



NextEra
Policy

THOUGHT PAPER

AI in Legal Practice

Transforming the Way Lawyers Work

A Thought Paper on Responsible AI
and the Future of Legal Services

Published by
NextEra Policy Innovation Foundation

August 2026



From the Director's Desk

Artificial intelligence is transforming the legal profession at a pace and scale not witnessed in generations. From knowledge work to client engagement, from research to risk assessment, AI technologies are reshaping how legal services are designed, delivered and experienced. Across jurisdictions, law firms, in-house legal teams, regulators and courts are experimenting, adapting and reimagining workflows while also grappling with questions of trust, accountability, ethics and human judgment. This paper is situated within that moment of transition.

At NextEraPolicy, we believe that responsible and future ready legal practice will depend on how well the profession understands, governs and integrates AI. This paper is intended as a contribution to that understanding.

This publication emerges from a comprehensive internal process that brings together domain expertise, policy insight and practical experience to examine the impact of AI on legal practice. It seeks to provide an independent, informed and balanced view of the opportunities, challenges and choices that lie ahead.

The purpose of this thought paper is twofold- to engage in conversation about the changing architecture of legal work in an AI enabled environment, and to offer a considered perspective that can inform decisions by practitioners, organisations, policymakers and other stakeholders.

The legal services sector stands at the confluence of tradition and transformation. It is a knowledge intensive, trust centric profession whose core mandate is to uphold rights, interpret rules and enable justice. At the same time, it is under pressure to become more efficient, accessible and inclusive. AI presents an opportunity to enhance these goals if adopted with deliberation and guided by strong ethical, regulatory and professional frameworks. The sector's ability to navigate this shift will shape not only its own future but also its contribution to the rule of law and to wider social and economic development.

RESEARCH BASIS

This paper draws on multiple sources of insight and engagement, including:

- Discussions with legal professionals, technologists, policymakers and industry leaders.
- Analysis of global and national policy developments, market trends and regulatory frameworks.
- Review of research studies, white papers, reports and academic literature.
- Internal reflections informed by practical experience across legal and policy domains.

The aim is to present a perspective that is evidence informed, practice aware and forward looking.

Executive Summary

Legal AI is moving from a question of whether lawyers should use artificial intelligence to a more consequential question: where should AI sit inside the chain of legal responsibility? That distinction matters because legal work is not a single activity. It combines information retrieval, pattern recognition, drafting, negotiation, factual assessment, professional judgment, advocacy and, in courts, the exercise of public power. AI can assist many of these functions; it cannot be allowed to blur who is answerable for the result.

Seven propositions for an AI-driven legal profession

- 1** AI will create its greatest near-term value by compressing high-volume cognitive work: finding, sorting, comparing, extracting, summarising and producing first drafts. The strategic benefit is not “automatic lawyering”; it is the reallocation of professional time toward judgment, negotiation, client counselling and accountability.
- 2** The unit of transformation is the legal task, not the legal job. A lawyer’s role should be decomposed into tasks and each task classified by legal consequence, confidentiality, reversibility, evidentiary impact and need for explanation.
- 3** Efficiency and decision quality are compatible only when verification is designed into the workflow. A faster wrong answer is not legal innovation. For legal propositions, the authoritative source must remain identifiable and independently checkable.
- 4** Law firms and corporate legal departments require different AI operating models. Firms must protect client duties across multiple matters and preserve defensible work product; in-house teams must additionally align AI use with enterprise data governance, procurement, compliance, cybersecurity and business decision-making.
- 5** Courts require the strictest boundary. India’s 2026 draft Regulations for Use of Artificial Intelligence in Courts place human primacy, judicial independence, transparency, accountability, auditability and data protection at the centre. The Supreme Court’s July 2026 judgment in *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.* reinforces a zero-tolerance approach to unverified AI-generated precedent. [S2] [S3]
- 6** The emerging international direction is not a blanket prohibition. Official frameworks in the European Union, England and Wales, and the United States increasingly distinguish between low-risk administrative assistance and AI use that can materially affect legal rights or judicial outcomes. [S7] [S8] [S10]
- 7** The durable competitive advantage will be institutional, not merely technological: disciplined source verification, secure data environments, clear accountability, measurable controls, skilled users and governance that can evolve as systems change.



Policy direction

The objective should be “responsible acceleration”; enable low-risk uses quickly, permit substantive assistance under documented human review, and prohibit or tightly control uses that substitute AI for final legal or judicial judgment.



Paper map

The paper proceeds from practice to policy: it first examines how AI changes the operating model of law firms and in-house legal teams; then addresses the judiciary and the boundary between assistance and adjudication; then develops risk, ethics and data-governance principles; and finally proposes a practical institutional framework and future-skills agenda.



LEGAL RESEARCH
& INFORMATION



RESPONSIBILITY
& ACCOUNTABILITY



LAW FIRMS &
IN-HOUSE TEAMS



COURTS &
JUDICIARY



GOVERNANCE
& CONTROLS



AI & TECHNOLOGY
ENABLERS

How AI Changes the Legal Operating Model

The most useful way to understand AI in legal practice is to separate “work production” from “professional responsibility”. Legal work contains large volumes of reading, classification, comparison and drafting. These are cognitively expensive but not always judgment-intensive. AI can reduce the time required to perform them. The lawyer’s professional obligation, however, does not shrink at the same rate. In many situations it increases, because the lawyer must now evaluate not only the underlying legal material but also the reliability, provenance and limitations of the AI-assisted process.

The task becomes the unit of redesign

Traditional legal teams are organised around matters, files and professional seniority. AI invites a second layer of organisation: tasks can be mapped according to whether they are information-seeking, generative, analytical, communicative, evidentiary or decisional. This enables an institution to decide where automation is useful and where it is inappropriate without attempting to produce one rule for an entire profession.

- 1 **Retrieval tasks** can often be AI-assisted if the underlying authority is visible and the lawyer independently confirms the source.
- 2 **Transformative tasks**—summarising, translating, converting a bundle into a chronology—can be accelerated, but the original record must remain available for comparison.
- 3 **Generative tasks**—drafting clauses, arguments, advice notes or correspondence—require explicit review because plausible language can conceal legal or factual error.
- 4 **Predictive or recommendation tasks** require stronger controls where they affect rights, litigation strategy, employment decisions, investigations, sanctions or settlement posture.
- 5 **Decisional tasks** that determine legal rights or exercise judicial authority should remain human-controlled; AI can inform but should not displace the accountable decision-maker.

Why “human in the loop” must mean more than a signature

A nominal human reviewer is not a safeguard if the reviewer lacks the time, expertise or source access needed to challenge the machine output. Effective human oversight therefore requires three conditions: authority to reject the output, competence to evaluate it, and enough visibility into the source material to perform meaningful verification. The 2026 draft Supreme Court regulations capture this idea by defining human-in-the-loop governance as mandatory review, supervision and verification with final authority and accountability resting with a human. [S2]

The strategic opportunity

Properly designed, AI can shift a legal function away from document production as its organising centre and toward decision architecture. Lawyers can spend less time generating the first version of a chronology, clause comparison or legal issue list and more time testing assumptions, identifying missing facts, explaining trade-offs, designing negotiation strategy and documenting why a particular course is legally and commercially defensible. This is not a smaller role for lawyers; it is a more concentrated one.



Practical implication

The correct implementation question is not “Which AI tool should we buy?” It is “Which legal tasks should change, what controls must travel with them, and what evidence will show that quality has not been traded for speed?”



REDESIGN BY TASK, NOT TITLE

Map legal work by task type to decide where AI adds value and where human judgment is essential.



OVERSIGHT WITH REAL AUTHORITY

Effective human oversight requires authority to reject, competence to evaluate and visibility into sources.



RESERVE DECISIONS FOR HUMANS

Decisions that determine rights or exercise judicial authority must remain human-controlled.



MOVE UP THE VALUE CHAIN

Shift from producing documents to designing and defending decisions with clarity, evidence and judgment.

Law Firms: From Production Capacity to Judgment Capacity

Where AI fits naturally

Law firms are information-intensive organisations. Their recurring work includes legal research, document review, due diligence, drafting, discovery, chronology building, contract comparison, knowledge retrieval and matter administration. AI can reduce friction across each of these activities, particularly where the task involves identifying patterns within a bounded document set or producing a draft that is subsequently reviewed by a lawyer.

The most defensible use cases share four features: the input material is defined; the desired output can be checked; the task is reversible; and the lawyer who relies on the output remains identifiable. For example, AI-assisted due diligence can surface unusual clauses for review, but should not silently determine whether a contractual risk is acceptable. AI-assisted research can produce candidate authorities and issue maps, but the lawyer must still read and validate the cited source before it enters advice or a pleading.

The economics of legal work will change before the duty of care does

AI is likely to reduce the amount of junior time required for some forms of first-pass work. That creates a business-model question: clients may become less willing to pay simply for the volume of hours consumed in tasks that are visibly automatable. Firms will therefore need to articulate value in terms of outcome, judgment, risk reduction, responsiveness and domain expertise rather than equating effort with value. At the same time, firms should resist a false efficiency in which an unverified AI draft is treated as a finished product merely because it looks professionally written.

The training paradox

There is a deeper institutional problem. Many lawyers historically learned judgment by performing the very tasks AI can now accelerate: reading full cases, preparing chronologies, drafting from first principles and reviewing large records. If those tasks disappear without a replacement learning model, firms could create a generation of lawyers who are faster at producing outputs but weaker at detecting when an output is wrong. Training therefore has to become more explicit: source comparison, adversarial review, red-teaming of arguments, explanation of why a clause matters, and supervised reconstruction of AI-assisted work.

A law-firm control model

- 1 **Approved systems and matter-level data rules**, including clear restrictions on entering privileged or confidential information into unapproved public tools.
- 2 **Mandatory authority checking** for every legal proposition generated or surfaced by AI before external use.
- 3 **Named human ownership** for each AI-assisted work product, with review intensity tied to consequence.
- 4 **Documented client-position rules**: when AI use should be disclosed, when client consent is appropriate, and how billing reflects automated assistance.
- 5 **Training** that treats AI competence as a combination of legal knowledge, system literacy, verification discipline and security awareness.



Firm-level principle

AI may compress production time, but it should not compress professional skepticism. The faster a draft appears, the more deliberate the verification stage must become.



NATURAL AI FIT

AI is well suited to defined, document-heavy work where patterns can be found and drafts reviewed by lawyers.



VALUE BEYOND HOURS

Clients buy outcomes, judgment and risk reduction—not the volume of time spent on automatable tasks.



TRAINING PARADOX

Removing traditional learning tasks without new models risks speeding outputs while weakening judgment.



CONTROL MODEL

Strong governance—approved systems, verification, human ownership and client rules—protects quality, trust and the profession.

Corporate Legal Departments: AI as an Enterprise Control Function

In-house legal teams sit at a different intersection from law firms. They are not only legal advisers; they are part of the enterprise operating system. Their AI decisions interact directly with procurement, information security, privacy, HR, finance, compliance, internal audit, product development and business strategy. As a result, the in-house legal function can become both a major user of AI and a governance hub for how the wider organisation adopts it.

High-value use cases

The strongest use cases are often repeatable workflows: contract intake and triage, clause extraction, first-pass review against playbooks, policy comparison, regulatory mapping, obligation tracking, legal request routing, investigation document organisation, board-material summarisation and generation of internal FAQs from approved sources. AI can also improve the consistency of issue spotting by forcing a structured first pass across a large set of agreements or communications before a lawyer applies context and commercial judgment.

The privilege and confidentiality problem is architectural

For in-house teams, the central risk is not merely whether a lawyer remembers to remove a client name from a prompt. Enterprise AI systems can interact with document repositories, email stores, HR systems and matter-management platforms. Once AI is connected to those environments, permission design, retention rules, logging, model training settings, vendor access and cross-border processing become part of legal-risk management. Privilege can be undermined not only by disclosure to an external system but also by over-broad internal access or poor separation between business and legal datasets.

Legal should govern by use case, not by tool brand

A tool-by-tool approval model becomes difficult to maintain because AI functionality is increasingly embedded inside ordinary productivity software, search, document management and workflow platforms. A more durable model is to approve classes of use. The legal department can define the permitted purpose, data category, reviewer requirement, external-use rule and evidence of verification for each class. This creates governance that survives product changes.

The in-house lawyer's emerging role

The in-house lawyer will increasingly act as a translator between legal consequence and technical design. That requires the ability to ask practical questions of a system: What data can it access? Is customer or employee data retained? Can prompts or outputs be used for model improvement? Is there an audit log? Can a user reproduce the source behind a conclusion? What happens when the model changes? Who is responsible for incident response? These are legal questions expressed through system architecture.

- 1 Build an AI-use register that records approved legal use cases, accountable owners, data classes and review obligations.
- 2 Require contract clauses for AI providers that address confidentiality, data use, security, change notification, audit rights and incident management.
- 3 Separate internal knowledge generation from external legal advice: a system may propose a draft, but accountable counsel should authorise the legal position.
- 4 Use legal operations metrics that measure quality and exception rates, not merely speed or volume.



SYSTEM-WIDE REACH

AI decisions interact across procurement, security, privacy, HR, finance, compliance, audit and business strategy.



PRIVILEGE AT RISK

Enterprise AI access to data environments makes permission design, retention, logging and vendor controls central to legal risk.



USE CASE GOVERNANCE

Approve classes of use with defined purpose, data category, reviewers and verification rules rather than individual tool approvals.



LEGAL AS ARCHITECT

In-house lawyers translate legal consequence into system design and govern questions of data, audit, change and accountability.

The Judiciary: India's Shift from Digitisation to Governed AI

India's judicial AI story is best understood as an extension of a longer digital transformation rather than a sudden move toward automated judging. The Supreme Court's November 2025 White Paper on Artificial Intelligence and Judiciary places AI within the broader e-Courts trajectory and identifies tools for research, translation, transcription, filing and administrative support. It also devotes substantial attention to hallucination, confidentiality, bias, verification and human oversight. [S1]

Officially identified judicial AI tools

The Supreme Court White Paper describes a progression from digital legal information toward more capable AI assistance, including SUPACE, SUVAS, AI-based transcription, Legal Research Analysis Assistant (LegRAA), and AI-assisted e-filing. A March 2026 Rajya Sabha answer from the Department of Justice confirms that LegRAA has been developed to assist judges with legal research and document analysis, while Digital Courts 2.1 is designed as a single-window platform for case information and tasks, including AI-supported transcription and translation functions. The same answer emphasises controlled pilots, internal use and data safeguards. [S1] [S4]

This architecture matters because it shows where public institutions currently see the strongest case for AI: not autonomous adjudication, but friction reduction around the judicial process. Translation, transcription, document access, search, filing scrutiny and case organisation are examples of functions where technology can improve reach and workflow while leaving the legal determination with the judge.

The White Paper's normative direction

The White Paper's recommendations emphasise human-in-the-loop review, accuracy and verification, confidentiality and privacy, fairness, specialised verification of translated text, and restricted deployment in routine administrative functions. Its suggested guidance for lawyers also warns against placing private, confidential or legally privileged information into unsuitable AI tools, treating large language models as substitutes for reliable legal research sources, or relying on one generative AI system to verify another. [S1]

From guidance to an institutional governance model

In June 2026, the Supreme Court published draft Regulations for Use of Artificial Intelligence in Courts. The draft is significant because it moves beyond general ethical advice toward institutional mechanisms: an Apex Body, judicial and technical committees, controlled-environment testing, an AI register, audits, an incident database, transparency requirements, an AI Content Verification Authority, procurement rules, data protection, cybersecurity, training and grievance redressal. [S2]



India's emerging model

The direction is neither "AI-free courts" nor "automated courts". It is governed assistance: encourage useful deployment, but make human primacy, verification, auditability, data protection and institutional accountability explicit.



AI AS JUDICIAL ASSISTANCE

AI supports research, translation, transcription, filing and administrative functions — the judge retains the final legal determination.



KEY TOOLS IDENTIFIED

SUPACE, SUVAS, LegRAA and AI-assisted e-filing under the e-Courts roadmap, with controlled pilots and data safeguards.



HUMAN-IN-THE-LOOP

Accuracy, verification, privacy, fairness and human oversight at the core of judicial AI deployment.



INSTITUTIONAL GOVERNANCE

Draft Regulations (June 2026) propose an Apex Body, AI register, audits, content verification authority and strong data protection.

Where Assistance Must Stop: AI and Adjudicatory Authority

The boundary between AI assistance and AI decision-making is most consequential in courts because adjudication exercises public authority over rights, obligations, property and, in some cases, liberty. A court can use technology to organise information without surrendering judgment; difficulty arises when a system begins to shape the reasons for a decision in a way that cannot be independently tested or meaningfully challenged.

The Supreme Court's 2026 draft regulations

The draft Regulations for Use of Artificial Intelligence in Courts, 2026 state that AI in court processes must remain strictly subservient to human judgment and judicial authority, that AI systems should operate solely in an assistive capacity, and that ultimate authority over matters of law, fact and justice remains with judicial officers. The draft also links higher-risk applications to heightened safeguards, including mandatory human-in-the-loop requirements and independent oversight. [S2]

Equally important, the draft rejects “the AI made me do it” as an accountability defence. Where an officer uses AI assistance, responsibility remains with the officer; opacity or hallucination cannot be invoked to avoid accountability. It also proposes technical and ethical impact assessments, audit requirements, disclosure rules, incident management and verification mechanisms. [S2]

Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.

The Supreme Court's judgment of 2 July 2026 provides a concrete illustration of why verification is not a procedural nicety. In *Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd.*, the Court dealt with tribunal decisions that relied on fake or hallucinated material represented as precedent. The Supreme Court set aside the decisions and declared a zero-tolerance approach toward producing, citing or using AI-generated precedents without verification.

It emphasised “total and absolute control over adjudication, with a human in the loop at every stage” and directed the Bar Council of India to consider guiding principles and disciplinary consequences for advocates who submit fake or hallucinated precedent. [S3]

A functional boundary

-  AI may assist with finding and organising authorities; the court or lawyer must verify the authority itself.
-  AI may help structure a record; it should not *decide* which disputed fact is true.
-  AI may generate a draft summary or outline; it should not become the unexamined source of judicial reasons.
-  AI may support translation or transcription; the system must allow correction and specialised verification where meaning is legally material.
-  AI may assist scheduling and administrative workflow; it should not autonomously determine substantive priority where the decision has rights-affecting consequences without a transparent and reviewable policy basis.

Constitutional intuition

The more closely an AI system approaches the application of law to contested facts, the less acceptable opacity and automation become. Efficiency cannot substitute for reasons, and reasons cannot be meaningful if the accountable decision-maker cannot defend them independently of the machine.



JUDICIAL PRIMACY

AI must remain strictly subservient to human judgment and judicial authority.



ACCOUNTABILITY REMAINS HUMAN

“The AI made me do it” is not a defence. Responsibility rests with the officer who uses AI.



VERIFICATION IS NON-NEGOTIABLE

All AI-generated content, including precedents, must be independently verified before use.



HUMAN-IN-THE-LOOP AT EVERY STAGE

High-risk applications require mandatory human oversight and independent review.



Decision Quality: Verification by Design

AI can improve decision-making by widening the field of information considered, surfacing inconsistencies, generating alternative formulations and reducing the cost of first-pass analysis. But decision quality improves only if the system is designed to expose uncertainty rather than conceal it. Generative AI is particularly persuasive because it produces fluent answers; fluency can create an illusion of authority. Legal workflows therefore need verification as an architectural step, not as an optional instruction to “double-check”.

A four-stage verification chain

Stage	Question	Minimum control
1	Source Can the proposition be traced to an authoritative primary source?	Do not rely on generated citations or quotations without opening the source.
2	Context Does the source actually support the proposition in its procedural and factual setting?	Read the relevant provision, paragraph and surrounding context.
3	Currency Is the authority still in force and applicable to the jurisdiction and date?	Check amendments, later judgments, commencement and territorial scope.
4	Ownership Which human professional accepts responsibility for the final use?	Record the reviewer for high-consequence work and preserve the verification path.



Verification must be independent

An important practical rule follows: one generative model should not be treated as independent verification of another model's answer. If both systems draw from similar training corpora or reproduce the same fabricated citation, cross-checking can create false confidence. The proper verification source is the authoritative legal text, official record, authenticated filing, or other reliable primary material. The Supreme Court White Paper expressly cautions against using one generative tool to authenticate another and requires independent verification before reliance. [S1]



When AI can improve legal reasoning

Used carefully, AI can be valuable as an adversarial assistant. A lawyer can ask it to identify counterarguments, missing factual assumptions, alternative clause interpretations or questions that a regulator, judge or opposing counsel might raise. The lawyer should then test those prompts against actual law and evidence. This use is different from asking the model to “give the answer”. It uses AI to widen the challenge set while keeping the conclusion grounded in accountable professional reasoning.



Decision records will become more important

For high-impact matters, institutions should preserve a concise record of how AI contributed: the approved tool or use case, the relevant source set, the human reviewer, material corrections, and any unresolved limitations. This need not become bureaucratic documentation for every low-risk task. The point is proportionality: the higher the consequence, the stronger the provenance trail.



EXPOSE, DON'T CONCEAL

Design workflows to expose uncertainty and limits, not to present fluent answers as final truth.



VERIFY AT THE SOURCE

Primary law, official records and authenticated documents are the only proper bases for verification.



HUMAN ACCOUNTABILITY

Technology assists; professionals decide. Ownership must be explicit and documented.



PROPORTIONALITY

Stronger verification for higher impact. Lower burden for low-risk tasks.



Ethics and Professional Responsibility

AI does not erase established professional duties. It creates new ways to breach them. Competence, confidentiality, candour, supervision, independence and responsibility for filed or advised work remain recognisable legal-professional principles. The challenge is translating them into operational rules for systems that generate plausible content, process large data volumes and can be embedded invisibly inside routine software.



Competence becomes socio-technical

A competent AI-using lawyer does not need to be a machine-learning engineer, but must understand enough about the system to know when it is unsuitable. That includes recognising hallucination risk, the difference between retrieval and generation, the limits of context windows and source grounding, the possibility of prompt injection or malicious documents, and the data consequences of using public versus enterprise-controlled tools. Technical literacy is now part of exercising legal judgment safely.



Confidentiality and privilege

Before confidential material enters an AI system, a lawyer should know who can access it, whether it is retained, whether it may be used for system improvement, where it is processed, how long logs persist, and whether the organisation can delete or export records. The Supreme Court White Paper and the Judiciary of England and Wales guidance both emphasise caution with confidential or private information in public AI systems. [S1] [S8]



Candour, evidence and authenticity

AI use raises special concerns when the output can alter evidence or create the appearance of personal recollection. Drafting a pleading or skeleton argument is different from drafting a witness statement. The England and Wales Civil Justice Council's 2026 work reflects this distinction: it indicates substantial confidence in existing professional frameworks for ordinary legal drafting while treating witness evidence as a distinct area where AI could reshape, embellish or influence the witness's account. [S9]



Supervision and accountability

An organisation should apply the same basic principle to AI that it applies to junior lawyers and external providers: delegation does not eliminate supervision. The human lawyer who approves a document, advice or submission should be able to explain the legal basis for it. In the judicial context, both India's draft regulations and the U.S. federal judiciary's interim guidance expressly preserve human accountability and independent verification. [S2] [S10]



Disclosure should be material, not theatrical

Blanket declarations that "AI was used somewhere" can generate noise without improving trust. Disclosure should instead be tied to materiality: whether AI generated content filed as evidence, materially assisted a rights-affecting process, created synthetic information, or played a role that a reasonable court, client, regulator or counterparty would need to understand. India's draft court regulations take a comparatively strong approach by proposing disclosure of AI-assisted court submissions and the steps taken to verify AI-generated content. [S2]



Ethical rule of thumb

A lawyer should never rely on AI in a way that would make the lawyer unable to defend the work if the AI system, prompt history or vendor disappeared tomorrow.



CORE DUTIES REMAIN

Competence, confidentiality, candour, supervision, independence and responsibility still apply—AI changes how they are exercised.



PROTECT CONFIDENTIALITY

Understand data handling, retention, access, location, logs and deletion before sharing any information with AI systems.



TRUTH AND AUTHENTICITY

AI can influence narratives. Differentiate legal drafting from witness testimony. Preserve truth, accuracy and authenticity.



HUMAN ACCOUNTABILITY

Supervision is non-negotiable. The approving lawyer is responsible and must be able to justify the legal basis.



DISCLOSE WHAT MATTERS

Disclose AI use only when it is material to the court, client, regulator or counterparty.



Data Governance, Privacy and Cybersecurity

Legal AI is often discussed as a model problem when it is equally a data-governance problem. A system's usefulness grows when it can access contracts, advice, pleadings, evidence, emails, policies and internal knowledge. Those same connections can increase the consequences of error, leakage, over-broad access and unauthorised reuse. The legal function therefore needs a data architecture for AI, not simply a usage policy.



India's privacy framework changes the baseline

The Digital Personal Data Protection Act, 2023 creates a statutory framework for processing digital personal data, while the Digital Personal Data Protection Rules, 2025 and commencement notifications establish a staged implementation path from November 2025. For AI governance, the operational lesson is straightforward: legal teams should be able to identify why personal data is being processed, what data is necessary, who can access it, how security is maintained, and when it should be deleted or no longer used. [S5] [S6]



India's 2026 draft court AI regulations explicitly incorporate purpose limitation, data minimisation, privacy by design, sensitive judicial data protection and cybersecurity. That is a useful model beyond courts because it connects abstract principles to the AI lifecycle: system design, procurement, deployment and audit. [S2]

A minimum control stack for legal AI



Data classification: distinguish public material, internal business information, personal data, confidential legal material, privileged communications, evidence and highly sensitive categories.



Purpose control: authorise the system for a defined use case rather than granting unrestricted access because the technology is useful.



Least-privilege access: an AI assistant should not be able to retrieve every matter simply because a user can ask a natural-language question.



Retention and training controls: know whether prompts, documents or outputs are retained or used to improve models; contractually prevent uses inconsistent with legal obligations.



Logging and audit: preserve enough technical and user activity information to investigate an incident without creating unnecessary repositories of sensitive prompts.



Change management: require notice and re-assessment when a provider materially changes the model, data handling, hosting, features or connected sources.



Fallback and incident response: define what happens when a system is unavailable, produces a serious error, leaks data or generates content that has already entered a legal process.



Cybersecurity is also an integrity issue

Legal AI creates new attack surfaces. Malicious or manipulated source documents can influence AI output; a compromised knowledge base can propagate incorrect information quickly; and over-permissioned agents can take actions beyond what a lawyer intended. Security review must therefore include not only confidentiality but integrity: can the organisation trust that the source set, model behaviour and connected systems have not been manipulated?



Governance principle

Do not give an AI system broader legal access than the human user needs for the approved task. Convenience is not a sufficient reason to collapse matter boundaries, privilege boundaries or data classifications.



LEGAL AI IS A DATA PROBLEM

The more data AI can access, the more valuable—and the greater the risk.



PRIVACY BY DESIGN

Purpose limitation, data minimisation and transparency are the foundation of trusted AI use.



CONTROL ACCESS

Limit access to what is necessary. Least privilege reduces exposure and misuse.



SECURE THE SYSTEM

Protect confidentiality and integrity through strong cybersecurity and continuous monitoring.



PREPARE AND RESPOND

Define fallback plans and incident response before issues occur—not after.

Comparative Official Direction: A Convergence Around Human Accountability

Although legal systems differ, current official materials show a striking convergence. The emerging policy question is not whether AI can be used in law; it is how to distinguish acceptable assistance from uses that create unacceptable risks to rights, fair process, confidentiality or institutional legitimacy.



European Union: risk classification and human oversight

The EU AI Act classifies certain systems used by or on behalf of judicial authorities to assist in researching and interpreting facts and law and applying law to a concrete set of facts as high-risk. Recital 61 expressly states that AI may support judicial decision-making but should not replace it and that final decision-making must remain human-driven. The Act then connects high-risk systems to lifecycle risk management, transparency, documentation and other obligations. [S7]



England and Wales: personal responsibility and secure use

The Judiciary of England and Wales updated its AI guidance in October 2025. The guidance highlights hallucination, bias and confidentiality, cautions against putting private information into public AI tools, and reinforces the personal responsibility of judicial office holders for material produced in their name. [S8]

The Civil Justice Council's 2026 work adds a useful distinction between categories of legal documents. Its June 2026 update indicates strong support for relying on existing professional responsibility frameworks for ordinary pleadings and advocacy documents, while continuing to examine additional safeguards for witness statements and expert evidence. This points toward contextual rather than blanket regulation. [S9]



United States federal judiciary: experimentation without delegation

The Administrative Office of the U.S. Courts established an AI Task Force in 2025 and developed interim guidance. The official annual report states that the guidance cautions against delegating core judicial functions, including decision-making and adjudication, recommends independent verification of AI-generated content, and emphasises accountability for work performed with AI assistance. [S10]



NIST: governance as a continuous risk process

NIST's AI Risk Management Framework provides a useful non-sectoral governance vocabulary—Govern, Map, Measure and Manage—while its Generative AI Profile focuses the framework on generative-system risks. For legal departments, its value is methodological: AI risk should be continuously identified, measured and managed across the lifecycle rather than treated as a one-time procurement questionnaire. [S11]

The common pattern



Human decision authority remains central where legal rights are affected.

Verification is a distinct duty, not a by-product of using a sophisticated system.

Transparency and explainability become more important as consequence increases.

Data protection and secure environments are prerequisites for sensitive legal use.

Risk management must continue after deployment through logging, audit, incident response and system review.

Low-risk administrative functions can be treated more permissively than substantive rights-affecting uses.



Comparative conclusion

The global regulatory trend supports differentiated governance. Legal institutions do not need a single "AI rule"; they need a risk ladder that becomes stricter as AI moves from administration to substantive legal reasoning and from advice to adjudication.



HUMAN DECISION AUTHORITY

AI supports; it does not decide.
Final authority for legal rights remains human.



VERIFICATION FIRST

Independent verification is a distinct duty, not an optional follow-up.



TRANSPARENCY MATTERS

As consequences rise, so does the need for transparency and explainability.



PROTECT DATA

Secure environments and strong data protection are the foundation for trustworthy legal AI use.



MANAGE CONTINUOUSLY

Risk management does not end at deployment—monitor, audit, respond and review.



RISK-BASED DIFFERENTIATION

Stricter rules for high-risk, rights-affecting uses; more flexibility for low-risk administrative tasks.

A Proposed Framework for Responsible Legal AI

The following framework is proposed as an institutional starting point. It is intentionally simple enough to become policy later without tying the organisation to a particular vendor, model or technical architecture.

A three-zone model for legal AI use

Zone	Typical use	Examples	Control level
GREEN Administrative / low consequence	Routine, non-substantive legal support tasks	Formatting, routing, scheduling, public-information summarisation, internal search over approved non-sensitive material	Approved tool; normal supervision; periodic sampling
AMBER Substantive assistance	Analytical or drafting support within legal responsibility	Legal research leads, contract drafts, due diligence, chronology generation, analysis of bounded records, internal legal summaries	Named human owner; primary-source verification; data controls; material-use log
RED Decision-substituting / integrity-sensitive	High-consequence or integrity-critical activities	Autonomous final legal advice, unverified precedent, AI-authored witness recollection, autonomous adjudication, unsupervised use of privileged or highly sensitive data	Prohibited or exceptional approval with heightened independent safeguards

Six principles for responsible legal AI

- Human accountability.** Every substantive AI-assisted legal output must have an identifiable human owner with authority and competence to reject it.
- Verified authority.** No generated legal proposition, quotation, citation or factual assertion should be relied upon externally without checking the authoritative source or record.
- Confidentiality by design.** Sensitive legal data should be processed only in environments whose access, retention, training, security and contractual terms are understood and approved.
- Risk-proportionate controls.** Review, disclosure, audit and approval requirements should scale with the consequence and reversibility of the task.
- Traceability and transparency.** High-consequence uses should leave a defensible record of the system used, the source basis, material human review and any required disclosure.
- Continuous governance.** Approval is not permanent. Models, features, integrations and legal rules change; therefore testing, training, incident review and re-approval must continue through the lifecycle.

Institutional ownership

A workable governance model should combine legal, technology, privacy, cybersecurity and business ownership. Legal determines the consequence of the use case; technology and security assess architecture and controls; privacy addresses data processing; business owners confirm operational need; and an AI governance function maintains the register, approval record, training requirements and incident process. Courts require an analogous institutional structure with judicial independence protected at its core—a direction already visible in India's 2026 draft regulations. [S2]



Design principle

Policy should make the safe path the easy path. If approved tools, source-grounded workflows and secure environments are cumbersome, professionals will route around them. Good governance is therefore partly a product-design problem.



LEGAL AI IS A PARTNER, NOT A REPLACEMENT

AI assists; humans decide.
Accountability never transfers.



VERIFICATION FIRST

Every legal proposition depends on authoritative sources and human review.



HUMAN AT THE CENTRE

Human judgment, professional ethics and legal responsibility remain non-negotiable.



SECURE BY DESIGN

Confidentiality, integrity and availability are built in from the start.



GOVERN FOR IMPACT

Governance, audit and learning ensure safe, effective and future-ready AI use.

Future Skills, Implementation Roadmap and Policy Conclusions

The lawyer of the AI era

The future lawyer will need deeper judgment, not less. As drafting and search become easier, differentiation will lie in framing the right question, identifying material facts, testing sources, recognising uncertainty, negotiating under ambiguity and explaining a defensible recommendation to a client, board, regulator or court.



Legal source literacy: moving from an AI proposition back to the statute, rule, judgment or official record.



Verification and adversarial review: testing for missing authority, contradictory facts, false confidence and unsupported inference.



Data and system literacy: understanding permissions, retention, audit logs, integrations and model change.



Workflow design: deciding what to automate, what to review and where proof of quality must be captured.



Client counselling: explaining assumptions, trade-offs and residual uncertainty behind AI-assisted analysis.



Governance leadership: translating professional duties into procurement, privacy, cybersecurity and audit controls.

A practical implementation sequence

Begin with tasks rather than tools. Classify recurring workflows under the Green–Amber–Red model; establish approved systems with clear data rules and named owners; create verification templates for high-consequence outputs; train users on realistic failure modes; maintain an incident-learning loop; and re-assess use cases whenever systems gain new capabilities or access to new data.

Policy recommendations for institutional release

- ✓ Adopt a written principle that AI augments but does not displace accountable legal judgment.
- ✓ Classify uses by consequence, enabling low-risk adoption while heightening safeguards for rights-affecting work.
- ✓ Require primary-source verification for legal authority generated or surfaced by AI before external reliance.
- ✓ Keep privileged, confidential and sensitive legal information out of unapproved AI environments.
- ✓ Require material disclosure where provenance affects evidence, adjudication, fairness or trust.
- ✓ Create multidisciplinary governance for approvals, audits, incidents, training and system changes.
- ✓ Make verification, source literacy, data governance and critical reasoning core professional-development skills.

Conclusion

AI will transform legal practice because it changes how cognitive work is allocated across an institution. That capacity should be welcomed where it removes friction and expands access to information. But legal legitimacy still depends on responsibility; courts remain responsible for judgments, lawyers for advice and submissions, and institutions for the systems and data architectures they deploy.

The policy objective is not to preserve every traditional workflow. It is to preserve accuracy, confidentiality, independence, fairness, reasoned decision-making and accountability while redesigning work around better tools. That is the path from experimental AI use to an AI-enabled legal profession worthy of institutional trust.



STRONGER JUDGMENT

AI handles the heavy lifting.
You bring the judgment that matters.



BUILD VERIFICATION HABITS

Good lawyers verify,
Great lawyers verify and challenge.



GOVERN DATA, PROTECT TRUST

Data discipline is the
foundation of responsible
AI use.



DESIGN BETTER WORKFLOWS

Automate the routine.
Review the critical.
Prove the quality.



LEAD WITH RESPONSIBILITY

Translate professional duties
into systems, training and
governance that last.



SHAPE THE FUTURE

Build an AI-enabled legal
profession worthy of
institutional trust.

Primary Public Sources and Research Notes

The paper deliberately excludes commercial market reports, proprietary legal-industry surveys, vendor white papers and competitor-sponsored statistics. The following primary public materials form the principal research base. Short-form identifiers [S1]–[S11] are used in the paper.



[S1] Supreme Court of India, Centre for Research and Planning, White Paper on Artificial Intelligence and Judiciary (November 2025). Official Supreme Court report. **Key topics:** Indian judicial AI tools; hallucinations; confidentiality; human-in-the-loop review; guidelines for courts, lawyers and law clerks.



[S2] Supreme Court of India, Draft Regulations for Use of Artificial Intelligence in Courts, 2026, published for stakeholder comments on 3 June 2026 (deadline subsequently extended). **Key topics:** human primacy; rule of law; fairness; transparency; accountability; auditability; purpose limitation; proportionality; AI registers; incident management; disclosure; data protection; cybersecurity; training.



[S3] Pooja Ramesh Singh v. Jammu and Kashmir Bank Ltd. & Anr., Civil Appeal No. 11950 of 2025, 2026 INSC 668, judgment dated 2 July 2026, Supreme Court of India. **Key topics:** fabricated or hallucinated precedents; zero tolerance; human control over adjudication; accountability and Bar Council directions.



[S4] Government of India, Ministry of Law and Justice, Department of Justice, Rajya Sabha Unstarred Question No. 3155, answered 19 March 2026, “Use of AI in Legal Framework”. **Key topics:** eCourts Phase III; LegRAA; Digital Courts 2.1; controlled pilots; privacy, security and ethical safeguards.



[S5] The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023), India Code, Ministry of Law and Justice / Legislative Department. **Key topics:** lawful processing of digital personal data; obligations and rights; statutory data-protection framework.



[S6] Government of India, Ministry of Electronics and Information Technology, Digital Personal Data Protection Rules, 2025, published 14 November 2025, together with commencement and implementation materials. **Key topics:** operational data-protection requirements and phased implementation.



[S7] Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), Official Journal / EUR-Lex. **Key topics:** high-risk AI; administration of justice; risk management; transparency; human oversight.



[S8] Courts and Tribunals Judiciary, England and Wales, Artificial Intelligence (AI) – Judicial Guidance, updated 31 October 2025. **Key topics:** hallucinations; bias; confidentiality; personal judicial responsibility; secure use.



[S9] Civil Justice Council, England and Wales, Use of AI for Preparing Court Documents: Interim Report and Consultation (2026) and Update on Consultation Findings (30 June 2026). **Key topics:** professional legal drafting; witness statements; expert evidence; authenticity and disclosure.



[S10] Administrative Office of the United States Courts, Annual Report 2025, Court Operations – “Developing Artificial Intelligence Policies”. **Key topics:** AI Task Force; no delegation of core judicial functions; independent verification; accountability and disclosure considerations.



[S11] U.S. National Institute of Standards and Technology (NIST), Artificial Intelligence Risk Management Framework (AI RMF 1.0) (2023) and Generative Artificial Intelligence Profile (2024, updated 2026). **Key topics:** lifecycle risk governance; mapping, measuring, managing and governing AI risk.



Research note on currency

This paper reflects public materials located and reviewed through 7 August 2026. The 2026 Supreme Court Regulations cited above are draft regulations published for consultation and should not be described as final binding rules unless and until formally brought into force. The paper treats them as evidence of emerging institutional policy direction, not as a statement of currently operative law. Similarly, the DPDP framework has staged commencement; organisations should verify which provisions are in force on the date of any specific implementation decision.



Suggested citation: “AI in Legal Practice: Transforming the Way Lawyers Work — A Thought Paper on Responsible AI and the Future of Legal Services” (Institutional Draft, 7 August 2026).