Comments on the Draft 'Regulations for Use of Artificial Intelligence (AI) in Courts, 2026'

Overview

Before our detailed comments, we wish to outline seven overarching observations:

First, the draft Regulations leave ***significant terms undefined*** or deferred, such as "*high-risk*" (Clause 7(3), 12(2) and 46(4)(h)), "*demonstrably biased*" data (Clause 14(2)), "*responsible AI adoption*" (Clause 16), and the disclosure threshold of "*materially assists*" (Clause 43(1)), even though each carry distinct obligations or consequences.

Second, the ***institutional architecture*** under Chapter IV separates key mandates, but could strain judicial capacity leading to delays in implementing these regulations. Implementation is contingent on establishing a cascade of bodies from the Apex Body to its five Committees, CoRE-AI, and the Supreme Court and every High Court's AI Committee and Secretariat [59 in total, assuming that tribunals and commissions do not have their own separate AI Committees] before the provisions can function in practice. We recommend adopting a phased-commencement schedule within the draft Regulations themselves, under which Chapter II's principles and Clause 20's prohibitions may take immediate effect, additional formats for reporting and impact assessments be done within six months and a further extended timeline for establishing all the bodies (especially where dedicated skill-set based recruitment must be done i.e. the AI Secretariats).

Third, several of the absolute prohibitions listed under Clause 20 are diluted through grant of "*prior approval of the Appropriate Authority*" or allowing the "*officer responsible and accountable… [to] for reasons to be recorded in writing, dispense with the requirement of verification*" under Clause 8(3), creating possible exceptions or conditional permissions. We recommend the draft Regulations retain certain unambiguous prohibited use-cases to guide remedial measures under Clause 21 and to protect judicial independence.

Fourth, the recommended timelines for verification (Clause 8(3)), audit (Clause 38), the compliance-review timeline for legacy systems (Clause 41), and individual officer discretion (Clause 40) are currently uniform regardless of an AI tool's actual risk profile. We would recommend that a ***tiered risk classification*** be created to operationalise Clause 12 (Proportionality) and Clause 18 (Appropriate Authority to identify use of AI in Courts), which would allow oversight actions to be specified as per risk levels.

Fifth, the draft Regulations take an ***undifferentiated treatment of various categories of AI systems.*** While several categories are defined such as "AI System" or "AI Tool" (Clause 3(i)), "Algorithmic Decision-Making" (Clause 3(j)), "Machine Learning" (Clause 3(ze)), "Large Language Model" (Clause 3(zd)), and "Generative Artificial Intelligence" (Clause 3(y)), the draft Regulations do not specify how each category would be treated. A deterministic, rules-based scheduling tool carries a different risk profile than a generative, probabilistic large language model used for judgment drafting. We recommend Clause 19's permissible-use list to detail the class or classes of system(s) that would be relevant for each use case, subject to modification by AI Committees.[1]

[1] Along with the class of AI systems, the regulations (or the Apex body under its powers) may wish to highlight the risks of dependence on global foundational models. Given the recent US government

Sixth, the procurement framework at Clause 46 provides a solid foundation but may be strengthened through further mandatory contractual provisions. These may include data-deletion certification, a subprocessor register, surviving confidentiality obligations, and vendor disclosure of ownership and funding. Additionally, the procurement section currently only covers AI products; however, supplementary service contracts such as confidentiality-protected third-party audits may be mentioned as well, since they may require external organisations (or consortiums) to be engaged.

Seventh, the draft Regulations seem to apply to a range of actors both within the court ecosystem (judicial officers, registry officials, etc) and externally (lawyers or litigants) and how their use of AI could impact judicial processes. For example, if the prohibition in clause 20(f) on profiling is applicable to lawyers or litigants as well - then the procedure for enforcement of such a clause would need to be well-defined. Explicitly establishing said scope is necessary to establish the intent and objectives of the Regulations.

directive to suspend access to new Anthropic models (Source), similar regulatory actions should not interrupt the “primary operational systems or networks of [the] Court”.

Part I: Comments on Institutional Structure

This Part examines the network of various policy-making bodies established under Chapter IV along with the AI Content Verification Authority constituted under Clause 44 of Chapter V.

Apex Body and its functions and powers

The permanent Apex Body at the Supreme Court under Clause 22 is calibrated to serve an overarching national governance role. Clause 3(l) defines "Appropriate Authority" as either the Apex Body or the relevant AI Committee, "*as the case may be*," depending on which body exercises administrative control over a given AI System. At the same time, Clauses 23(g) and 24 give the Apex Body a second-order role, hearing and recommending action on grievances and setting the standards for compliance. Where the Apex Body is the Appropriate Authority for a nationally administered tool, it would simultaneously serve as the first-instance approver of that deployment and the body reviewing grievances arising from it. We **recommend the removal of the Apex Body's standing as an Appropriate Authority for nationally administered AI systems** and retain it as a standard-setting, coordination and grievance-recommending body.

Committees of the Apex Body and the Centre of Research and Excellence on Artificial Intelligence

Additional committees/bodies to be constituted under the Apex Body include:

- the Judicial Committee and its Sub-Committee of Rules (Clause 26),
- the Technical Committee (Clause 27),
- the Committee on Infrastructure and Finance (Clause 28),
- the Case and Data Management Committee (Clause 29), and
- the Cyber Security Committee (Clause 30)

Along with CoRE-AI under Clause 32, the AI Committees under Clause 33, the AI Secretariats under Clause 34, and the AI Content Verification Authority under Clause 44, this establishes a complex oversight network.

This sheer number of bodies raises a concern over the availability and bandwidth of judges, who may be simultaneously serving on other committees under each High Court.[2] On CoRE-AI (Clause 32), we highlight a duplication between its evaluation function of AI tools and prototypes under Clause 32(3)(b) and the "Technical and Ethical Impact Assessment" to be done prior to deployment and submitted to the AI committees under Clause 33(3)(a). If these are to be undertaken at different stages of co-development, experimentation/testing, and approval or for different purposes, this may be specified.

AI Committees and AI Secretariat

Clause 33(5) requires AI Committees at the Supreme Court and all High Courts to meet at intervals of not more than three months, with membership drawn from sitting Judges under

[2] A 2019 analysis observed that the number of committee posts at a high court may be greater than the number of judges, implying that judicial time is not only a significant capacity constraint in the courtroom, but also for judicial administration. (Source)

Clause 33(2). This is an additional responsibility layered onto already demanding judicial workload.

Further, there is a tension between empowering each AI Committee to “approve AI Systems for use in Court processes within its jurisdiction, following a Technical and Ethical Impact Assessment” yet placing “establishment of infrastructure related to Artificial Intelligence, financial aspects of AI adoption in Courts including budget allocations, cost-benefit evaluations and the audit of expenditure under these regulations” under the Committee on Infrastructure and Finance and the Apex Body. Even if selection of AI systems is done by high courts, they are dependent on a centrally-run budget allocation mechanism; if they wish to supplement this, they must seek funds from the state governments. This limitation will constrain high courts from exercising their full discretion in developing plans for their specific needs or driving experimentation as they may wish.

AI Content Verification Authority

We address Clause 44 alongside other institutional provisions, though it appears separately. Given that the Technical Committee under the Apex body is already tasked with monitoring advances in AI research and assessing emerging technological risk, **we recommend that Clause 44’s mandate (continuous updating of verification standards for generative-AI content in court process) be subsumed under the Technical Committee’s functions.** Under the present language, each "Appropriate Authority" shall constitute a separate AI Content Verification Authority, needlessly creating extraneous institutional bodies.

Part II: Points of Clarification

This Part proceeds sequentially through chapters of the draft Regulations and raises points which may be inconsistently interpreted or require further explanation.

A. Definitions (Clause 3)

"Administrative function" – Clause 3(b)

The definition lists examples from case management such as filing, scheduling, notices, records, certified copies but is unclear on whether these regulations would apply to the use of AI within the Court's human resource management (postings, transfers and leave management), finance, or infrastructure monitoring. Clause 3(q) defining "court process" covers "any work, function, or activity relating to the adjudicatory, quasi-adjudicatory, or administrative functions of a Court" but also does not mention these 'back-end' functions.

We recommend a clarification that these regulations apply only to AI used in connection with case management-related operations and thereby exclude AI used in the Court's internal administration.

"AI Service Provider" – Clause 3(h)

This definition requires a contractual relationship ("by virtue of a contract"), which risks excluding informal pilots and unsanctioned or "shadow," AI use from the accountability chain the definition is meant to establish. As per a 2024 UNESCO Global Judges' Initiative survey, 44% of judicial operators surveyed globally had already used AI tools for work-related tasks, frequently without institutional sanction; the same dynamic may be observed in Indian courts.

We recommend Clause 19(2)'s power to specifically approve 'permissible' use-cases be read together with a clarification that any unauthorised or undisclosed use of AI in connection with a court process shall be considered a breach of service rules.

This could form an initial basis for the development of capacity building materials (Clause 49(3)), along with other resources such as:

- Dos and Don'ts under White paper on Artificial Intelligence and Judiciary (2025)
- Annexure 2: Court User Guidance for AI - AI for Justice: Ethical, Fair and Robust Adoption in India's Courts (2026)
- AI and Rule of Law (UNESCO MOOC)

"Anonymisation" – Clause 3(k), read with Clause 19(1)(i)

The definition correctly treats anonymisation as irreversible, but Clause 19(1)(i) permits anonymisation specifically "for publication", narrower than other internal-processing contexts in which the Digital Personal Data Protection Act, 2023 and its 2025 Rules separately contemplate reversible de-identification or pseudonymisation.

We recommend distinguishing between (i) irreversible anonymisation for public-domain publication of judgments and orders under Clause 19(1)(i), from (ii) a separate, reversible de-identification standard, consistent with the DPDP framework, for internal processing where re-identification may legitimately be required, for instance, in family, juvenile, or sexual-violence matters, where established identity-protection procedures already operate and should not be displaced by a blanket irreversibility requirement.

"Court data" – Clause 3(p)

The definition correctly includes within its ambit data generated or in the possession of the court.

We recommend that the definition also include any data generated by way of "court process" to prevent any ambiguity in the scope "*of data … generated by the court*".

"Ethical Impact Assessment" – Clause 3(v)

The definition refers to a "structured" evaluation "as specified by the Appropriate Authority," with the operative content deferred.

We recommend that the standard format to be prescribed under Clause 35(2) draw on the four-part Institutional Readiness, Risk, Technical and Ongoing Assessment framework set out in Part II.D below, which we have developed for the Indian court context. We would welcome the opportunity to share these instruments directly with the AI Secretariat and CoRE-AI for adoption or adaptation.

"High-Risk AI Tools" – referenced at Clauses 7(3), 12(2) and 46(4)(h), but not defined

"High-risk" uses require heightened scrutiny under Clause 7(3), stronger safeguards under Clause 12(2), and mandatory explainability documentation in vendor contracts under Clause 46(4)(h), yet this is only defined as AI "affecting personal liberty or any lawful right of a person". This is a threshold broad enough to describe many kinds of judicial AI use and may not serve as a useful classification in practice.

Our report cautions against a binary "high-risk/low-risk" division (Section 3.1 of our accompanying report), observing that risk-first approaches limited to two tiers tend to understate cumulative and long-term harms. Instead, we propose a four-tier framework calibrated to the degree to which a use case might directly influence or steer judicial outcomes. We recommend the regulations and AI Committee(s) adopt a similar classification directly within Clauses 3 and 18, as set out below:

Indicator: how the use case influences judicial functioning	**Description**	**Proposed risk classification**
May decide or strongly steer case outcomes	Replaces or significantly influences judicial discretion or final decisions (e.g., sentencing determination).	**High risk — recommend prohibition**
Advises on substantive legal issues	Assists with legal reasoning or interpretation where explanations may be unclear or incomplete.	**High risk**
Could potentially influence judicial or staff reasoning	Outputs may shape thinking, even though judges or staff retain full discretion.	**High risk**
Provides advisory legal assistance	Helps understanding of law or cases but is explicitly non-binding and optional.	**Medium risk**

Supports legal work without affecting decisions	Assists with drafting, summarising, or searching case materials, including real case text.	**Medium risk**
Assists with routine, non-adjudicatory tasks	Supports optimising scheduling, case tracking, or procedural workflows with limited internal data.	**Medium risk**
Handles administrative or procedural outputs only	No legal analysis or decision support.	**Low risk**
AI used for peripheral or general purposes	Unrelated to judicial work or decision-making.	**Low risk**

We recommend that uses falling in the first row above be treated similarly to Clause 20's existing, non-derogable prohibitions, and that the remaining tiers determine the intensity of verification, audit frequency, and disclosure obligations.

"Synthetic Data or Synthetic Information" – Clause 3(zi), read with Clause 43(5)

The definition refers to AI-generated content that mimics real data and distinguishes it from that "directly collected from real-world events or persons" or "produced by established mathematical or scientific modelling that is not AI-assisted". Synthetic data is typically used to refer to proxy data used for model training purposes, where real world data is not available.

However, Clause 43(5)'s disclosure obligation for the use of "Synthetic Data or Synthetic Information… in any judicial proceeding" is unclear.

- It could apply to both AI-generated material submitted by advocates or parties that is "fabricated, false, misleading, or inaccurate" as mentioned under Clause 43(5), against which the Court may pass orders. But it could also apply to use of AI-generated synthetic exhibits – for instance, synthetically generated audio, video or images used to model or refine evidence,[3] which may be disclosed and utilised.
- Submitting any AI-generated outputs without "full and transparent disclosure of its AI-generated character" is already a prohibited use under 20(1)(h). Therefore, we recommend either removing Clause 43(5) and (6) considering the established prohibited use or re-wording it to specifically highlight the nature of AI-generated synthetic exhibits.

B. General Principles (Chapter II)

Accountability and the verification waiver – Clause 8(3)

The first proviso to Clause 8(3) permits the responsible officer to dispense with verification of AI-generated output "for reasons to be recorded in writing" and the second proviso exempts administrative tools "certified…to have established reliability" from verification on a class basis.

Two modifications may be required. First, "established reliability" should not rest on a one-time certification: generative and adaptive systems drift and degrade in performance over time, a concern our report addresses in detail under "model drift" in its post-deployment mitigation guidance. We recommend tying "established reliability" to periodic re-certification, aligned with the audit cycle under Clause 38, rather than a single determination.

[3] https://www.ncsc.org/resources-courts/ai-generated-evidence-guide-judges

Second, the discretion to waive verification should itself be risk-calibrated: waiver is more defensible for tools classified low-risk under the tiering system described above than for tools that may materially affect case-management decisions. We recommend the Regulations provide that verification may not be waived for any tool classified medium-risk or above.

Proportionality, read with Clauses 38, 40 and 41

Clause 12(2) calls for "correspondingly heightened safeguards" for higher-risk applications, but the uniform standards actually imposed elsewhere, a single one-year audit interval under Clause 38(1), unqualified individual-officer discretion under Clause 40, and a uniform one-year compliance-review period for legacy systems under Clause 41 do not vary by risk.

We recommend each of these provisions be expressly tied to the four-tier classification proposed in Part II.A, so the most rigorous safeguards apply where risk is highest and lower-risk administrative tools are not weighed down with verification burdens.

Data integrity – Clause 14(2)

The prohibition on deploying AI Systems trained on "demonstrably biased" datasets presumes a mechanism capable of identifying and establishing bias, without specifying what counts as "demonstrable" or who bears the burden of demonstrating it.

We recommend the standard Technical and Ethical Impact Assessment format under Clause 35 requires sharing/disclosure of bias-testing methodology and documentation through formats such as data cards, model cards, and bias-audit records, of the kind our report's Technical Assessment recommends vendors maintain.

Presumption in favour of responsible AI adoption – Clause 16

"Responsible AI" has been variously and at times inconsistently defined across academic, civil-society and industry usage internationally; left undefined here, the presumption in Clause 16(1) risks being invoked to justify deployments that a more rigorous, rights-based standard would not support.

We recommend the presumption be operationalised through a mechanism (even if it is yet to be developed) such as the Technical and Ethical Impact Assessment. This way, "responsible adoption" may be understood within the context of the regulations and not draw on broader, possibly inapplicable, concepts defined in part by big tech companies building global foundational AI models.

Innovation over Restraint – Clause 17

The phrase "innovation over restraint" has some tension with the cautious, rights-protective character of Chapter II. However, its substantive content including the qualifying language ("all other things being equal") and its requirement that innovation be pursued "with due regard to the other general principles" is considerably more measured than the heading suggests.

We recommend the principle be operationalised concretely by directing that innovation be pursued specifically through the Controlled Environment Testing mechanism at Clause 36, so the principle expands the pathway for responsible context-driven experimentation without diluting the substantive protections set out elsewhere in this Chapter.

C. Permissible and Prohibited Uses (Chapter III)

Legal research and citation verification - Clause 19(1)(d)

The use of any form of machine learning for citation verification negates the principle of human primacy. Verification of citations generated by machines and otherwise cited must be completed by humans and not left to the judgement of machines. This is especially relevant in the context of Clause 43(5), which holds any person submitting AI-generated content that may be "fabricated, false, misleading, or inaccurate" shall "bear full responsibility" - if an advocate relies on a legal research database or citation verification tool that produces errors, then the Court "may pass orders" against them.

Conversational AI assistants and chatbots – Clause 19(1)(f)

These are permitted "*subject to human oversight of their functioning*" but oversight of a system's general "functioning" is distinct from oversight of each individual interaction with a litigant, who may not be able to recognise when incorrect legal information has been provided.

We recommend approval under Clause 18 for any such assistant be conditional on:

- restricted purpose to merely guide navigation on a court website or mobile application;
- demonstrated adequacy of training data for the specific legal domain and language(s) of deployment;
- continuous monitoring with periodic sampling review of actual responses, not solely of system functioning;
- mandatory incident reporting under Clause 39 specific to chatbot misinformation;
- clear and conspicuous labelling that the assistant provides legal information only, and not legal advice or emotional support; and
- an accessible grievance-redressal pathway signposted within the interface itself.

The Clause 19(2) approval pathway and shadow use

Clause 19(2)'s requirement of prior written approval for any use "not enumerated" does not, by itself, address unauthorised or undisclosed use, which proceeds without seeking that approval.

We recommend Clause 21's remedial mechanism be triggered by discovery of any undisclosed AI use; in addition, capacity-building modules under Clause 49 should outline what shadow use is and address its risks, rather than relying solely on ex post enforcement.

Prohibited uses – Clause 20, and the dilution of "prohibition" by conditional approvals elsewhere

Several uses framed as absolute prohibitions are, on a fuller reading of the draft Regulations, conditionally permitted once approvals and human oversight are obtained. Clause 20(1)(a)'s prohibition on training AI using personal data "without the prior approval of the Appropriate Authority" is, by its own terms, an approval-conditioned restriction rather than an absolute bar; similarly, adjudication and sentencing functions barred under Clause 20(1)(c) remain permissible "with mandatory Human-in-the-Loop." This sits uneasily with Clause 20(2)'s statement that the prohibitions "shall be absolute."

We recommend the Regulations distinguish between provisions that are genuinely non-derogable – such as Clause 20(1)(d) (the bar on Risk Scoring) and Clause 20(1)(f) (the bar on predicting future conduct), which contain no qualifying language and those that are, more accurately, conditional permissions subject to approval and oversight, such as Clauses 20(1)(a)

and (c). This distinction bears directly on how Clause 21's remedial mechanism for "violation of any prohibition" will be enforced.

On the practical enforceability of Clause 20(1)(f)'s bar on predicting future conduct, particularly against parties, witnesses or counsel using tools outside any Court's direct oversight. We recommend further operational guidance, including illustrations of prohibited use, be developed in coordination with State Bar Councils and integrated into the mandatory training contemplated under Clause 49.

Qualifier for Algorithmic Decision-Making - Clause 20(1)(b)

The language of this clause - that no judicial outcome can be determined 'solely' or by Algorithmic Decision-Making (ADM) 'alone' risks creative misinterpretation. In such a case, as long as the human judicial authority reviews a decision generated entirely through an AI-assisted process, they would meet the requirement of this section, which does not seem to the intent or spirit of the prohibition on ADM in the judicial context.

D. Mandates, Oversight and Disclosure (Chapter V)

Technical and Ethical Impact Assessment – Clause 35

Clause 35(2) gives the Appropriate Authority six months from commencement to prescribe a standard format, and Clause 35(3) lists six minimum evaluation criteria. We recommend this standard format adapt as needed the four-part assessment architecture set out in our accompanying report:

- an Institutional Readiness Assessment (covering regulatory, infrastructural and capacity preparedness);
- a Risk Assessment (envisioning use-case-level risk and rights impact);
- a Technical Assessment (vendor-side due diligence covering training-data provenance, explainability documentation, bias-testing and harms-modelling, and audit-trail and sandbox-testing commitments); and
- an Ongoing/Continuous Assessment (post-deployment monitoring for model degradation and incident reporting).

We further recommend that the "purpose" criterion under Clause 35(3)(a) be operationalised through a objectives-justification list, so that Clause 11's purpose-limitation principle has a clear basis at the point of approval. We propose the following objectives, drawn from our report as a start:

No.	Illustrative objective(s) for considering an AI-based tool for a given use case
1.	To improve accuracy, completeness, or reliability
2.	To streamline or optimise tasks or reduce time required from court personnel
3.	To support greater standardisation, uniformity, and procedural fairness in administrative tasks
4.	To shift workload from repetitive tasks to tasks requiring judgment, discretion, and human engagement

5.	To optimise judges' time available for adjudicative and evaluative duties
6.	To support analytics or dashboards to inform planning, supervision, and accountability
7.	To improve accessibility, including multilingual access, audio support, or plain-language explanations
8.	To improve transparency and strengthen litigants' experience through better information, status updates, and guidance
9.	To support innovation of judicial operations in line with emerging global standards, provided adequate safeguards and oversight mechanisms are in place

Controlled Environment Testing – Clause 36

This term does not draw a clear distinction between a short pilot and a sandbox testing arrangement involving live, but isolated, data.

- 36(2) refers to conditions usually applied to pilot test phases, specifically "compatibility with existing Court processes" which may not be discernible through a truly isolated sandbox.
- However, 36(3) includes the qualification that the Controlled Environment Testing "shall not affect, interfere with, or compromise the integrity, security or functioning of the primary operational systems or networks of any Court, nor shall the outputs of such testing be used in any actual adjudicatory or administrative decision in a Court process."

We recommend that the governing agreement for any Controlled Environment Testing arrangement be made explicit and recommend that all pre-deployment piloting, however informally initiated, be required to proceed through this Controlled Environment Testing pathway, giving the "innovation over restraint" principle a clear institutional pathway.

Audits – Clause 38

Clause 38(2) requires audits to be conducted exclusively "in-house," with source code, algorithms and datasets never shared with any third party "for an audit outside the Court premises." This creates a conflict.

Our report's institutional-readiness findings indicate that most courts currently rely on a small number of staff with limited technical qualifications rather than a dedicated, professionally recruited technology cadre, and an in-house-only audit mandate may exceed the practical technical capacity of many High Courts, particularly for frequently updated generative AI models. Standard international audit practice and our report's Technical Assessment guidance recommend periodic third-party audits as the norm rather than the exception.

We recommend Clause 38(2) be modified to permit independent third-party audits under strict confidentiality, data-localisation and non-disclosure safeguards – modelled on the audit and confidentiality clauses proposed for Annexure 1 of our report – rather than mandating in-house-only review irrespective of available capacity.

Discretion in supervising AI Systems – Clause 40

Clause 40 grants the nominated officer "full discretion to accept, modify, or reject" AI-generated recommendations or outputs "within the matters falling under his charge." Where an AI tool is

deployed Court-wide, a case-allocation or scheduling tool, for instance, a single nominated officer cannot meaningfully review every output generated for every user.

We recommend the verification frequency implied by Clause 40 be calibrated proportionately to the risk classification in Part II.A so the provision remains practically enforceable at scale:

- full, individualised review for high-risk tools, and
- sampling-based or aggregate review for medium- and low-risk tools.

Review of AI Systems already in use – Clause 41

The one-year window for the AI Secretariat to review legacy AI Systems for compliance is, perhaps, too long, given that legacy systems are already operating without having been assessed against this framework.

We recommend this review be conducted within a quarter from commencement, prioritised by the risk classification i.e. starting from any legacy system that would be classified high-risk.

Transparency and disclosure – Clause 43

Clause 43(1) limits the disclosure obligation to AI use that "materially assists" case management, document analysis, or judicial administration "that may affect the conduct of their proceedings" – a standard that is subjective and capable of inconsistent application across High Courts.

We recommend the "materially assists" proviso be removed, instead establishing a disclosure obligation for any AI use falling within Clause 19's permitted categories, which is consistent with Clause 43(2)'s statement that the obligation "shall apply in all permitted uses of AI specified in Clause 19," which suggests disclosure is intended to be comprehensive.

On the Annexure I and II declaration formats contemplated under Clause 43(3) (*addressed further in Part IV below*), we recommend the prescribed formats capture, at minimum:

- the name and version of the AI System or Tool used;
- its approved purpose and risk classification;
- the steps taken to verify accuracy; and the
- identity of the Designated Officer responsible for that use.

E. Annual Reporting (Clause 45)

Clause 45 introduces an Annual Transparency Report to be submitted by "[e]very High Court, Tribunal and Commission" but does not specify the intended audience, format, or purpose: whether it functions as an internal compliance instrument or a public-facing accountability document.

In addition,

- the Apex Body is to publish an Annual Governance Report on the state of AI in Indian Courts and
- every AI Committee shall prepare an annual report on the state of AI adoption within its jurisdiction and submit to the Chief Justice, which shall also be shared with the Apex Body

This creates a great administrative burden in developing these annual reports. We recommend the report's content be standardised around the AI Register's own fields under Clause 37, so that the nature of reporting and the formats of two instruments remain consistent.

F. Procurement (Chapter VI, Clause 46)

On Clause 46(2)'s reference to "ethical standards" as an evaluation criterion, we recommend this be reframed in terms of a vendor's demonstrated capacity to avoid or safeguard against specific, named categories of harm, bias and discrimination, privacy violation, loss of explainability, and improper data use – rather than left to an undefined "ethical standards" assessment.

Clause 46(3) vests procurement procedure-setting in "such procedures as the Chief Justice may determine," which is a departure from the institutional architecture established under Chapter IV, in particular, the Committee on Infrastructure and Finance's express mandate over "financial aspects of AI adoption" under Clause 28(2).

We recommend that the procurement procedure be settled by, or in close consultation with:

- the Apex Body on the Committee on Infrastructure and Finance's recommendation, or
- the AI Committees at high courts where state government may be asked to supplement the financial resources required

rather than vested in the Chief Justice's determination alone, for consistency with the institutional design adopted broadly in these regulations.

Clause 46(4)'s mandatory contractual provisions are a solid foundation, but it could be supplemented by several specific protections recommended in Annexure 1 of our report. We recommend the following additions:

- a defined data-retention and secure-deletion obligation, including a formal deletion certificate upon completion of an engagement;
- a current subprocessor register, with prior written approval required for any subcontracting that creates new data exit or entry points, and contractual flow-down of confidentiality, localisation and audit obligations to subcontractors;
- confidentiality obligations that survive termination of the contract indefinitely, rather than for a fixed term;
- a standardised incident-reporting timeline – for instance, within six hours of a serious incident, consistent with the CERT-In Directions already binding on regulated entities – in place of the present "without delay" standard at Clause 46(6), which is open to inconsistent interpretation; and
- mandatory vendor disclosure of ownership, investors, funders, and any political or commercial affiliations bearing on judicial independence and neutrality, a safeguard currently absent from Clause 46(4) despite its evident relevance to a court's confidence in a vendor's institutional neutrality.

On Clause 46(4)(k)'s prohibition on retraining or fine-tuning models using Court data without AI Committee approval, which we support, we recommend the regulations specify (or identify which body will later specify) the technical grounds on which such approval may be granted or withheld – including disclosure of the retraining objective, the safeguards against

re-identification of anonymised inputs, and whether the retrained model will be deployed exclusively for the relevant Court or more broadly – so the approval is not discretionary.

G. Grievance Redressal (Chapter IX)

Clause 52(1) requires an aggrieved party to "file an application… to the Court" in which the relevant AI System was used.

We recommend that the regulations establish a dedicated administrative grievance channel for AI-specific complaints - this could be overseen by the AI Secretariat at first instance, with the Apex Body prescribing potential sanctions for degree of violation.

Further, a grievance may arise for a broad range of reasons and is not limited to "direct or indirect effect of prohibited use of AI" - limiting the remedies available under the grievance redress mechanism to this category excludes other potential harms.

Part III: Further Opportunities

This Part sets out further opportunities the AI Committee may wish to consider, several of which draw directly on our accompanying report's assessment framework and "Moving Forward" recommendations.

Reframing institutional objectives at the point of approval (Clause 18)

Clause 18 tasks the Appropriate Authority with identifying the areas of Court process in which AI may be used, but the draft Regulations do not explain how a given use case will be considered for the use of AI. The threshold question our report poses is: *is AI necessary, or the only means, for this use case*? We recommend Clause 18 be read together with a purpose-justification requirement, using the objectives checklist proposed in Part II.D above, so that the AI Committees have a documented basis for distinguishing a genuine efficiency or access-to-justice gain from adoption driven by the perceived novelty of the technology itself.

A common Judicial Information Policy

We recommend the regulations contemplate, alongside the AI Register under Clause 37, a broader Judicial Information or Data Policy governing data definitions, collection, processing and sharing standards across Courts. The quality of any AI system deployed under these regulations depends fundamentally on the quality, structure and completeness of underlying Court data, much of which remains incomplete or inconsistently digitised across jurisdictions.

Admissibility of AI-translated documents – Clause 19(1)(c)

Translation of judgments, orders and pleadings is permitted "*subject to human verification of accuracy and fidelity to the original*," but the regulations do not state what legal weight a translated document carries once verified.

Till the time the matter is settled by legislation or case law, we recommend the regulations clarify that an AI-translation functions only as an accessibility aid. Where the admissibility of a document is in question and it is not in the official language of the court, evidentiary value must only be attributed to certified official translations, unless otherwise established by law after due consideration of the accuracy of AI-translated and human validated documents.

Document authenticity verification and interoperability (Clause 19(1)(h))

Clause 19(1)(h) permits document authenticity verification and fraud detection as a permissible administrative use, but its practical value depends on interoperability with external record systems, land records, banking systems for cheque-bounce verification, and similar registries that the eCourts system cannot request information from across India.

We recommend the AI Committee treat this as a sequencing dependency: the underlying data-interoperability infrastructure should precede or proceed alongside AI deployment for this use case, rather than the AI tool being expected to substitute for that limitation.

A central, cross-jurisdictional Judicial AI Register

Beyond the per-Court AI Register under Clause 37, we recommend the Apex Body maintain a central Judicial AI Register documenting all AI initiatives nationally (Controlled Environment Testing, pilots, proposed projects, discontinued efforts, and full deployments), recording vendors, funding models including any public-private partnerships, contractual terms, sandbox

outcomes, and incident history. A public-facing version, excluding sensitive fields, would strengthen the transparency objectives already embedded in Clause 45, while an internal version would enable cross-court learning and benchmarking of institutional maturity.

Broadening the Risk Scoring prohibition (Clause 20(1)(d))

Clause 20(1)(d) prohibits Risk Scoring "for any purpose in Court processes." We recommend this prohibition be read, and if necessary made explicit, to extend to analytics or dashboard tools permitted under Clause 19(1)(j) that use aggregate or large-scale data analysis to produce comparable scoring or ranking outputs about individuals – for instance, litigant or counsel "reliability" scores derived from case-history analytics – since the underlying rights concern (an opaque, individualised probability judgment substituting for human evaluation) is the same whether or not the tool is formally labelled "Risk Scoring."

Minimum institutional safeguards across all risk tiers

Drawing on our report's "Moving Forward" recommendations, we propose the regulations or accompanying guidance embed certain baseline safeguards across every use case and risk tier, regardless of classification such as:

- domain-expert consultation during tool development;
- clear public disclosure;
- Human-in-the-Loop review involving relevant legal or substantive expertise;
- mandatory verification of AI-generated summaries, translations or recommendations before reliance;
- data minimisation and privacy-by-design;
- an opt-out or restriction mechanism for litigants regarding AI use in their own proceedings;
- a dedicated [administrative] complaint and redress channel;
- ongoing cybersecurity risk assessment;
- bias and fairness auditing; and
- user-level audit logs of AI interactions, retained for accountability and possible review.

This list is intended as a floor, not a ceiling, and AI Committees should be encouraged to adopt additional safeguards as capacity allows.

Part IV: Miscellaneous Notes

Annexures I and II (Clause 43(3))

Annexures I and II, prescribing the declaration and certificate formats for party-initiated and Court-initiated AI use respectively, are referenced in the draft Regulations but not provided.

We recommend, at minimum, the fields proposed in Part II.D above – tool name and version, approved purpose, risk classification, verification steps taken, and the responsible Designated Officer – be incorporated into both formats, which should be differentiated as follows:

- Annexure I, for party- or counsel-initiated AI use, should include an attestation comparable to the disclosure and certification requirements that have emerged in other jurisdictions' court rules on generative AI in filings – for instance, California's Rule of Court 10.430, effective September 2025, and analogous federal district court orders; and
- Annexure II, for Court-initiated use, should additionally record the approving authority and the date of approval under Clause 19.

Standard format for Controlled Environment Testing (Clause 36)

We note the absence of a standard format for the governing agreement contemplated for Controlled Environment Testing arrangements under Clause 36, and recommend this be prescribed contemporaneously with the Technical and Ethical Impact Assessment format under Clause 35(2), rather than left open-ended.

Phased commencement (Regulation 1(2))

Clause 1(2) permits the regulations, or any provision, to come into force on dates separately notified by the Chief Justice of India and each High Court's Chief Justice – contingent, in practice, on the Apex Body, its five Committees, CoRE-AI, and every High Court's AI Committee and Secretariat first being constituted.

We recommend either a phased-commencement schedule within the regulations themselves, under which the foundational principles of Chapter II and the non-derogable prohibitions at Clause 20 take immediate effect, with the institutional chapters (IV and V) rolling out over a defined period, or, at minimum, a published implementation timeline accompanying the final regulations, so the protective core of the framework is not held in abeyance pending the completion of institution-building.

Concluding Remarks

We reiterate our appreciation for the Supreme Court AI Committee's initiative in bringing forward these draft Regulations – among the most comprehensive judicial AI governance frameworks proposed anywhere to date – and our comments above are offered in a spirit of strengthening its underlying objectives. Our principal submission is that the framework would benefit from precision around the handful of terms that could carry legal consequences – "high-risk," "demonstrably biased," "responsible adoption," and the disclosure threshold under Clause 43 – and from a more unified institutional architecture that avoids overlapping mandates while ensuring adequate, clearly benchmarked technical capacity at every level of the system.

We would be glad to share, in whatever format is most useful to the AI Committee, the assessment instruments referenced throughout this submission and remain available for any further consultation the Committee may consider useful as these regulations are finalised.