, strategic positioning, evidentiary evaluation, and oral advocacy, may be moving from mid-career milestones to entry-level prerequisites.

Displacement Concerns

The More Practitioners Use AI, the More Aware They Are of the Stakes



The Relationship Between Usage & Displacement

The study's empirical data reveals that those closest to AI are slightly more concerned about its impact: respondents who never use AI report the lowest concern of any group (1.59 out of 5.00), while regular and daily users report the highest (1.80 and 1.75, still relatively low). One likely explanation is exposure. Practitioners who use AI regularly have seen firsthand what it can do and how it has evolved, which may give them a clearer sense of how much of their work is within its reach. On this reading, their concern reflects familiarity with the technology's trajectory rather than unease with the unfamiliar.

Displacement Concern by Role

Broken down by role, displacement concern is mostly uniform. Partners (1.71), arbitrators (1.67), and associates (1.50) all cluster near the low average, which fits a field that sees AI absorbing tasks rather than roles. One group breaks this pattern: in-house counsel reported an average concern of 3.20 and, while still moderate, it is the highest of any role or segment in the data.

That figure rests on a smaller pool of respondents and, as such, is a signal rather than a conclusion. The review, research, and drafting work that corporate legal departments increasingly handle in volume and increasingly look to automate are exactly the tasks most exposed to AI disruption.

Strategic Implications for Practice Management

Taken together, these findings point to a workforce in transition. The survey shows respondents expect AI to absorb the task-heavy work; document review, first-pass research, proofreading, while the skills they believe will matter most (strategic judgment, substantive expertise, advocacy) are precisely the ones AI cannot currently replicate. The shift is in the work itself, away from routine tasks and toward the judgment and expertise that define the role.

Two things follow for how the profession should prepare:

1. **Shift Focus from Labor Volume to Analytical Discernment:** The market value of legal services will increasingly detach from billable hours dedicated to manual processing. Sustainable competitive advantage will belong to those who leverage automated tools to compress research and review cycles, while redirecting billable capacity toward strategic advice and evidentiary presentation.
2. **Re-Architect Legal Training Programs:** As routine drafting and research increasingly lean on AI for a first pass, the profession's distinguishing skills (verification, judgment, risk assessment, and advocacy) become more central. This has implications for how junior lawyers are trained, too: the apprenticeship model will need new ways to develop those instincts as AI takes on the groundwork.

The Institutional Access Divide


The data reflects an organizational divide on AI engagement, trust, and perceived ROI. While this disparity may be attributed to multiple factors, the data reveals workflow asymmetry as a possible driver behind this difference.

Infrastructure Support vs. Operational Isolation

Respondents who identified as working in law firms (global, mid-size, or boutique) and those who identified as independent arbitrators (solo practice) operate in starkly different environments:

AI Policy and Adoption Plans:

Firm-Based Respondents vs. Independent Arbitrators



Firm-based respondents are more than twice as likely to work under a clear AI policy (38% versus 15%), while solo arbitrators are nearly three times as likely to have no adoption plans (29% versus 11%). Tellingly, 43% of arbitrator respondents report no AI policy at all, more than double the 20% at firms, though for a solo practitioner that figure reflects the absence of a policy they would have to set for themselves rather than one handed down by an organization.

Additionally, among independent arbitrators, only 1% measure ROI, 10% name cost as a primary obstacle, and 29% report no adoption plans whatsoever.

Law firms often provide an institutional ecosystem: clear usage parameters, centralized software budgets, peer-learning networks, and dedicated IT security audits. Independent arbitrators must assemble this infrastructure individually, evaluating data privacy standards, funding software licenses out-of-pocket, and establishing prompt protocols without peer support.

Hanish described what that infrastructure looks like from inside a firm:

“My firm has more than 600 lawyers, and we’re probably at 600 different stages of adopting AI. Some lawyers have been leading the firm’s evaluation of platforms for at least two years while others have relied on the firm to choose providers and provide training, which is more recent and ongoing. The firm uses secure, enterprise-grade technology to support our attorneys and they remain accountable for every deliverable.”

Workflow Positioning Along the Judgment Spectrum

The disparity in AI adoption between law firm respondents and independent arbitrator respondents is not merely an infrastructure deficit; it is an inherent reflection of their respective roles in the arbitration ecosystem:

1. **Law Firm Respondents (High AI Susceptibility):** Workflows are dominated by pre-adjudicative processing. These high-volume, labor-intensive tasks sit squarely in the zone of high AI acceptance as well as susceptibility to AI automation, delivering clear time savings and high ROI.
2. **Independent Arbitrators (Low AI Susceptibility):** While some independent arbitrators do leverage AI tools for more rote arbitral tasks, important aspects of the arbitrator’s mandate reside beyond the “Judgment Line.” Sole neutrals and tribunal members are personally tasked with weighing conflicting testimony, conducting oral hearings, deliberating, and drafting final awards, functions where respondents currently prefer AI to not have primary responsibility for work product (58%).

Consequently, an arbitrator’s lower daily usage frequency and modest reported ROI are not necessarily indicators of technological lag; rather, they may reflect a practice model where many feel the ultimate task, taking ultimate responsibility for the decision, remains non-delegable.

Read together, these findings reframe the access divide. Part of the gap is structural and fixable: independent arbitrators lack the enterprise licenses, budgets, training, and peer networks that firms provide, and those are gaps that guidance, affordable tools, and institutional support can close. But part of it is inherent to the work itself. Lower adoption among independent arbitrators, then, reflects two different things at once; a resource gap worth closing, and a role that is, by design, less exposed to AI in the first place.

AI Adoption by Organizational Profile

AI adoption, trust, realized benefit, and organizational policy vary across different segments of the arbitration field and we profiled each segment separately based on respondent data.

Segment Scorecard: How Each Firm Type Compares

	GLOBAL LAW FIRM	BOUTIQUE FIRM	MID-SIZE FIRM	INDEPENDENT ARBITRATOR
TRUST IN AI (0-5)	2.83	2.45	2.48	2.02
REPORT TIME SAVINGS	50%	66%	51%	37%
SEE NO BENEFIT YET	21%	16%	25%	44%
CLEAR, ENCOURAGING POLICY	83%	40%	29%	15%
DAILY AI USE	21%	16%	8%	7%
NEVER USE AI	8%	3%	14%	27%
VERY SATISFIED W/ APPROACH	54%	36%	33%	41%

Global Law Firms: the Most Structurally Advanced

Global law firms rank at or near the top of nearly every institutional measure in the survey. Eighty-three percent of respondents at those firms operate under a clear policy that encourages AI use, more than double the next-highest group. They lead on formal ROI measurement (21%), daily AI usage (21%), and satisfaction, with 54% reporting they are “very satisfied” with their organization’s approach, the highest rate of any segment. Global firm respondents also express the greatest confidence in their competitive standing: 17% consider their organization “far ahead” of the market, and not a single respondent in this segment indicated their organization lacks a defined AI strategy.

This profile reflects the structural advantages available to large, well-resourced organizations: dedicated budgets for enterprise tools, staff capacity to develop and enforce policy, and the scale required to measure returns systematically. The data confirms a consistent pattern throughout this report: where organizations provide clear guidance, adequate resources, and institutional permission to use AI, trust and adoption follow.

Boutique Firms: the Highest Realized Value

Boutique firms yielded the most notable finding, one that runs contrary to the assumption that AI adoption scales with firm size and resources.

Boutique firms' respondents reported the highest levels of realized value of any segment in the survey. Two-thirds (66%) cite significant time savings from AI use, compared with 50% at global firms and 51% at mid-size firms. Only 16% report no meaningful benefit to date, the lowest rate of any group. Boutique firms also record the lowest rate of non-use (3%) and the highest share of daily users among all firm types.

While only 40% report having clear AI policies, well below global firms (83%), and 17% report having no defined AI strategy, their agility affords them an advantage. A boutique firm may operate at a scale where tool adoption decisions can be implemented quickly, shared across a small team, and refined through direct feedback without the procurement cycles, committee approvals, and phased rollouts that can slow adoption at larger organizations. This agility may be the reason why the gap between boutique and global law firms is less pronounced than it otherwise might be.

Mid-Size Firms: Steady Adoption

Mid-size firms represent the stable middle of the field's AI adoption landscape. Half report significant time savings (51%), their average AI trust score (2.48) sits just below the overall firm average, and a higher share of mid-size firm respondents describe their organization as "on pace" with the market than any other group.

The data points to a segment where adoption is proceeding at a measured rate without a clearly defined strategic framework to support it. The rate at which mid-size firms report having clear AI policies (29%) is the lowest among firm-based respondents, which may indicate that the primary challenge for mid-size firms is not practitioner reluctance but organizational structure. Usage is already occurring; what is lacking is the formal policy, consistent governance, and measurement practice

that would convert that usage into deliberate, accountable strategy. This may come down to administrative infrastructure challenges, financial resource constraints, or a lack of client messaging on this topic. This gap represents the clearest near-term opportunity for this segment.

Independent Arbitrators: Limited Institutional Support

Independent arbitrators, the majority of survey respondents, report the lowest levels of institutional support of any group. That lack of support likely results in their consistent scores: the lowest average trust score (2.02), the lowest rate of significant time savings (37%), the highest rate of no meaningful benefit (44%), and the highest rate of non-use (27%).

Despite these indicators, 41% of independent arbitrators report being "very satisfied" with their current approach to AI, a higher rate than boutique or mid-size firms. This figure might reflect the autonomy inherent in solo practice: without an organizational framework setting pace or expectations, there is no institutional friction to generate dissatisfaction.

This distinction is consequential. Independent arbitrators in this survey hold relatively consistent views on AI's potential value and identify concrete, practical barriers (cost and access to training) as the primary impediments to their adoption of AI. The satisfaction reported by this segment should not be read as an absence of need. It sits alongside a genuine and well-documented resource gap.

Professional Role as a Lens on Adoption


Law firm partners report trust levels and usage patterns broadly aligned with the overall firm averages: an average trust score of 2.51, low concern about role displacement (1.71), and a relatively even distribution between general-purpose AI tools (50%) and legal-specific platforms (44%). This profile reflects a segment that has incorporated AI into practice at moderate rates without significant anxiety about its implications.

Associates and in-house counsel appear in smaller numbers within the dataset, but each presents a finding of note. Associates report the highest rate of legal AI use of any role (62%), consistent with the broader generational pattern identified in this report: AI adoption is more prevalent among respondents earlier in their careers.

In-house counsel have an average displacement concern of 3.20, nearly double the overall survey average and the highest of any professional role or organizational segment, paired with the heaviest overall AI usage rate. This combination suggests that the individuals most embedded in AI-assisted workflows are simultaneously the most attuned to its potential implications for their roles.

Implications Across the Field

Considered together, the segment profiles reveal that AI advantage in arbitration practice is not principally a function of firm size or market standing. Global firms lead on structural maturity; boutique firms lead on realized value; mid-size firms are advancing steadily within a governance gap they are well-positioned to close; and independent arbitrators face barriers that may be more practical rather than attitudinal.




The Maturity Model

To move beyond rigid demographic cuts, the study evaluated respondents based on core operational behaviors: trust, usage frequency, AI policy, impact and ROI measurement, and future expansion commitments.

Combining these metrics produces our AI Maturity Model for Arbitration, which categorizes the field into four distinct operational archetypes: Leaders, Adopters, Experimenters, and Skeptics.

Four Levels of AI Maturity



The Four Maturity Archetypes

Because AI adoption across arbitration is still at an early stage, the factors separating the most advanced practitioners from one another are, for now, relatively few. In our data, what distinguishes Leaders from Adopters comes down to two key factors: structure and scale. Leaders apply formal systems to their AI use and deploy it more widely; Adopters use AI consistently but manage it more informally and on a smaller scale.

That the difference is this contained is itself a sign of how early the field is; it is early enough that even the most advanced firms haven't yet built out the full apparatus that will eventually distinguish them from peer organizations. As practitioners and firms mature in their use of AI, Leaders are likely to develop a wider set of differentiators: allocated roles, dedicated oversight, purpose-built workflows, and more. As those capabilities accumulate, the distance between the tiers may grow: the categories may become more distinct, and the value Leaders derive from AI is likely to expand. What today is a modest gap of structure and scale is, in other words, an early snapshot of a separation that seems likely to widen as the whole field advances along the maturity curve.

Skeptics and Experimenters are separated by the most basic threshold in the model: whether an individual uses AI at all. Skeptics have largely stayed on the sidelines, most have never used the tools, and they report the lowest trust and benefit in the survey. Experimenters have begun to use AI, but occasionally and without structure, testing it rather than relying on it. The step from that experimental use to the consistent, deliberate use that defines Adopters is where trust and value climb most sharply; it is the largest jump in the model.

With that trajectory in mind, the four tiers as they stand today:

Leaders (13% of Respondents)

- **Profile:** Everything Adopters do, plus the formal systems and scale they lack: structured training, formal ROI tracking, firm-wide governance, and defined plans to expand AI investment, sometimes led by a dedicated AI or innovation team. The structure that is partial and informal for Adopters is deliberate and practice-wide for Leaders.
- **Metrics:** Highest trust score (3.32); 65% report time savings; 100% plan to expand their use.
- **Key Finding:** Leaders' time savings (65%) are nearly identical to Adopters' (66%). Today, the two tiers differ more in how systematically and formally they manage AI than in what they get from it, a reflection of how early the field still is. As Leaders build out the deeper structures described above, that is the gap most likely to widen, and to translate into a growing advantage in the value they extract.

Adopters (22% of Respondents)


- **Profile:** Consistent use but managed informally. Governance is partial: a few usage guidelines, some approved tools, basic risk controls, applied case by case and ad hoc. They have not yet built the formal training, ROI measurement, or expansion plans that mark Leaders.
- **Metrics:** Composite trust reaches 2.78; 66% report significant time savings; 90% express satisfaction with their organization's approach.
- **Key Finding:** The move from occasional use to structured use produces the single largest jump in realized value in the entire model.

Experimenters (33% of Respondents)

- **Profile:** Active ad-hoc users operating without formal organizational guidelines or closed enterprise infrastructure.
- **Metrics:** Active tool deployment, only 3% are non-users; trust score rises to 2.28; 46% report tangible time savings.
- **Role in Market:** Represents the field's primary transition group: those who have accepted the utility of AI but lack institutional structure to optimize its use.

Skeptics (31% of Respondents)


- **Profile:** Low usage, high security anxiety, and minimal institutional support.
- **Metrics:** 64% have never deployed AI; composite trust score sits at 1.15; only 14% report deriving any operational time savings.
- **Composition:** 36% of independent arbitrator respondents fall under the Skeptics category. This is the segment furthest from AI adoption.



Organizational Distribution Across Archetypes

Mapping these four archetypes across organizational types reinforces the structural dynamics identified throughout the report.

Maturity Isn't About Size: Where Each Firm Type Lands




Global law firms dominate the Leader category (50%) due to extensive financial and administrative resources, followed by boutique practices with 60% falling into the top two tiers (Adopters and Skeptics). While not reaching the maturity level of global firms, boutiques accelerate their maturity likely through operational agility rather than capital expenditure, bypassing corporate bureaucracy to deploy enterprise legal platforms rapidly.

Independent arbitrators concentrate at the bottom, with 36% Skeptics and only 8% Leaders, which is the maturity-model expression of the access divide.

Accelerating AI Adoption

When respondents were asked to rank the initiatives that would most effectively accelerate responsible AI adoption across dispute resolution, their preferences revealed a clear consensus: respondents cited continued authoritative governance and domain-specific tooling rather than cost reductions or commercial metrics.

What Would Most Accelerate Responsible AI Adoption



The Demand for Continued Institutional Guidelines Over Commercial Metrics

The two highest priorities, Institutional Guidelines (30%) and Domain-Specific Tools (24%), account for more than half of all primary responses. Conversely, conventional commercial drivers such as lower licensing costs (3%) and ROI evidence (4%) rank at the bottom of practitioner priorities.

This ordering underscores a fundamental characteristic of the sector: dispute resolution is driven by risk mitigation and professional duty. What ranks highest is not lower cost or better technology, but clarity, clear ethical parameters, confidentiality safeguards, and continued institutional guidance. While cost is a genuine barrier for some practitioners, it is not the primary factor holding back AI adoption.

A Two-Stage Adoption Strategy

When practitioners were asked what would most accelerate responsible AI adoption, continued guidance and purpose-built tools ranked well ahead of lower cost, proof of return, or client demand. This is a field asking to be shown what responsible AI use looks like and to be given tools equal to the work.

The profession's two largest groups, Skeptics and Experimenters, are held back by distinct barriers and require distinct forms of support. A uniform push toward adoption would address only one of them. The data instead points to a sequenced, two-stage strategy.


Phase 1: Guidance and frameworks for Skeptics: Skeptics need a clear standard for what responsible use looks like. Continued guidance from arbitral institutions, covering disclosure, data privacy, and acceptable use cases, provides that standard and gives practitioners the confidence to begin experimenting. Arbitral institutions, including the AAA, have started this work by developing frameworks addressing responsible AI use in arbitration.

Phase 2: Specialized tooling for Experimenters: Once practitioners are actively using AI, their constraint shifts from guidance to capability. Experimenters rank purpose-built arbitration tools as their top requirement, because general-purpose AI is not designed for the specific demands of dispute resolution work. Arbitration comprises many distinct workflows, each of which a specialized tool can be engineered to serve. In arbitration research, for instance, a single question may require weighing treaty provisions, institutional rules, and conflicting national laws at once, an interplay a purpose-built tool can navigate in ways general-purpose systems cannot.

For those already ahead: Structured Adopters and Leaders, together a third of the profession, have largely cleared both hurdles. Their task is no longer adoption but consolidation: turning consistent use into measured, well-governed practice, and sustaining it as tools and expectations evolve. As the maturity model showed, the two groups capture nearly identical value today; what distinguishes Leaders is not better results but greater rigor, formal measurement, training, and a definitive plan to scale. For this segment, the priority is protecting and extending an advantage already in hand rather than building one from scratch.

In sequence, the logic is straightforward. Skeptics need clarity before they will begin; Experimenters need better tools once they have that clarity; Adopters and Leaders, already ahead, need the disci-

pline to sustain and scale what they have built. A profession moving at four different speeds is not served by a single approach. Meeting each group where it stands may move the greatest number of practitioners forward.



Looking Ahead

The findings in this report describe a field navigating an early and uneven moment with AI. This closing section translates those insights into strategic direction: what the data anticipates next, and the concrete measures individuals and firms can adopt today.

Where This is Heading

Three underlying patterns in the data illuminate the likely future trajectory of the field.

1. Accelerating, yet Asymmetric, Adoption

Adoption will likely continue its upward trajectory, but at different rates for different segments. Early-career respondents deploy AI at more than twice the rate of their senior peers, and this will likely continue as AI is increasingly employed for entry-level tasks, and by the early-career segment. The Skeptics, currently at 31%, will likely decrease in number over time. However, if the access divide continues, this rising tide may not lift everyone equally. Without the right support, investment and, infrastructure, firm-backed practitioners likely will continue to pull ahead of the solo majority, maintaining or widening the resource gap captured in today's data rather than closing it.

2. The Evolution of the Judgment Line

The Judgment Line is currently holding but it is unclear whether individuals will grow more comfortable delegating award drafting or witness credibility determinations in large, complex arbitration cases to AI as software capabilities improve. Their current resistance is not a function of technological limitations; instead, it is a matter of accountability, specifically, who arbitration holds responsible for the final decision in a case. That boundary may prove durable as underlying models advance, or it may lessen as individuals weigh time and cost against increased AI capabilities.

3. The Shift Toward Continued Authoritative Guidance

Continued institutional guidance is necessary, and respondents are demanding it. The single clearest priority identified in the survey is for institutional standards and best practice directives. That demand has been met already by some arbitral institutions, and bar associations will continue to be important as AI usage increases in the field.

Recommendations for Practitioners

For individual practitioners across the dispute resolution spectrum, the data points to a series of concrete strategic considerations rather than a generic call to “embrace AI.”

Initiate usage on closed case records: The highest-value step in the entire maturity model is simply beginning to use AI, and the lowest-risk way to begin is on work where the outcome is already known or easily verifiable. For example, deploy the tool against a finalized matter where every factual assertion and authority can be verified against an established record. This calibrates user judgment without putting live client work at stake.

Institutionalize verification as a permanent habit: The flawed filings appearing in sanction rulings originate from a myriad of professionals from across the field. It is important to verify outputs even after the tool has earned your trust.

Invest where judgment resides: The survey is specific about which capabilities will become more important as AI absorbs high-volume task work: strategic judgment, substantive legal expertise, and the ethical sense to apply both. These core competencies appear likely to remain indispensable, particularly in large, complex cases. Leverage AI to clear away administrative volume so a greater share of your time is dedicated to the judgment a human practitioner can provide.

Recommendations for Firms

For firm leadership, the maturity model offers an actionable playbook, one that begins by recognizing that different cohorts within an organization require distinct interventions.

Target adopter status across the fee-earner base: The data demonstrates that nearly all realized operational value is captured by the time a practitioner achieves Adopter status: characterized by consistent usage, guided by policy, and evaluated with light metrics. Firms should focus on moving their fee-earners reliably past the two transitions that matter most: starting usage and adding structure.

Align your strategy with the archetype: Skeptics and Experimenters need different things: Skeptics need clear standards and a low-risk way to begin; Experimenters need domain-specific tools and better workflows. A single, blanket AI policy addresses neither well and tends to leave both groups underserved.

Do not mistake organizational scale for AI maturity: The boutique firms in this survey reached the top maturity tier at nearly double the rate of mid-size firms. For smaller practices, this suggests that an AI advantage is genuinely within reach; for larger institutions, it underscores that organizational scale and the strategic will to execute must work in lockstep.

Conclusion

A profession encountering AI is often expected to divide into camps: those who embrace it and those who resist, those who adapt and those who fall behind. This survey did not find that division. On the questions that matter most, where AI helps, where it falls short, and where its use currently may be considered more controversial, arbitration practitioners are broadly aligned, regardless of firm, role, or years in practice.

That agreement centers on a clear distinction: between the work that prepares a decision and the act of drafting a final award. Practitioners have readily adopted AI for the former and have currently been more reluctant to adopt it as the primary driver for the latter. The reason they give is not the technology's current limitations but accountability, who answers for a decision once it is made. Whether this line holds as tools improve and institutional norms develop remains to be seen. For now, it is the clearest signal in the data of how this profession intends to adopt AI: on its own terms, and without ceding the judgment at its core.

The principal risk in this picture is unequal access to AI. The respondents gaining the most from these tools are generally those with an organization behind them, one that provides the software, the training, and the freedom to experiment. Independent arbitrators have the least of that support, the core work of their work is also closer to the judgment line, and they report the lowest usage and the least benefit as a result. This gap is neither inevitable nor permanent, and how the profession addresses it, through guidance, affordable tools, and shared standards, will largely determine who benefits from what comes next.

This report is a snapshot of a profession early in a long transition. The findings will change, as the technology, the rules, and the practitioners themselves change, and this is why we will continue updating this report annually. But the direction is already visible: arbitration is adopting AI thoughtfully, though unevenly, and on terms that keep human judgment where the profession believes it belongs.

About Jus Mundi

Founded in 2019 and recognized as a mission-led company, **Jus Mundi** is a pioneer in the legal technology industry dedicated to powering global justice through artificial intelligence. Headquartered in Paris, with additional offices in New York, London and Singapore. Jus Mundi serves over 150,000 users from law firms, multinational corporations, governmental bodies, and academic institutions in more than 90 countries. Through its proprietary AI technology, Jus Mundi provides global legal intelligence, data-driven arbitration professional selection, and business development services.

About the American Arbitration Association®

The American Arbitration Association is the largest provider of alternative dispute resolution (ADR) services in the world. Marking its centennial in 2026, the American Arbitration Association has transformed how legal issues are resolved for better since 1926, turning disputes worldwide into opportunities for understanding and progress. A not-for-profit organization, the American Arbitration Association's mission is to deliver ADR services with integrity, transparency, and innovation. For more information, visit www.adr.org.

About the AAA-ICDR Institute™

The AAA-ICDR Institute is the thought leadership and R&D arm of the AAA, focused on enhancing equity, efficiency, and accessibility in dispute resolution. Through initiatives that promote access to justice, court-connected ADR, the use of artificial intelligence in ADR, and fairness in arbitration, the Institute collaborates with courts, academic institutions, nonprofits, and technology partners to develop equitable, data-driven solutions to shape the future of ADR. Learn more at adr.org/aaa-icdr-institute.

