

॥ न्यायस्तत्र प्रमाणं स्यात् ॥



Comments on the

Draft Regulations for Use of Artificial Intelligence (AI) in Courts, 2026

To the Artificial Intelligence Committee
Supreme Court of India

July 2026





About NLU Delhi

National Law University Delhi (NLU Delhi), one of India's premier institutions of legal education and research, was established by the Government of the National Capital Territory of Delhi under Act No. 1 of 2008 on the initiative of the High Court of Delhi. As a publicly funded university, it has the Chief Justice of India, or a nominee Judge of the Supreme Court, as its Visitor and Chair of the Governing Council, while the Chief Justice of the High Court of Delhi serves as its Chancellor. In less than two decades, NLU Delhi has emerged as a pre-eminent centre of legal scholarship, securing the second position in the Law category of the National Institutional Ranking Framework (NIRF) for eight consecutive years and recently making its debut in the QS World University Rankings by Subject for Law and Legal Studies.

The University is committed to advancing rigorous, interdisciplinary and socially responsive legal education grounded in constitutional values, professional ethics and the rule of law. Alongside its undergraduate, postgraduate, doctoral and professional programmes, its specialised centres undertake doctrinal, empirical and interdisciplinary research across law, justice and governance. The University extends the impact of this work through research submissions, consultations, policy advice and collaborations with courts, government institutions, universities, international organisations, think tanks, civil-society organisations and the private sector.

At a time when artificial intelligence is poised to transform legal education and practice, public administration, adjudication and the delivery of justice, the University is increasingly examining the implications of AI and other emerging technologies for law and governance, including the ways in which they are reshaping the judiciary and the wider justice system. Its work seeks to ensure that the adoption of technology strengthens the efficiency, accessibility and responsiveness of the justice system while remaining firmly grounded in fairness, transparency, accountability, constitutional rights and the rule of law.

About IGAP

The Indian Governance and Policy Project (IGAP) is a research-led policy organisation working at the intersection of public policy, governance, technology and strategic affairs. Founded by alumni of National Law University Delhi, IGAP draws upon a strong tradition of interdisciplinary legal and policy research. Its work is informed by a close understanding of geopolitics, India's developmental priorities, political economy, sectoral regulation and institutional practice.

IGAP brings together multidisciplinary and sector-specific expertise to advance rigorous, evidence-based policy thinking. Its approach combines analytical depth with a careful appreciation of institutional and operational realities, enabling it to produce research and recommendations that are balanced, contextually grounded and relevant to real-world decision-making.

IGAP develops policy and regulatory solutions that advance legitimate national security, economic and governance priorities while remaining firmly anchored in constitutional values, the rule of law and institutional accountability. Through its research, policy engagement and collaborative initiatives, IGAP seeks to contribute to the development of effective, responsible and future-oriented public policy.

TABLE OF CONTENTS

Summary of Recommendations	01
1. Overview	10
2. Chapter I - Key Definitions	13
3. Chapter II - General Principles to Govern Adoption, Deployment and Use of AI Systems	23
4. Chapter III - Permitted, High-Risk and Prohibited Uses	32
5. Chapter IV - Institutional Architecture and Operational Requirements	37
6. Chapter V - Oversight, Audits and Incident Management	46
7. Chapter VI - Procurement and Private Sector Engagement	51
8. Chapter VII - Data Protection and Cyber Security	54
9. Chapter VIII - Capacity Building, Training and Best Practices	59
10. Chapter IX - Grievance Redressal and Remedies	63
11. Chapter X - Miscellaneous Provisions	66
12. Operational Readiness and Implementation Sequencing	67
13. Annexure - I	71

SUMMARY OF RECOMMENDATIONS

The Draft Regulations are strong in their normative design. Their emphasis on human primacy, transparency, auditability, data protection, accountability and responsible adoption is well placed. The codification of hallucination as a court-specific risk, the prohibition on AI performing adjudicative or sentencing functions, the creation of institutional oversight structures, and the recognition of Sensitive Judicial Data are thoughtful interventions. In some respects, they go further than comparable judicial AI instruments in other jurisdictions.

The principal challenge is implementation. A judicial AI framework cannot rest only on high-level principles. Courts will need to know which systems are covered, how risk is classified, who approves deployment, what records must be kept, when systems must be audited, how incidents are escalated, and what remedies are available when AI use causes harm. The recommendations below are directed to that limited but essential task. They seek to preserve the ambition of the Draft while making it clearer, more workable and more capable of consistent application across the Supreme Court, High Courts, district courts, tribunals, statutory commissions and specialised forums.

1. Scope and Key Definitions

- **Materiality threshold for “AI System.”** The definition should not be triggered merely because a general purpose tool, such as a word processor, search tool or case management system, contains an embedded AI feature. Regulation should attach, and the intensity of impact assessment, registration, audit and disclosure obligations should scale, according to whether the AI functionality materially shapes, assists, influences or automates a court process, not according to the incidental presence of a feature in routine software.
- **“External System.”** The scope of this term should be clarified, distinguishing commercial third party systems from government supported infrastructure, such as NIC systems, court authorised hosting, sovereign cloud, MeitY empanelled services and STQC audited environments. A coordination clause addressing the relationship between this definition and applicable data localisation and procurement requirements would reduce inconsistent application across jurisdictions.
- **“Black Box.”** A functional definition, describing the effect (an inability to provide an adequate explanation of a given output) rather than the underlying mechanism, would be more legally precise and durable, drawing on the transparency framing in Articles 13 and 14 of the EU AI Act. Explainability obligations should similarly be recalibrated to what is technically feasible and sufficiently informative for judicial use and

meaningful oversight, rather than requiring disclosure of internal reasoning, factors and weightings.

- “AI Incidents” and “substantial risk.” The threshold for a “substantial risk” incident should be clarified by specifying relevant factors, including the likelihood and nature of harm, with periodic guidance identifying categories of events that would ordinarily meet the threshold as systems evolve.
- “Court”. The scope of this definition should be clarified, including whether it extends to bodies such as Juvenile Justice Boards. Since covered institutions perform materially different adjudicatory functions and operate under distinct statutory confidentiality regimes, minimum standards and implementation requirements should permit forum specific adaptation rather than a uniform approach.

2. Risk Classification Framework

The Draft invokes proportionality and heightened safeguards but does not yet provide an operative risk classification framework. A four tier structure is recommended, with each tier carrying pre-assigned obligations so that classification, rather than a fresh determination in each case, settles the intensity of assessment, oversight and audit.

- Low risk: administrative functions not touching the merits, such as scheduling, cause list preparation, defect scrutiny, routine notices, certified transcription, record management and anonymisation for publication, subject to expedited approval, simplified assessment, basic logging and periodic audit.
- Intermediate risk: tools shaping material placed before the court without determining outcomes, such as legal research, translation, summarisation, evidence organisation and citation retrieval, requiring structured assessment, mandatory human verification before reliance, defined performance standards, revalidation and use logs.
- High risk: systems capable of influencing findings of fact or law, evidentiary assessment, judicial discretion, sentencing related assistance, liberty, rights or access to justice, requiring comprehensive impact assessment, continuous validation, full traceability, meaningful human review, independent audit and stricter incident reporting.
- Prohibited: uses under Regulation 20, which should remain outside any approval pathway.

This tiering should also anchor other obligations that currently apply uniformly, including data sensitivity classification, the frequency of business continuity

testing, fallback protocol requirements, and the intensity of vendor evaluation and contractual safeguards in procurement.

3. Prohibited and Permitted Uses

- Prohibitions (Regulation 20). The Regulations should distinguish expressly between non-derogable prohibitions, such as AI adjudicating, sentencing, determining outcomes, scoring flight risk or recidivism, profiling parties or witnesses, and compromising deliberations, and conditional restrictions, such as use of personal data for training or surveillance uses authorised by law, grouping them separately. It should be clarified that Human in the Loop review is not a gateway to AI adjudication. If the prohibition on future conduct assessment is intended primarily to protect criminal justice and personal liberty, this should be stated expressly and preserved as non-derogable, with separate, narrower consideration given to civil or commercial analytics that do not affect rights, liberty, credibility or access to relief. The consequences of violation, including disciplinary consequences for responsible officers and whether affected proceedings may be challenged or reviewed, should also be specified, and the relationship between Regulations 7(3) and 20(1)(e), covering permitted research as against prohibited outcome prediction, should be clarified.
- Permitted uses (Regulation 19) and Regulation 19(1)(f). The Regulations should adopt either an exhaustive list model, under which all additional uses require formal approval, or objective criteria allowing functionally similar low risk uses to proceed without separate approval each time. A process for the Appropriate Authority to periodically review and update permissible uses, linked to the Apex Body's standard setting function, is recommended. For conversational or litigant facing tools under Regulation 19(1)(f), minimum safeguards should be considered, including pre-deployment accuracy testing, clear disclosure that the user is interacting with an automated system, a real time escalation pathway to a human officer, and clarity on remedies for reliance on defective guidance.

4. Human Primacy, Oversight and Operational Logging

- Giving content to human primacy. For any system whose output informs a judicial determination, the impact assessment under Regulation 35 could require a short statement of how the reviewing officer is expected to understand, question and, if necessary, set aside that output, paired with the training contemplated in Chapter VIII, so that oversight does not settle into routine endorsement.
- Operational logs. A requirement to generate and retain operational logs would make accountability under Regulation 8, auditability under Regulation 9 and explainability under Regulation 7 verifiable rather than a

matter of assertion. Logging should be risk calibrated, with the fullest record for systems informing adjudication, litigant facing guidance or rights affecting administration, and a lighter record for routine administrative use. Who may access logs, and for how long they are retained, should be specified at a high level, consistent with the confidentiality of the proceedings concerned.

- Bias assessment. Recording the outcome of bias assessments in the AI Register would make a system's position transparent to affected persons and reviewable over time.

5. Audit Standards

The audit obligation under Regulation 38 requires substantive content, namely who may conduct audits, minimum auditor qualifications, scope of assessment, reporting format, and consequences of adverse findings. These minimum requirements should be set out in a dedicated Schedule or a binding Audit Protocol issued by the Apex Body, aligned with recognised standards such as ISO/IEC 42001, covering accuracy, bias, robustness, cybersecurity, privacy, data minimisation, explainability, logging, human review and procurement compliance. For higher risk systems, controlled independent or empanelled audits under confidentiality safeguards, including secure room review, restricted datasets and court supervised access, can preserve judicial confidentiality without reducing audit independence to a formality.

6. Sensitive Judicial Data, Privacy and Security

- Category based definition. Sensitive Judicial Data should be defined by enumerated categories, such as health information, financial records, data concerning children and survivor testimony, retaining the current harm based test only as a residual category for data not otherwise listed. This classification should link to the risk classification framework so that security obligations correspond to actual risk.
- Statutory confidentiality regimes. The Regulations should expressly recognise statutory confidentiality and record management obligations governing specialised proceedings, including, where applicable, obligations relating to anonymity, restricted disclosure and the statutory deletion of records, so that safeguards under the Regulations operate consistently with existing requirements. This includes, for example, the mandatory destruction of records under Section 24(2) of the Juvenile Justice Act, 2015, where Juvenile Justice Boards fall within scope.
- Security baseline. The requirement of commensurate security should be anchored to a recognised technical baseline, such as the ISO/IEC 27001 control families, to make the obligation measurable, verifiable and auditable.

- DPDP Act interface. Administrative and technological court functions fall outside the Section 17(1)(b) exemption and attract data fiduciary obligations. Before deployment for such functions, courts should document the categories of data processed, the purpose of processing and applicable safeguards, and the Regulations should require a structured data protection impact assessment addressing data volume and sensitivity, automation risk, and the adequacy of minimisation and anonymisation measures.
- Data mapping exercise. Courts should undertake a structured data mapping exercise before implementing minimisation and anonymisation obligations, identifying categories of personal data across processes, information essential to traceability and evidentiary value, and applicable statutory confidentiality obligations.
- Sovereign or on premise deployment. The terms “sovereign” and “on premise” should be defined, with clarity on whether sovereignty requires physical localisation, legal control, or both, and deployment requirements should be calibrated to risk classification rather than applied uniformly. Certified, audited infrastructure already operating within the Indian public digital ecosystem, such as MeitY empanelled and STQC audited environments under the GI Cloud framework, should be recognised as compliant alternatives. More broadly, the requirements of in house audit, sovereign or on premise deployment and restricted external transfer should be redrafted to permit certified or independently audited environments meeting defined jurisdictional and security criteria, rather than confining courts to wholly self hosted arrangements.
- Court data and ownership. The scope of “court data” and “court resources,” the threshold for triggering ownership or licensing rights, and whether such rights persist where a system is subsequently fine tuned or retrained on data other than court data, should be clarified.

7. Institutional Architecture

- Executive and technical coordination. Executive technical input, for example from MeitY, may be valuable on cybersecurity, infrastructure and standards, but where the Apex Body exercises approval or oversight functions, the scope of such participation should be bounded to preserve judicial autonomy. This includes specifying whether the MeitY representative's role is advisory or voting, and prescribing quorum and decision making thresholds. A mechanism representing the operational contexts of tribunals and statutory commissions in Apex Body deliberations should also be considered.
- CoRE-AI and the Technical Committee. A clearer functional demarcation is needed between bodies with overlapping mandates, with CoRE-AI

focused on research and capability development and the Technical Committee on deployment oriented assessment and approval support.

- Technical input to decision making. Structured technical input into Apex Body and AI Committee deliberations, including designated technical advisers without requiring formal membership, would strengthen decision quality while preserving judicial authority over final outcomes.
- AI Secretariat. Minimum staffing and technical capacity requirements should be specified, or technical support arrangements provided to supplement in house capacity.
- AI Content Verification Authority. The Authority's mandate should be framed carefully. No authority can reliably certify the truth of all AI generated content or detect every instance of GenAI use, and overclaiming this function risks false confidence. It should instead function as a standards and assurance body, prescribing disclosure formats, provenance standards, verification trails, citation checking protocols, document authenticity processes and consequences for false or misleading filings, distinguishing verification obligations for court generated, lawyer generated and litigant generated content. The Authority's composition and timeline should be specified before commencement.
- The Apex Body's dual role. Given the definitional split between the Apex Body's standard setting function and its approval and oversight functions, either a dedicated clarifying provision or a refined definition is needed to identify which functions fall within its remit.

8. Procurement, Business Continuity and Compliance Standards

- Vendor evaluation and contractual obligations in procurement should be calibrated to the risk profile of the system procured, preserving rigorous safeguards for high risk systems while reducing burden for lower risk ones.
- The frequency of business continuity testing and acceptable recovery times should scale with risk classification, and a tested fallback protocol should be a condition of approval for higher risk systems rather than a parallel obligation.
- The compliance standard under Regulation 41 should be published within the same timeframe as the Minimum Mandatory Standards under Regulation 24, so that the one year review period runs against a defined, uniform benchmark from commencement.

9. Grievances and Remedies

- Permitted use harms. The grievance framework should expressly cover harms arising from defective, inaccurate, undisclosed or improperly supervised use of permitted AI systems, such as incorrect chatbot guidance, transcription errors, translation inaccuracies or failed anonymisation, and not only prohibited uses. It should identify timelines, the responsible authority, standards of review, correction of records, escalation to the AI Committee, and preservation of other legal remedies.
- Disclosure and access. Baseline requirements on disclosure of AI use in proceedings, and access to information necessary to assess and file a grievance, would support meaningful application of the “earliest opportunity” standard under Regulation 52(2). Parties should have clarity on access to relevant audit records or system logs for remedies under Regulation 53, with consideration given to differentiated procedures for permissible use harms.
- Escalation and timelines. An escalation pathway linking grievances under Regulation 52 to the Apex Body's function under Regulation 23 should be established for systemic concerns, with minimum procedural timelines for acknowledgment and disposal specified in the Regulation or through Apex Body standards. The interaction between Regulation 54(2) and enhanced safeguards for Sensitive Judicial Data, transparency requirements for directions under Regulation 55, oversight of the relaxation power under Regulation 56, and a periodic review cycle under Regulation 57 should also be clarified.

10. Training and Capacity Building

Regulation 49 should be operationalised through a phased model beginning with priority groups, with training that is role specific, integrated into existing judicial education structures, and delivered through modular programmes combining foundational literacy with periodic refresher and advanced modules, linked to dedicated budgetary allocation and institutional partnerships for continuity and scalability. A competency based framework prescribing role specific learning outcomes and appropriate combinations of theoretical and practical instruction is recommended, along with minimum standards for multilingual training materials and institutional responsibility for their periodic review. Standardised templates for recording incidents, case studies and best practices, with uniform reporting and review requirements, would support a coherent national repository.

11. Implementation Readiness and Phased Rollout

Beyond drafting refinements, the effectiveness of the Regulations will depend on whether the institutions to which they apply can actually implement them.

Digitisation under the e-Courts Phase III project is significantly behind schedule and unevenly distributed across jurisdictions. Underlying case record data suffers from documented quality issues. Hardware and connectivity remain constraints, as recognised in relation to SUPACE, and vacancy levels in district courts, High Courts and tribunals raise real questions about capacity to discharge new oversight obligations. Certain categories of courts, including Family Courts (given the consent requirements recognised in *Santhini v. Vijaya Venketesh*), POCSO and Fast Track Special Courts, and Juvenile Justice Boards, present additional, forum specific constraints not currently addressed in the Draft.

- The Apex Body should conduct a baseline readiness assessment before commencement, covering digitisation, data quality, network capacity, hardware, staffing, training, cybersecurity posture and forum specific confidentiality requirements. This is not an argument for delay, but for properly sequencing implementation.
- Implementation should proceed through controlled pilots and phased rollout, beginning in courts with stronger digital infrastructure and administrative capacity, evaluated for accuracy, reliability, bias, security, auditability and grievance and incident patterns, following the model of institutional experimentation demonstrated by SUVAS and SUPACE. Institutions meeting defined readiness criteria could implement the full framework immediately, with others following a phased pathway tied to measurable capacity improvements.
- Minimum data quality standards should be established before Technology Enabled Impact Assessments are relied upon, and the TEIA format itself should be published at or before commencement rather than deferred. It should require identification of affected persons and specific harms, and should operate as a continuing obligation with defined review triggers rather than a one time exercise.

12. Regulatory Instruments and Standards

Operationalising this framework will require a substantial number of subordinate instruments, protocols, schedules and standards, for example the Audit Protocol, the Minimum Mandatory Standards, the TEIA methodology, and the classification criteria discussed above. Our draft identifies the full set of such instruments and standards in a consolidated table, together with the body responsible for issuing each and a suggested sequencing logic, on the basis that instruments which are preconditions for other obligations should precede the obligations that depend on them. Readers are referred to that table rather than to a restatement here.

Submission on the Draft Regulations for Use of Artificial Intelligence in Courts, 2026

AI Committee | Supreme Court of India

We at National Law University Delhi (NLU Delhi) and the Indian Governance and Policy Project (IGAP) thank the Artificial Intelligence (AI) Committee of the Supreme Court of India for inviting comments on the Draft Regulations for Use of Artificial Intelligence in Courts, 2026. The Committee's publication of the draft framework for stakeholder consultation reflects a welcome commitment to transparency, institutional dialogue and careful regulatory design in an area that will shape the future administration of justice.

The governance of artificial intelligence (AI) in courts is not merely a question of technology adoption. It is a question of institutional design. It concerns the conditions under which digital systems may assist the justice system without displacing judicial authority, weakening due process, obscuring accountability or diminishing public confidence in adjudication.

The Draft Regulations proceed from the correct premise that AI may support courts, but must remain subordinate to the constitutional function of judging. This is an important and timely intervention. AI is already entering the environment in which courts function, including through legal research, translation, transcription, scheduling, document management, defect scrutiny, litigant assistance and administrative support.

Several features of the Draft Regulations deserve particular recognition. The emphasis on human primacy and judicial independence places the judge, and not the system, at the centre of adjudication. The treatment of hallucination as a court-specific risk is a valuable contribution to AI governance in legal settings. The prohibition on AI performing adjudicative or sentencing functions provides an important constitutional safeguard.

The Draft also raises emerging questions that will require continued attention. Questions around AI privilege and the legal protection of AI-assisted communications, outputs and workflows are increasingly being discussed, particularly where such systems are used in legal advice, court processes or litigation preparation. In India, where low-cost internet and mass smartphone use may make AI a first port of call for legal information, such tools can widen legal awareness but must be clearly bounded so that general guidance is not mistaken for legal advice, judicial direction or court-verified assistance.

Our comments are offered in that constructive spirit. They seek to strengthen the Draft Regulations by improving legal clarity, reducing interpretive uncertainty, aligning safeguards with risk and ensuring that implementation is

calibrated to the varying capacities of courts, tribunals and statutory commissions. The aim is not to slow responsible adoption. It is to make responsible adoption more durable, accountable and institutionally credible.

1. Overview

The Draft Regulations are a comprehensive attempt to develop a dedicated governance framework for the use of artificial intelligence within a judicial system. Drawing upon a range of international instruments, the Draft Regulations seek to adapt emerging principles of AI governance to the distinctive constitutional, institutional, and operational requirements of the Indian judiciary.

As a matter of regulatory design, the Draft Regulations are notably wide in their reach. They bring together high level principles, institutional structures, restrictions on use, procurement controls, data governance obligations, audit mechanisms, incident management procedures and grievance redressal. At the same time, the Draft Regulations reveal a recurring tension between policy intent and operational implementation. Across multiple chapters, the Draft identifies appropriate regulatory objectives but leaves critical aspects of implementation to future standards, guidelines, formats, protocols, or institutional determinations. Several obligations spanning fairness, explainability, proportionality, accountability, cybersecurity, and risk-based oversight, lack sufficient clarity at the foundational level. This creates uncertainty not only for compliance and enforcement but also for the development of the standards and guidance that are intended to operationalise them.

Three themes emerge consistently throughout this submission. First, while the Draft Regulations articulate the appropriate governing principles, they frequently defer the standards, metrics or thresholds necessary for their implementation to future instruments. Second, several provisions presented as bright-line rules, including the prohibitions under Regulation 20 and the treatment of opaque AI systems under Regulations 7 and 20(1)(e), are subject to internal qualifications that reduce their practical certainty. Third, the institutional framework established under Chapter IV allocates overlapping responsibilities across multiple bodies without providing a clear mechanism for coordinating decision-making or resolving differences between them.

The principal challenge is therefore not one of regulatory intent but of clarity. Effective governance of judicial AI requires more than the articulation of desirable principles. It requires clear allocation of institutional responsibility, measurable standards, defined decision-making pathways, and implementation timelines that correspond to institutional capacity. This challenge is particularly acute in the Indian judicial context, characterised by substantial variations in technological maturity, administrative capacity, and digital infrastructure across jurisdictions.

This submission also addresses a broader implementation question (discussed in detail in Section 12 on operational readiness) that runs alongside the analysis of individual provisions. The Draft Regulations apply uniformly across the full range of courts to which they extend, from the Supreme Court to subordinate courts, tribunals and statutory commissions, each of which starts from a materially different baseline in terms of digital infrastructure, data quality, hardware capacity, staffing and training. The experience of SUPACE and SUVAS illustrates the Indian judiciary's preference for phased AI adoption.¹ SUVAS has expanded from the Supreme Court to several High Courts, while SUPACE remains in a limited pilot phase. Together, they demonstrate an iterative approach in which AI systems are deployed incrementally under institutional supervision, with wider adoption following demonstrated reliability rather than preceding it.

This submission recommends that the Apex Body conduct a structured readiness assessment across court tiers before commencement, and that implementation proceed through a phased rollout with pilot projects at courts with established digital infrastructure before extending obligations to those where foundational conditions are still being built. This is not a reservation about the Regulations' objectives but a submission that those objectives are best achieved by building institutional capacity alongside the regulatory framework designed to govern it.

Against this background, this submission evaluates the Draft Regulations through four interrelated lenses. These are legal clarity and interpretive coherence, institutional design and the allocation of responsibility, operational feasibility and readiness and the sequencing of implementation.

While the Draft provides a strong normative foundation, its long-term effectiveness will depend upon whether these principles are translated into a sufficiently clear, risk-sensitive, and operationally workable regulatory framework.

¹ Press Information Bureau, Ministry of Law & Justice, Use of Artificial Intelligence in Legal Field (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202159&lang=1®=3>; Press Information Bureau, Ministry of Law & Justice, Promotion of Hindi in Higher Courts (Feb. 2, 2024), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2001863>.

Chapter and Regulation Map

Ch.	Title	Regs.	Primary Subject Matter
I	Preliminary	1–3	Commencement; application to the SC, HCs, Tribunals, and Commissions; 35 defined terms
II	General Principles	4–17	14 general principles: human primacy, rule of law, fairness, transparency, accountability, auditability, data protection, purpose limitation, proportionality, inclusivity, data integrity, cyber security, presumption of adoption and innovation
III	Permissible & Prohibited Uses	18–21	Illustrative list of 11 permitted uses (Reg 19); 9 absolute prohibitions (Reg 20); remedial measures by AI Committee (Reg 21)
IV	Institutional Mechanism	22–34	Apex Body (Reg 22–24); 5 sub-committees (Reg 25–30); CoRE-AI (Reg 32); AI Committees at HC level (Reg 33); AI Secretariat (Reg 34)
V	Oversight & Incident Management	35–45	TEIA requirement (Reg 35); controlled environment testing (Reg 36); AI Register (Reg 37); audits (Reg 38); AI Incident Database (Reg 39); discretion of supervising officer (Reg 40); review of existing systems (Reg 41); emergency protocols (Reg 42); transparency/disclosure (Reg 43); AI Content Verification Authority (Reg 44); Annual Transparency Report (Reg 45)
VI	Procurement	46	Private sector engagement; 12 mandatory contract clauses; expedited AI Secretariat approval pathway (Reg 46(7)); court IP retention (Reg 46(9))
VII	Data Protection & Cyber Security	47–48	Compliance with DPDP Act 2023 and IT Act 2000; sensitive judicial data handling; data minimisation; anonymisation before training; cybersecurity audit cycle
VIII	Capacity Building	49–51	Structured AI training for judges, advocates, and court staff; best practices repository; biennial review of training adequacy
IX	Grievance Redressal	52–53	Application-based grievance to the court using the AI system; High Court to prescribe format; other legal remedies preserved
X	Miscellaneous	54–57	Non-derogation from existing laws; directions by Chief Justice; relaxation power (excluding Reg 20 prohibitions); periodic regulatory review

2. Chapter I: Key Definitions

2.1 Overall Assessment

The definitions in Chapter I draw selectively from international frameworks including the EU AI Act², the GDPR³, the OECD AI Incident Framework⁴, and comparative judicial governance instruments such as Brazil's CNJ Resolution No. 615/2025⁵. However, the Draft Regulations do not reproduce these models verbatim. Instead, it adapts them to the institutional context of the judiciary, translating market-oriented or general-purpose AI governance concepts into court-specific safeguards and procedural controls. The definitions section carries significant structural weight because every substantive obligation in the Regulations is anchored to a definition, and imprecision at this stage propagates into every chapter that follows. The comments below are directed at a few definitions where greater precision or minor amendments would strengthen the Draft Regulations.

2.2 Defining the Technology and its Use Artificial Intelligence and AI System (Regulations 3(1)(m) and 3(1)(i))

The Draft takes a considered step that the EU AI Act does not. Where the EU AI Act defines only an 'AI system,' Regulations 3(1)(m) and 3(1)(i) separately define 'Artificial Intelligence' and 'AI System,' distinguishing the underlying technology from its use in court processes. Unlike the EU framework, which regulates AI across sectors and can address the technology and its use together, the Draft adopts a separate definition to connect AI specifically to its deployment within the court system. This reflects a narrower and sector-specific focus.

Regulation 3(1)(m) defines 'Artificial Intelligence,' and excludes general purpose software or digital tools unless they are specifically embedded with, augmented by, or functionally dependent upon artificial intelligence. Regulation 3(1)(i) defines an 'AI System' to include "any software, platform, application, device or process that employs Artificial Intelligence in connection with a court process."

Read with the definition of 'Court Process' under Regulation 3(1)(q) which includes "any work, function or activity relating to the adjudicatory, quasi-adjudicatory or administrative functions of a court," and names among them filing, scheduling, hearing management, evidence handling, legal

² Regulation (EU) 2024/1689 (Artificial Intelligence Act) (2024) OJ L 1689/1, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>.

³ Regulation (EU) 2016/679 (General Data Protection Regulation) (2016) OJ L 119/1, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>.

⁴ OECD, Overview and Methodology of the AI Incidents and Hazards Monitor, <https://oecd.ai/en/incidents-methodology>.

⁵ Conselho Nacional de Justiça (Brazil), Resolution No. 615/2025 (11 March 2025, in force from about 14 July 2025), English translation hosted by the Council of Europe. <https://rm.coe.int/resolution-6152025/1680b51b66>.

research, drafting, translation, transcription and record management, it creates a risk of over inclusion.

The carve out in Regulation 3(1)(m) turns on whether a tool is embedded with, augmented by, or functionally dependent upon AI. Ordinary office software increasingly meets that description by default. Word processors, search tools, case management systems and document readers now ship with AI features as standard. On a literal reading, the moment such a tool is embedded with AI, the whole tool becomes an AI System, and attracts the full apparatus of impact assessment under Regulation 35, registration under Regulation 37, audit under Regulation 38 and disclosure under Regulation 43. Without a threshold of materiality, the definition risks reaching a great deal of routine software that the framework was never intended to govern.

For instance, Microsoft Word is widely used in courts across the country, and recent versions arrive with an AI assistant that can draft, summarise and rewrite text.⁶ Under that reading, use of that feature could bring Word within the definition of an AI System, since drafting a note, transcribing a record or preparing a filing is covered as a court process under Regulation 3(1)(q). An AI System should be understood to arise where the artificial intelligence functionality of a tool is material to the court process for which it is used, and not wherever such a feature merely happens to be present in widely used software.

Recommendation: The definition of an AI System should be qualified by a threshold of materiality, so that a general purpose tool is regulated only where its AI functionality is material to the court process for which it is used. The intensity of the obligations that follow should correspond to the significance of the AI's role, so that impact assessment, registration, audit and disclosure attach to systems that materially shape a court process, and not to the incidental presence of an embedded feature in routine software.

Generative AI (Regulation 3(1)(y))

The definition of 'Generative AI' (Regulation 3(1)(y)) embeds operational obligations, specifically disclosure and verification, within the definitional clause itself. Definitions ordinarily describe the meaning of terms rather than impose duties, and combining the two creates interpretive uncertainty. For instance, it may be unclear whether the disclosure obligation applies only in provisions that expressly invoke the defined term 'Generative AI', or whether it applies

⁶ Microsoft Support, Welcome to Copilot in Word, <https://support.microsoft.com/en-us/word/welcome-to-copilot-in-word>.

independently wherever a court officer uses a generative AI tool, regardless of how that use is described elsewhere in the Regulations.

A further concern arises from the reference to "a designated centralised mechanism" for GenAI content verification, which remains undefined. Given that Regulation 44 separately establishes an AI Content Verification Authority, it would be useful to clarify whether this mechanism and that Authority refer to the same institution, and if so, to reflect that consistently across both provisions.

Recommendation: It may be worth considering the placement of obligations of disclosure and verification within the substantive provisions governing GenAI use rather than within the definitional clause. The "designated centralised mechanism" may also benefit from being defined or operationalised before the Regulations come into force, with its relationship to the AI Content Verification Authority under Regulation 44 expressly clarified.

External System (Regulation 3(1)(x))

The definition of 'External System' is framed broadly and may extend to cloud-hosted judicial infrastructure, third-party integrations, and shared technology arrangements with government providers. The practical significance of this definitional question is immediate. Under the e-Courts Project, the judiciary's digital infrastructure has been substantially developed through government-supported systems, with the National Informatics Centre playing an important role in judicial technology applications.⁷ Courts do not, as a general matter, operate AI-independent server infrastructure. Any broad reading of 'External System' that captures existing NIC-hosted or government-managed cloud infrastructure would create an unintended compliance burden for the existing arrangements that are, in many cases, the only feasible option available.

The implications of this breadth for data localisation and procurement must also be considered. For example, where a High Court uses a cloud-hosted AI transcription service operated by a private provider, it may be unclear whether that arrangement constitutes an 'External System' for the purposes of these Regulations, what additional obligations are thereby triggered, and how those obligations interact with the DPDP Act's provisions on data fiduciaries and cross-border data transfers.

⁷ Press Information Bureau, Ministry of Law & Justice, Use of Artificial Intelligence in Legal Field (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202159&lang=1®=3>; Press Information Bureau, Ministry of Law & Justice, Strengthening of E-Courts System (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202162&lang=1®=3>.

Recommendation: Clarification of the intended scope of 'External System' would be helpful, particularly in relation to cloud infrastructure and government technology arrangements. An express coordination clause addressing the relationship between this definition and applicable data localisation and procurement requirements would reduce the risk of inconsistent application across court jurisdictions.

2.3 AI Risks and Safeguards

Black Box (Regulation 3(1)(n))

The definition of 'Black Box' systems departs from the EU AI Act's indirect approach, which regulates opacity through transparency and human oversight obligations rather than definitional categorisation. By directly naming opacity as a regulatory concern, it strengthens conceptual clarity but also risks fixing a technically fluid concept that may evolve rapidly with model architectures.

Hallucination (Regulation 3(1)(z))

The decision to define 'Hallucination' as a legally cognisable category is one of the more significant moves in this chapter. Technical discourse treats hallucination as a known limitation of generative AI i.e. an output that is confident, coherent, and incorrect. The Draft does something more deliberate by adapting the technical concept to the judicial context, and identifying the specific forms the failure takes in a legal setting. The definition expressly includes "fabrication, misstatement or erroneous representation of legal facts, evidence, case precedents, statutory provisions, rules or legal principles". By naming these legal outcomes, the Draft signals that an epistemic failure in an AI output is not an operational risk to be managed at the user's discretion, but a defined category of harm to which legal consequences attach.

Explainability (Regulation 3(1)(w))

The objective is sound, but the formulation may be read as requiring a full reconstruction of a model's internal reasoning, including factors and weightings. For many modern machine-learning systems, particularly large generative models, that level of explanation may not be available in a form that is both technically faithful and intelligible to a non-specialist.

Rather than treating explainability as disclosure of internal mechanics, the Regulations could require explanations that are technically feasible, decision-relevant, and sufficient for meaningful human review and party challenge. This would preserve the Draft's accountability objective while avoiding a standard that capable systems cannot realistically satisfy.

Algorithmic Decision-Making (Regulation 3(1)(j))

The definition of 'Algorithmic Decision-Making' (Regulation 3(1)(j)) improves on the GDPR's Article 22 by capturing algorithmic outputs that inform decisions "whether in whole or in part."⁸ This formulation appropriately extends the regulatory scope to encompass systems that influence, rather than determine, judicial outcomes, reflecting the increasingly hybrid nature of human-AI interaction in judicial decision-making.

AI Incident (Regulation 3(1)(e))

The definition of 'AI Incident' broadly aligns with the OECD AI Incident Framework but departs from it by collapsing the conceptual distinction between realised harms ('incidents') and precursor events ('hazards').⁹ By including events that "create a substantial risk of harm", it adopts a precautionary stance that is suited to judicial environments, where even potential disruption may have systemic implications. However, the phrase "substantial risk of harm" is undefined, and this creates a serious operational problem. Reporting obligations and remedial mechanisms under Chapter V are triggered by 'AI Incidents.' Without a threshold for "substantial," court administrators and AI system operators would apply divergent standards.

Recommendation:

Consider a functional definition of 'Black Box' that describes the *effect* (inability to provide an adequate explanation of how a given output was produced) rather than the underlying mechanism. This approach is both more legally precise and more durable. The EU AI Act's transparency requirements in Articles 13 and 14 offer a useful template for this functional framing.

Retain the objective of explainability, but revise the definition to reflect the present technical capabilities of AI systems. Rather than requiring disclosure of a system's internal reasoning, factors and weightings, the Regulations should require explanations that are technically feasible, sufficiently informative for judicial use, and capable of supporting meaningful human oversight.

The definition of AI incidents would benefit from clarifying the meaning of "substantial risk" by specifying the factors relevant to its assessment, including the likelihood and nature of harm. The Regulations may also provide for periodic guidance identifying categories of events that would

⁸ GDPR, art 22(1).

⁹ OECD.AI, Overview and Methodology of the AI Incidents and Hazards Monitor, <https://oecd.ai/en/incidents-methodology>.

ordinarily meet this threshold as AI systems evolve.

2.4 Human Oversight and Audit

Human in the Loop (Regulation 3(1)(zb))

The definition of Human in the Loop (HITL) appropriately ensures that final decision-making authority, responsibility, and accountability remain with a human.

AI governance frameworks generally distinguish between **Human in the Loop (HITL)**, **Human on the Loop (HOTL)**, involving ongoing monitoring and the ability to intervene, and **Human in Command (HIC)**, where humans retain overall authority over the deployment and use of AI systems. The Draft Regulations use these concepts interchangeably. For example, Regulation 4 reflects Human-in-Command, Regulation 19(f) refers to “human oversight” of conversational AI systems, while adjudicative functions require Human-in-the-Loop under Regulation 20(1)(c). Treating these different forms of oversight under a single label may create ambiguity regarding the level of human involvement expected across different categories of AI systems. Once Human-in-the-Loop is defined with this precision, the use of other, undefined terms for human involvement should be avoided.

Recommendation: Use the defined term Human-in-the-Loop consistently, and define or replace looser terms such as human oversight, so that the level of human involvement expected is clear for each category of use. The distinctions between Human-in-the-Loop, Human-on-the-Loop and Human-in-Command should be reflected where different intensities of oversight are intended.

AI Audit (Regulation 3(1)(c))

The definition of 'AI Audit' establishes the obligation to conduct audits but does not specify the minimum elements of the audit process. Unlike structured audit regimes under instruments such as ISO/IEC 42001 or conformity assessment frameworks under the EU AI Act, it does not indicate who may conduct audits, the scope of assessment, the criteria to be applied, or the consequences of an adverse finding. This risks inconsistent implementation across courts and jurisdictions, and may limit the effectiveness of audit obligations under Regulation 38.

Recommendation: The minimum requirements for AI audits should be specified, covering the qualifications of auditors, scope of assessment, reporting format, and follow-up actions upon adverse findings. These requirements may be set out in a dedicated Schedule or elaborated through a binding Audit Protocol to be issued by the Apex Body, aligned with established standards such as ISO/IEC 42001.

Court (Regulation 3(1)(o))

The Regulation defines Court to include the Supreme Court, the High Courts, all courts, and tribunals and statutory commissions performing adjudicatory functions. The breadth of this definition gives rise to a number of interpretive and practical questions, some of which are discussed below.

Juvenile Justice Boards: A Juvenile Justice Board is constituted under Section 4 of the Juvenile Justice (Care and Protection of Children) Act, 2015 (JJ Act), and comprises a Principal Magistrate sitting with two social worker members. Although the Board exercises the powers of a Magistrate under the BNSS, it is not constituted as a conventional court.

The JJ Act itself draws a distinction between a Board and a Court. Section 2(20) defines a Children's Court as a court under the Commission for Protection of Child Rights Act, 2005, a Special Court under the POCSO Act, 2012, or otherwise a Court of Session. It is the Children's Court, rather than the Board, that tries a child transferred for trial as an adult under Sections 15 and 18. Section 2(23) separately defines a court, for adoption and guardianship matters, as a civil court. Throughout the Act, the institution is consistently described as a Board rather than a Court.¹⁰

It is therefore not clear whether a JJ Board falls within the expression “courts, tribunals and statutory commissions performing adjudicatory functions” in Regulation 3(1)(o).

Fast Track Special Courts: Fast Track Special Courts, established under the centrally sponsored scheme for the trial of rape and POCSO offences, are Sessions Courts and therefore fall within Regulation 3(1)(o). These courts are among those most likely to benefit from AI-enabled tools that support case management, transcription, translation and other administrative functions, given the volume and urgency of the matters they handle. At the same time, their continued operation depends upon periodic extension of the underlying scheme, the latest confirmed sanction for which is understood to run until 31

¹⁰ Juvenile Justice (Care and Protection of Children) Act, No. 2 of 2016, §§ 2(20), 2(23), 4, 15, 18, 24(2), 74 (India).

March 2026.¹¹ It would therefore be helpful to clarify whether the Regulations are intended to apply to such courts notwithstanding the contingent nature of the scheme under which they operate.

The breadth of Regulation 3(1)(o) also warrants careful implementation. The definition brings together institutions whose jurisdiction, records and procedural requirements differ significantly. A POCSO Court, a Family Court and a Debt Recovery Tribunal each deal with different categories of records, operate under distinct confidentiality obligations and follow different procedures.

A dataset, procedure or template that is well suited to one forum may be ill suited, or even unsafe, in another. It is submitted that any minimum standards set under the Regulations should therefore leave room for forum specific calibration, so that a single standardised approach is not applied uniformly across institutions whose data and procedures are not alike.

Recommendation: Clarify the scope of the definition of Court in Regulation 3(1)(o), particularly whether it extends to institutions not explicitly included such as Juvenile Justice Boards. The Regulations should also recognise that the institutions covered by the definition perform materially different adjudicatory functions, handle different categories of judicial records and operate under distinct statutory confidentiality regimes. Accordingly, any minimum standards, assessment methodologies or implementation requirements set by the Appropriate Authority should permit forum-specific adaptation rather than prescribe a uniform approach across all courts and adjudicatory bodies.

Sensitive Judicial Data (Regulation 3(1)(zh))

The creation of a distinct category of 'Sensitive Judicial Data' addresses a legislative gap. The DPDP Act, 2023 does not retain the category of sensitive personal data that existed under the Information Technology (Reasonable Security Practices) Rules, 2011. Courts regularly process categories of information, including health records submitted in evidence, financial disclosures, and testimonies of minors and survivors of sexual offences, that the DPDP Act's general framework does not adequately protect in the AI context. The creation of enhanced safeguards tailored to this environment is therefore a welcome development.

A few aspects of the definition may benefit from further refinement. First, the operative threshold, namely personal information whose unauthorised

¹¹ Press Information Bureau, Ministry of Law & Justice, Scheme of Fast Track Special Courts (Aug. 8, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2154077>.

disclosure "may cause harm", is broad enough to potentially encompass most personal data processed by courts. For example, even a party's contact address, if disclosed without authorisation, may cause harm in certain proceedings. Over-inclusive classification risks diluting the value of the enhanced protections the category is intended to provide. A standard that cannot be practically applied also tends to be applied inconsistently across courts and jurisdictions.

Second, the relationship between this regime and the DPDP Act has not been addressed. Where an AI system processes personal data in the course of a court function, obligations under both instruments may be engaged simultaneously. In proceedings involving court-compelled disclosures, for instance, the DPDP Act's consent-based framework may be difficult to apply, and the applicable rule as between the two instruments remains unclear.

Third, the definition assumes particular significance in the judicial context, where several classes of proceedings are already governed by statutory confidentiality obligations. There is a need to account for the sensitivity and confidentiality of data handled in these categories of proceedings/courts. The following examples illustrate this point:

POCSO Courts: A Special Court designated under Section 28 of the POCSO Act, 2012 is required by Section 33(7) to ensure that the identity of the child is not disclosed at any stage of the investigation or trial. The expression "identity" has been interpreted broadly to include the child's family, school, neighbourhood and other identifying particulars. Section 37 requires proceedings to be conducted *in camera*, while Section 23 of the POCSO Act, read with Section 72 of the Bharatiya Nyaya Sanhita, 2023, restricts disclosure of a victim's identity.¹²

These statutory protections have been reinforced by the Supreme Court in *Nipun Saxena v. Union of India*,¹³ which directed that identifying material be sealed and that a victim's name should not appear even in cause lists. More recently, in *Sheetal Vasant Thakur v. Chirag Arora*,¹⁴ the Supreme Court emphasised that a child's disclosures and therapeutic or evaluative records should remain confined to what is strictly necessary for adjudication and should not be reused in collateral proceedings. These statutory and judicial safeguards underscore the need for the definition of Sensitive Judicial Data to adequately capture records generated in child protection proceedings.

Family Courts: Proceedings before Family Courts routinely involve matrimonial disputes, custody proceedings, allegations of domestic violence, medical records and communications made during conciliation. Much of this material is treated as confidential by statute or established judicial practice. Courts also routinely protect particularly sensitive material through sealed-cover procedures and

¹² Protection of Children from Sexual Offences Act, No. 32 of 2012, §§ 23, 28, 33(7), 37 (India); Bharatiya Nyaya Sanhita, No. 45 of 2023, § 72 (India).

¹³ *Nipun Saxena v. Union of India* (2019) 2 SCC 703.

¹⁴ *Sheetal Vasant Thakur v. Chirag Arora* (2026) SCC OnLine SC 1110.

anonymise judgments where necessary, especially in matters involving children. These established confidentiality practices indicate that records generated in Family Court proceedings warrant recognition as Sensitive Judicial Data.

Juvenile Justice Boards: Section 74 of the Juvenile Justice (Care and Protection of Children) Act, 2015 protects the identity of every child involved in proceedings under the Act. Section 24(2) further requires the relevant records of a child to be destroyed after the expiry of the appeal period or within such a reasonable period as may be prescribed.¹⁵ This is not merely an obligation to limit collection or processing, but a statutory obligation to destroy records after a prescribed period. While the draft Regulations address data minimisation and anonymisation, they do not presently provide for a corresponding obligation to ensure permanent deletion where another law expressly requires it. The Regulations would therefore benefit from recognising statutory record-destruction requirements applicable to specialised judicial proceedings.

Recommendation: A category-based approach to defining 'Sensitive Judicial Data', enumerating specific classes such as health information, financial records, data concerning children, and testimonies of survivors, with the effects-based threshold retained as a residual for data not falling within a named category, may provide greater clarity and more consistent application in practice.

In relation to sensitivity and confidentiality attached to data in certain courts, the Regulations should expressly recognise statutory confidentiality and record-management obligations governing specialised proceedings. This should include, where applicable, obligations relating to anonymity, restricted disclosure and the statutory deletion of records, so that the safeguards prescribed under the Regulations operate consistently with existing judicial and legislative requirements.

A detailed mapping of each definition against its international source is set out in Annexure I.

¹⁵ Juvenile Justice (Care and Protection of Children) Act, No. 2 of 2016, §§ 2(20), 2(23), 4, 15, 18, 24(2), 74 (India).

3. Chapter II: General Principles to Govern Adoption, Deployment and Use of AI Systems

3.1 Overall Assessment

Chapter II reflects a thoughtful engagement with both constitutional values and international best practices. Its emphasis on judicial independence, fairness and accountability is well placed. Certain provisions merit particular recognition. The retention of the Bangalore Principles of Judicial Conduct anchors the use of artificial intelligence within an established ethical framework. The explicit identification of protected grounds introduces a constitutional specificity that is often absent in comparable instruments. The prohibition on systems trained on unlawfully obtained or demonstrably biased data reflects a clear and commendable normative stance.

The remaining task is to ensure that these principles can be realised in practice. The comparative experience suggests that this transition from principle to practice is best achieved through the use of structured frameworks, defined metrics and verifiable records outlined in Section 5 of this submission. The Draft would benefit from adopting similar mechanisms.

Table: Where the Draft stands, and the comparator clauses worth following

Principle	What the Draft already does well	The comparator clause that shows how to operationalise it	Where the Draft stands today
Human primacy (Reg. 4)	It makes AI strictly subservient to human judgment, confines it to an assistive role, and vests final authority in the judge.	CEPEJ Principle 5 ¹⁶ keeps the user in control. EU AI Act Article 14(4) requires that the human overseer be able to set the output aside, override it or stop it, and be alert to automation bias. ¹⁷	At or above best practice as a statement of principle. The capacity of the judge to exercise that primacy in fact is not yet spelt out.

¹⁶ European Commission for the Efficiency of Justice (CEPEJ), European Ethical Charter on the Use of Artificial Intelligence in Judicial Systems and their Environment (Council of Europe 2018), <https://rm.coe.int/ethical-charter-en-for-publication-4-december-2018/16808f699c>.

¹⁷ Regulation (EU) 2024/1689, art. 14(4).

Fairness and non discrimination (Reg.6)	It names the protected grounds and the vulnerable groups, and applies the duty to design, training and deployment.	EU AI Act Article 10(2) requires systems to be examined for bias and biases to be mitigated. Brazil Article 8 ¹⁸ adds the missing consequence, namely that a system whose bias cannot be cured must be withdrawn.	A strong duty that is difficult to prove. There is no required bias test against the protected grounds, no fairness metric, and no stated consequence.
Accountability (Reg. 8)	It places responsibility on the officer and refuses to accept the AI output, the black box or the hallucination as an excuse.	EU AI Act Articles 12 and 26(6) require automatic logs and their retention for at least six months. Brazil Article 13(II) ¹⁹ requires a record of the degree of human supervision.	A well drafted rule that lacks the record needed to prove, after the event, that the judge reviewed rather than merely adopted the output.
Data protection (Reg. 10)	It treats sensitive judicial data as deserving the highest protection, and supports this with anonymisation and minimisation in Chapter VII.	EU AI Act Article 10 read with the GDPR, and Brazil Articles 4 and 7 ²⁰ , define the standard they invoke and apply it to AI processing.	The phrase highest standard is undefined, and the cross reference is to a DPDP Act that recognises no category of sensitive data.
Proportionality (Reg. 12)	It states that safeguards should rise with risk, and requires human review for the most sensitive uses.	EU AI Act Article 6 with Annex III(8)(a) and Article 9 grade risk and escalate obligations by tier. Brazil Chapter III and Article 14 ²¹ require an impact assessment for high risk systems.	Asserted but not yet operational. There is no risk framework, no definition of high risk, and no body charged with grading.
Regulation 15. Cyber security	It applies the confidentiality,	EU AI Act Article 15 requires accuracy,	Sound and convergent in its

¹⁸ Regulation (EU) 2024/1689, art. 10(2); CNJ Resolution No. 615/2025, art 8.

¹⁹ Regulation (EU) 2024/1689, arts. 12, 26(6); CNJ Resolution No. 615/2025, art 13(II).

²⁰ Regulation (EU) 2024/1689, art. 10; CNJ Resolution No. 615/2025, arts 4 and 7.

²¹ CNJ Resolution No. 615/2025, art 14.

	integrity and availability triad, with measures that are robust, layered and kept current.	robustness and cybersecurity. CEPEJ Principle 3 ²² requires a secure technological environment.	wording. The terms sensitivity and commensurate are undefined, and no security baseline is named.
--	--	--	---

3.1.1. Translation of the Principles into Operational Frameworks

The Apex Body is empowered to set mandatory standards under Regulation 24(1)(a), supported by the Technical Committee and the Centre of Research and Excellence (CoRE-AI). However, the Regulations do not specify the technical standards, metrics or methodologies that will govern implementation.

U.S. court-sector guidance issued by the National Center for State Courts, together with the NIST AI Risk Management Framework, provides useful operational reference points for performance, fairness, bias and robustness.²³ Brazil recommends adopting ISO and NIST standards as reference points.²⁴ These are settled, internationally recognised frameworks that already supply the measurable parameters the Draft now lacks.

Recommendation: The Regulations should require the Apex Body, within a fixed period of its constitution, to publish: (i) a risk classification framework for judicial AI systems; (ii) the technical standards and metrics to be applied; and (iii) a standardised impact assessment methodology. This would provide courts with a common operational framework and reduce implementation uncertainty.

3.1.2. Human Primacy and Judicial Independence (Regulation 4)

Regulations 4(1) to 4(3) collectively establish that AI must remain subordinate to human judgment, operate solely in an assistive capacity and never displace judicial authority. Taken together, this is one of the clearest statements of human primacy in a judicial AI instrument. It is more emphatic than the EU

²² CEPEJ Ethical Charter, Principle 3.

²³ Andrea L. Miller et al., National Center for State Courts, AI Readiness for the State Courts 4–5, 10–15 (2025), <https://www.ncsc.org/sites/default/files/media/document/AI%20Readiness-for-the-State-Courts-2025.pdf>; National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0) (Jan. 2023), <https://www.nist.gov/itl/ai-risk-management-framework>.

²⁴ Conselho Nacional de Justiça [CNJ], Resolução No. 615, de 11 de março de 2025, art. 31, parágrafo único, IV (Braz.), <https://atos.cnj.jus.br/atos/detalhar/6001>.

language of human oversight and the conduct based formulations of responsibility found in the United Kingdom²⁵ and Singapore²⁶.

However, human primacy cannot rest on formal authority alone. It requires judicial officers to have the practical ability to understand, scrutinise and, where necessary, reject AI-generated outputs. This is reflected in Article 14(4) of the EU AI Act, which treats human oversight as a set of capabilities rather than a purely supervisory role. Those responsible for oversight must understand the system's limitations, remain alert to automation bias, and be able to disregard, override or suspend its outputs where appropriate. Judicial training and impact assessment processes should therefore equip judicial officers not only to use AI systems, but also to exercise informed and independent judgment when evaluating their outputs.

Recommendation: To give practical content to human primacy, the impact assessment under Regulation 35 could require, for any system whose output informs a judicial determination, a short statement of how the reviewing officer is expected to understand, question and if necessary set aside that output. Pairing this with the training contemplated in Chapter VIII would help ensure that the authority reserved to the judicial officer is matched by the capacity to exercise it, and that oversight does not, in practice, settle into routine endorsement.

3.1.3. Fairness and Non-Discrimination (Regulation 6)

The principle of fairness is appropriately articulated but remains difficult to operationalise. While Regulations 35 and 38 require consideration of bias during impact assessments and audits, the Draft Regulations do not specify how fairness should be measured or what consequences follow where bias is identified. The only present footholds are the evaluation of risks of bias in the impact assessment under Regulation 35(3)(c) and the duty of audits to detect biases under Regulation 9(1).

Consistent with the publication of frameworks and standards recommendation in Section 3.1.1, a workable approach would require documented bias assessments prior to deployment, the use of recognised methodologies and fairness metrics, periodic revalidation, and a clear consequence where

²⁵ Courts and Tribunals Judiciary (England and Wales), Artificial Intelligence (AI) Guidance for Judicial Office Holders (updated 31 October 2025), <https://www.judiciary.uk/guidance-and-resources/artificial-intelligence-ai-judicial-guidance-october-2025/>.

²⁶ Supreme Court of Singapore, Guide on the Use of Generative Artificial Intelligence Tools by Court Users (2024), https://www.judiciary.gov.sg/docs/default-source/news-and-resources-docs/guide-on-the-use-of-generative-ai-tools-by-court-users.pdf?sfvrsn=3900c814_1.

discriminatory bias cannot be adequately mitigated. For instance, the EU AI Act requires systems to be examined for bias under Article 10(2), while ISO/IEC TR 24027 and the Measure function of the NIST framework provide established methodologies for assessing and monitoring bias. Without such measures, the obligation risks remaining aspirational rather than enforceable.²⁷

Recommendation: Consideration may be given to recording the bias assessment and its outcome in the AI Register, so that the position of a given system is transparent to those affected by it and capable of review over time.

3.1.4. Accountability (Regulation 8)

The provision rightly places responsibility on the human decision-maker and rejects reliance on system outputs as a defence. In doing so it addresses the failure that is now drawing judicial criticism elsewhere, namely the fabricated citation and the unverified output. This is a necessary safeguard. The duty to treat AI output as advisory and to verify it, under Regulation 8(3), tempered by a sensible pathway of deemed verification for administrative tools certified by the Secretariat, strikes a fair balance between rigour and workability.

To make this principle effective, it must be supported by a requirement to maintain operational records. Systems should record their outputs, the manner in which those outputs are used and the extent of human supervision. Such records are essential for audit, review and accountability.

Comparative frameworks reinforce this approach. The European Union treats logging as a foundational requirement. High-risk systems must generate logs automatically under Article 12, and deployers are required to retain them for a defined period under Article 26(6). Brazil similarly requires a record of the degree of human supervision that contributed to the outcome.

Recommendation: A requirement to generate and retain operational logs would render accountability under Regulation 8, auditability under Regulation 9 and explainability under Regulation 7 capable of verification, rather than resting on assertion alone.

A duty to generate and retain operational logs would be most effective if calibrated to risk rather than applied uniformly, so that the fullest record

²⁷ ISO and IEC, ISO/IEC TR 24027:2021, Bias in AI systems and AI aided decision making, <https://www.iso.org/standard/77607.html>. See also ISO/IEC 42001:2023 (AI management systems), ISO/IEC 23894:2023 (AI risk) and ISO/IEC 27001 (information security).

attaches to systems that inform adjudication and a lighter record suffices for routine administrative use.

It would also assist to specify, even at a high level, who may access these logs and for how long they are retained, so that they serve the purposes of audit and review under Regulations 9 and 38 while remaining consistent with the confidentiality of the proceedings to which they relate.

3.1.5. Data Protection and Privacy (Regulation 10)

Regulation 10 provides that sensitive judicial data shall be accorded the highest standard of protection in accordance with the Digital Personal Data Protection Act, 2023.²⁸ The DPDP Act, however, does not recognise a distinct category of sensitive personal data or prescribe graded standards of protection based on data type. Its framework is uniform, though it permits higher obligations to arise under other applicable laws. In that context, the Draft Regulations assume significance. It expressly carves out a category of sensitive judicial data and, read with Chapter VII, attaches to it a set of enhanced safeguards.

These safeguards are substantive. They require anonymisation before data is used for training and testing, restrict transfers to external systems without authorisation, prefer data-minimising system design, impose least privilege access, and mandate controlled or sovereign infrastructure for sensitive datasets. These safeguards, and their interaction with the DPDP Act and the broader legal framework governing data protection and cyber security, are examined in greater detail in Section 8 of this submission, which comments on Chapter VII of the Draft Regulations.

3.1.6. Proportionality (Regulation 12)

Regulation 12(2) provides for correspondingly heightened safeguards, including mandatory human review and independent oversight, for higher risk applications. However, the Draft Regulations contain no risk classification framework against which to apply that premise. Its only gradation is binary - an expedited pathway for low risk administrative tools under Regulation 34(4) and 46(7), and a single, undifferentiated impact assessment under Regulation 35 for everything else.

Comparative frameworks address this through structured risk tiering. The European Union classifies systems assisting judicial authorities as high risk under Annex III(8)(a), with obligations that escalate accordingly. Brazil operates a high and low risk categorisation under a risk classification annex in Chapter III,

²⁸ Digital Personal Data Protection Act 2023, No 22 of 2023 (India), <https://www.meity.gov.in/static/uploads/2024/06/2b1f0e9f04e6fb4f8fef35e82c42aa5.pdf>.

and only high risk systems must undergo an algorithmic impact assessment under Article 14 and enhanced, continuous governance under Article 13.²⁹ Both approaches provide clarity on how risk is defined, assigned and regulated. A similar structure would strengthen the Draft Regulations. The task of designing it may be assigned to the Apex Body under Regulation 24(1)(a), with technical support from the Centre of Research and Excellence and the Technical Committee.

In the judicial context, a four-tier structure would be workable:

1. **Lower risk** uses would include administrative functions that do not affect adjudication, such as scheduling, cause list preparation, defect scrutiny, transcription with human certification, and anonymisation for publication. These may remain subject to expedited approval, simplified assessment and periodic audit.
2. **Intermediate risk** tier should capture tools that shape the material placed before the court without determining outcomes. This would include legal research, summarisation, translation of pleadings, and evidence organisation. While not determinative, such tools influence the record and may introduce error or bias at an earlier stage. These uses should be subject to structured impact assessment, mandatory human verification before reliance, defined performance standards, periodic revalidation, and logging of use. Audit requirements may be calibrated accordingly.
3. **Higher risk** uses would include systems that directly affect adjudication or are capable of influencing findings of fact or law. These warrant comprehensive impact assessment, continuous validation, meaningful human oversight, independent audit and full traceability.
4. **Prohibited uses** under Regulation 20 should remain outside any approval framework.

It may assist the Committee to see how the comparator tiers translate to judicial tasks. Drawing on the logic of EU Annex III and the Brazilian annex, a workable schedule for Indian courts might place uses in the following four broad bands:

Indicative tier	Examples of court uses	Calibrated safeguards
Lower risk	Administrative and ministerial tools that do not touch the merits, such as scheduling, cause list preparation, defect scrutiny, transcription with human	Expedited approval, as the Draft already allows. Simplified impact assessment. Periodic audit. Human review of output.

²⁹ CNJ Resolution No. 615/2025.

	certification, and anonymisation for publication.	
Intermediate risk	Tools that shape the material placed before the court without determining outcomes, including legal research and summarisation, translation of pleadings, and evidence organisation. These influence the record and may introduce error or bias at an earlier stage.	Structured impact assessment. Mandatory human verification before reliance. Defined performance standards and fairness metrics. Periodic revalidation. Logging of use. Calibrated audit.
Higher risk	Systems that directly affect adjudication or are capable of influencing findings of fact or law, such as any tool whose output informs a finding, the assessment of evidence, or sentencing.	Comprehensive impact assessment. Continuous validation. Meaningful and capable human oversight. Independent audit. Full traceability and operational logging.
Restricted or prohibited	Uses already prohibited under Regulation 20, such as decision making by AI alone, risk scoring, and profiling of parties or witnesses.	Outside any approval pathway. Non derogable.

Two further clarifications would be necessary. First, the expression “integrity of judicial outcomes” should be defined to include any use capable of influencing findings of fact or law, the assessment of evidence or the exercise of judicial discretion. Second, the Draft Regulations should require that the depth of assessment, intensity of oversight and frequency of audit follow the assigned risk tier.

Recommendation: The Regulations should adopt a defined risk classification framework along the lines of the four tiers proposed above, with each tier carrying a pre-assigned set of obligations so that classification settles the intensity of assessment, oversight and audit without a fresh determination in each case.

3.1.7. Cyber Security (Regulation 15)

The provision is broadly aligned with international practice and reflects a sound understanding of baseline security principles. Its emphasis on confidentiality, integrity and availability is well placed. Our reservation mirrors the one under Regulation 10. The duty is qualified by the phrase commensurate with the sensitivity of the data and the nature of the court process. Yet neither element is defined. The Draft Regulations specifies no sensitivity levels, and commensurate has no scale against which proportionate security may be judged. The result is that the obligation, while correctly framed in principle, remains difficult to apply, measure or audit in practice.

Recommendation: The Regulations should define categories of data sensitivity. These may be linked to the definition of sensitive judicial data under Regulation 10 and aligned with the risk classification framework we have proposed under Regulation 12. A structured classification would ensure that security obligations correspond to the actual risk profile of the data and the process in which it is used.

The requirement of commensurate security should be anchored to a recognised technical baseline. Reference to established standards, such as the ISO/IEC 27001³⁰ control families or comparable frameworks, would convert a general obligation into one that is measurable, verifiable and capable of audit. It would also promote consistency across courts and reduce interpretive uncertainty.

³⁰ ISO and IEC, ISO/IEC 27001:2022, Information security, cybersecurity and privacy protection - Information security management systems - Requirements, <https://www.iso.org/standard/27001>.

4. Chapter III: Permitted, High-Risk and Prohibited Uses

4.1 Overall Assessment

The Regulations adopt a tiered approach to AI deployment by distinguishing between permissible uses, high-risk applications requiring enhanced safeguards, and prohibited uses. This is an appropriate framework for judicial contexts, as it seeks to harness AI for administrative efficiency, accessibility, legal research and case management, while preserving judicial autonomy and protecting individual rights.

Several aspects of the framework, however, require further refinement to ensure internal coherence and operational certainty. The distinction between permitted and prohibited uses is not always clear, key concepts relating to risk categorisation remain underdeveloped, and certain provisions create structural tensions that may lead to inconsistent application across court jurisdictions.

4.1.1. Identification of AI Use

Regulation 18 entrusts the Appropriate Authority with determining the court processes in which AI systems may be utilised. While this provides the necessary institutional responsibility, the Regulations do not indicate how this determination is to be undertaken on an ongoing basis.

Recommendation: Consider outlining a process for the Appropriate Authority to periodically review and update its determination of permissible AI uses, supported by defined reporting mechanisms and linked to the Apex Body's standard-setting function under Regulation 24. This would ensure that decisions on the use of AI remain current and are based on regular institutional oversight.

4.1.2. Regulatory Certainty in the Permissible Uses Framework

The drafting of Regulation 19 creates uncertainty regarding the approval pathway for new AI applications. While Regulation 19(1) describes the list of permissible uses as “illustrative and not exhaustive”, Regulation 19(2) requires prior written approval for any use that is not expressly enumerated or otherwise approved. As drafted, the illustrative character of Regulation 19(1) serves little practical purpose. If all unenumerated uses require approval, the framework effectively operates as a closed-list system rather than an open-ended catalogue of examples. Retaining the term “illustrative” may therefore create unnecessary ambiguity regarding the scope of institutional discretion.

Recommendation: The Regulations would benefit from adopting one of two approaches. Either the list should be treated as exhaustive, with all additional uses requiring formal approval, or the Regulations should establish objective criteria that allow functionally similar low-risk uses to proceed without requiring a separate approval process in every instance.

4.1.3. Absolute and Conditional Prohibitions

Regulation 20 describes all prohibited uses as “absolute and non-derogable.” This is appropriate for prohibitions that admit no exception. However, the prohibition on AI in adjudicative and sentencing functions is not simply a prohibition on AI adjudication without HITL oversight. The Draft Regulations prohibit AI from performing the adjudication function itself. HITL operates as a requirement governing AI assistance in adjudicative processes, not as a condition whose satisfaction would permit AI-driven adjudication. This distinction should be clearly understood when implementing the framework, as framing HITL as a gateway to AI adjudication would be inconsistent with the intent of Regulation 20.

Two other provisions within Regulation 20 create a structural inconsistency. The prohibition on using personal data for AI training is qualified by the possibility of prior approval, and the prohibition on AI-enabled surveillance is subject to authorisation under applicable law. These are conditional restrictions rather than absolute ones, and describing them alongside genuinely non-derogable prohibitions may generate interpretive difficulty in practice.

A related concern arises under Regulation 8(3), which permits a court officer to waive the verification requirement on recording reasons in writing. Where the waiver touches an adjudicatory function, this threshold may warrant closer consideration.

Recommendation: The Regulations should distinguish expressly between prohibitions that are genuinely non-derogable and those that are conditional on prior authorisation or applicable legal authority, grouping them separately within Regulation 20 or through a clarificatory provision. The threshold for waiving the verification requirement under Regulation 8(3) should be addressed in the Apex Body's implementation guidelines, specifying that where a waiver touches an adjudicatory or quasi-adjudicatory function, a higher standard of justification applies than where it concerns an administrative function.

4.1.4. Opaque Systems, Outcome Prediction and Behavioural Profiling

Regulation 7(3) permits opaque AI systems in high-risk contexts subject to heightened scrutiny, while Regulation 20(1)(e) prohibits their use wherever lawful rights or personal liberty may be materially affected. The two provisions cover overlapping categories of systems but prescribe different consequences, and their relationship is not explained.

Separately, the prohibitions on behavioural profiling and outcome prediction in Regulation 20 leave a boundary question unaddressed. An AI legal research tool that surfaces statistically likely outcomes based on precedent patterns may not be readily distinguishable from prohibited outcome prediction. The Regulations draw no line between the two.

Recommendation: Clarification of the relationship between Regulations 7(3) and 20(1)(e) would assist consistent application. Guidance on the distinction between permitted AI-assisted legal research and prohibited outcome prediction would also be helpful, particularly for courts and technology providers seeking to deploy research tools within the permitted framework.

4.1.5. Risk Scoring

The prohibition on risk scoring under Regulation 20(1)(d) applies "for any purpose in Court processes." In criminal justice contexts, this prohibition should remain firm. AI scoring for bail, sentencing, witness credibility, recidivism, flight risk, personal liberty, or prediction of future conduct should remain outside any approval pathway. The definitional question arises only at the margins: whether the same prohibition is intended to cover civil, commercial or financial analytics that do not score an individual's future conduct in a manner that affects rights. The National Company Law Tribunal, for instance, is a forum in which AI-assisted analysis of insolvency data or debt-recovery patterns may arise as an administrative or analytical use case. The Regulations would benefit from clarifying this boundary without diluting the absolute prohibition in criminal justice and liberty-related contexts.

Recommendation: If the prohibition is intended primarily to address criminal justice, personal liberty and future conduct assessments, this should be stated expressly, and preserved as a non-derogable safeguard. Separate consideration may be given to civil or commercial analytics only where the tool does not score an individual's future conduct in a way that

affects rights, liberty, credibility, entitlement or access to judicial relief.

4.1.6. AI Chatbots and Conversational Assistants

Regulation 19(1)(f) permits conversational AI assistants to interact directly with litigants and other stakeholders, subject to ‘human oversight of functioning.’ As presently framed, the provision does not indicate the nature or extent of such oversight. It is unclear whether this requires review of individual interactions, periodic supervision, or only system-level monitoring.

This assumes particular importance in the judicial context. For a litigant, especially one appearing without legal representation, a conversational system may function as the first point of engagement with the court. Incorrect guidance on limitation, jurisdiction or filing requirements may cause procedural consequences that are not easily reversed. Further, pre-deployment accuracy testing should be conducted in the language of the court in which the system is to be deployed, given that AI tools calibrated primarily on English and Hindi judicial text may perform unevenly in regional language environments. The Regulations also do not address whether a litigant who suffers harm from relying on defective chatbot guidance has a remedy under the grievance mechanism in Chapter IX.

Recommendation: The level of human oversight required under Regulation 19(1)(f) may benefit from further specification. Minimum safeguards worth considering include pre-deployment accuracy testing, clear disclosure that the user is interacting with an automated system, and a real-time escalation pathway to a human court officer. The Regulations may also wish to clarify whether reliance on defective AI guidance gives rise to a remedy under Chapter IX.

4.2 Enforcement of the Prohibited Uses Regime

The most significant gap in Chapter III is the absence of defined consequences for violations of Regulation 20. Regulation 21 provides that violations are referred to the AI Committee for “such remedial measures as it deems appropriate.” For prohibitions described as absolute and non-derogable, this is a notably open-ended response. The Regulations do not address whether violations attract disciplinary consequences for responsible officers, whether affected litigants may challenge proceedings tainted by prohibited AI use, or whether orders or convictions made in such proceedings are subject to review.

The enforcement gap is most acute in courts where the consequences of prohibited AI use would be most severe. POCSO Courts and Fast Track Special Courts operate under statutory confidentiality regimes, including Section 33(7) and Section 37 of the POCSO Act and Section 366 of the BNSS, that independently criminalise unauthorised disclosure of sensitive data. A violation of Regulation 20 in these courts, may simultaneously constitute a criminal offence under the POCSO Act. Regulation 21's referral of violations to the AI Committee for 'such remedial measures as it deems appropriate' does not address how violations that independently engage statutory offences are to be treated, or whether the AI Committee's remedial powers operate concurrently with or in displacement of those statutory consequences. The Regulations should address this interaction expressly, at minimum through guidance issued by the Apex Body under Regulation 24.

Recommendation: The Regulations may benefit from specifying the consequences of violating Regulation 20, including disciplinary consequences for responsible officers and whether affected proceedings may be challenged or reviewed. Clarity on these points is essential to giving the prohibited uses regime the practical force that its framing as absolute and non-derogable implies.

5. Chapter IV: Institutional Architecture and Operational Requirements



5.1 Overall Assessment

The proposed institutional architecture establishes a comprehensive governance framework for judicial AI, centred on the Apex Body, specialised committees, CoRE-AI, AI Committees at each High Court, and AI Secretariats. While the framework reflects an effort to balance judicial oversight with technical expertise, its current design risks creating fragmentation, overlapping mandates, and operational bottlenecks that may undermine effective implementation.

5.1.1. Independence and Decision Making

At the apex level, the composition of the Apex Body raises questions about institutional independence and decision-making robustness. Executive technical

coordination may be valuable, particularly on cybersecurity, infrastructure, standards and public digital systems. However, where the Apex Body exercises approval or oversight functions affecting judicial AI deployment, the scope of executive participation should be carefully bounded to preserve judicial autonomy. The inclusion of a MeitY representative as a voting member should therefore be accompanied by clearly articulated safeguards governing participation in approval, oversight and standard-setting decisions.

Equally significant is the absence of prescribed quorum requirements, voting thresholds, or procedural safeguards for decision-making within the Apex Body. Given its central role in approving AI systems and shaping judicial AI policy, procedural specificity would strengthen predictability, continuity and perceived institutional legitimacy.

Tribunals and statutory Commissions, brought within the framework through deemed AI Committee status under Regulation 33(4), have no dedicated seat on any of these bodies. Their specialised processes and institutional contexts may therefore not be adequately reflected in standards set at the Apex level.

Recommendation: Executive technical coordination may be valuable, particularly on cybersecurity, infrastructure, standards and public digital systems. However, where the Apex Body exercises approval or oversight functions affecting judicial AI deployment, the scope of executive participation should be carefully bounded to preserve judicial autonomy. The Regulations may therefore specify whether the MeitY representative's role is advisory or voting in different classes of decisions, and should prescribe quorum and decision making thresholds. Consideration may also be given to a mechanism by which the operational contexts of tribunals and statutory commissions are represented in Apex Body deliberations.

5.1.2. Fragmentation and Overlapping Mandates

A broader challenge lies in the distribution of functions across the various committees and governance entities. Several bodies are assigned responsibilities that intersect substantially without a corresponding coordination mechanism. Under Regulations 28 and 32, both the Technical Committee and CoRE-AI engage in technical evaluation of AI systems, without a clear functional distinction between research, capability development, and deployment-focused assessment. Similarly, the Judicial Committee and the Case and Data Management Committee both have responsibilities touching adjudication-support tools. The Regulations do not specify decision-making primacy, escalation pathways, or mechanisms for resolving conflicting recommendations across these bodies.

Recommendation: A clearer functional demarcation between bodies with overlapping mandates would reduce duplication and improve decision quality. In particular, the respective roles of CoRE-AI and the Technical Committee may benefit from being distinguished expressly, with CoRE-AI focused on research and capability development and the Technical Committee on deployment-oriented assessment and approval support.

5.1.3. Technical Expertise in Oversight Structures

The Apex Body is responsible for approving AI systems and setting technical standards under Regulation 24, but the Regulations do not specify how technical assessments are to be weighted in final decisions or how expert body recommendations are to be reflected in Apex Body deliberations.

AI Committees at the High Court level, which exercise important approval and oversight functions under Regulation 33, have no technical expertise requirement in their standing composition. By comparison, Brazil's Sinapses platform is administered by the National Council of Justice (CNJ), a body with dedicated technical staff and an independent governance structure, rather than by an officer-level secretariat with no prescribed technical capacity.³¹ Decisions on AI adoption necessarily involve assessments of model performance, limitations, and risk that extend beyond judicial expertise alone.

Recommendation: Structured technical input into both Apex Body and AI Committee deliberations would strengthen decision quality while preserving judicial authority over final outcomes. AI Committees may also benefit from access to designated technical advisers, without requiring formal membership, to support their oversight functions.

5.1.4. Capacity and Operational Sustainability

The Secretariat is assigned an extensive portfolio: AI Register (Reg 37), AI Incident Database (Reg 39), receiving reports from the supervising officer (Reg 39(3)), initiating remedial measures (Reg 39(4)), coordinating audits (Reg 38), overseeing Controlled Environment Testing (Reg 36), developing training programmes (Reg 49), granting expedited approvals (Reg 34(4)), and activating emergency fall-back protocols (Reg 42). This concentration of operational, administrative, and quasi-regulatory functions in a single institution with no

³¹ CNJ Resolution No. 615/2025 and CNJ documentation on Sinapses, <https://www.cnj.jus.br/sistemas/sinapses/>.

defined resource base risks creating a structural bottleneck as AI adoption scales across the judicial system.

Recommendation: The Regulations may benefit from specifying minimum staffing and technical capacity requirements for the AI Secretariat, or providing for technical support arrangements to supplement its in-house capacity.

5.1.5. AI Content Verification Authority

Although Regulation 44 is placed in Chapter V, this section addresses it here because the principal concerns are institutional design, composition, reporting line and operational mandate. The same recommendations should also be read with the transparency and disclosure obligations in Regulation 43 and the annual reporting obligation in Regulation 45.

Regulation 44 requires the Appropriate Authority to constitute an AI Content Verification Authority, but does not specify its composition, governance structure, reporting line, or implementation timeline. This creates a disconnect between an operative legal obligation under Regulation 3(y), which prohibits the filing of GenAI-generated content without verification through a designated centralised mechanism, and the institutional infrastructure needed to discharge it.

The scope of this obligation is broader than it may appear. "Court process" covers all adjudicatory, quasi-adjudicatory and administrative functions, and GenAI extends to text, images, audio, video and code. The verification obligation therefore applies equally to AI-generated content produced by court systems, by lawyers, and by litigants. These are materially different contexts involving different actors, accountability relationships and evidentiary consequences, and the Regulations apply a single standard to all three without acknowledging this distinction.

A more fundamental difficulty is that the verification function contemplated by Regulation 44 may not be reliably achievable through a single post-hoc detection mechanism. Existing methods such as provenance standards, watermarking and fingerprinting can assist, but they operate differently and have different limits. Content Credentials, for instance, are designed to establish the origin and edit history of digital content rather than certify the truth of the content itself.³² The draft EU Code of Practice on AI-generated content similarly concludes that no single detection technique is sufficient and that multiple

³² Coalition for Content Provenance and Authenticity (C2PA), Content Credentials, <https://c2pa.org/>.

approaches must be combined.³³ The EU AI Act addresses this by placing disclosure and marking obligations on AI system providers at the point of generation, rather than creating a central post-hoc verifier.³⁴

The Authority should therefore not be expected to certify truth or detect every instance of AI-generated content. Its more workable role would be to set standards for disclosure, provenance, verification trails, citation checking, authenticity protocols, record retention and consequences for false or misleading filings. This would align the institution with what current verification technologies can support, while preserving the Court's ability to require accountability from users of GenAI material.

Recommendation: The Regulations should distinguish between verification obligations applicable to court-generated, lawyer-generated and litigant-generated content, with differentiated standards for each. The Authority's composition and timeline should be specified before commencement.

Institutional Framework: Functions and Key Concerns

Body	Role (as drafted)	Key concerns
Apex Body (Reg. 22–24)	National standard-setter: minimum mandatory standards, system approvals, policy, liaison with MeitY/CERT-In/Data Protection Board	• Composition is heavily judicial (6 of 12 members are judges). • Representation of finance and cybersecurity expertise is limited. • Unclear to what extent Technical Committee's recommendations bind or influence Apex Body decision-making on approvals.
CoRE-AI (Reg. 32)	Research and evaluation body; Conducts research, evaluates AI tools and prototypes, provides technical advice	• Significant functional overlap with the Technical Committee (Reg. 27) • Risks parallel assessments leading to divergent recommendations to the Apex Body.

³³ European Commission, Code of Practice for General-Purpose AI Models and AI-generated Content, <https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content>.

³⁴ EU AI Act, art 50, read with the draft EU Code of Practice on AI generated content (December 2025), <https://artificialintelligenceact.eu/transparency-rules-article-50/>.

AI Content Verification Authority (Reg. 44)	Oversees verification standards and protocols for GenAI content in court processes	• Unclear whether this extends to content generated by courts, litigants, and lawyers equally. • Authority's composition, governance and implementation timeline are unspecified, leaving the obligation operationally unanchored.
AI Committee, SC and each HC (Reg. 33)	Jurisdictional governance body responsible for approving systems, monitoring compliance, reviewing incidents, and issuing supplementary directions	• Verification obligation applies to all GenAI content filed or produced before a court, but does not distinguish between content • Absence of structured technical decision thresholds raises questions about consistency in evaluating complex AI systems under TEIAs across jurisdictions.
AI Secretariat, each HC (Reg. 34)	Primary operational and administrative body maintaining the AI Register, incident database, audit functions, training, fallback protocols, expedited approvals, and support for CET processes	• Heaviest operational portfolio in the framework but has no prescribed staffing, technical expertise, or resourcing requirements. • Expedited approval powers allow the Secretariat to bypass the standard AI Committee approval pathway without defined triggering conditions, safeguards etc.
Tribunals & Commissions (Reg. 33(4))	Brought within the framework by 'deemed' AI Committee status under the relevant HC or Supreme Court AI Committee	No dedicated representation on any decision-making body. Specialised tribunal processes (NCLT, ITAT, CAT, etc.) may not be well understood by a general-purpose HC AI Committee.

5.2. Operationalisation Requirements

The absence of deadlines is compounded by a sequencing failure: several implementing instruments are preconditions for other obligations that run from commencement. The AI Register is an operative obligation from the moment the Draft Regulations take effect but its format is undated. The annual audit cycle (Regulation 38) begins at or shortly after commencement but the audit protocol does not yet exist. The legacy-system review (Regulation 41) must be completed within one year but the compliance standard against which legacy systems will be reviewed has no deadline.

The EU's experience is instructive. The EU AI Act entered into force on 1 August 2024 and follows a phased implementation timeline. In June 2026, the Council gave final approval to the Omnibus VII simplification package, including delayed application dates of 2 December 2027 for stand-alone high-risk AI systems and 2 August 2028 for high-risk AI systems embedded in products.³⁵ This experience illustrates the importance of ensuring that implementing standards, assessment protocols and SOPs are in place before substantive obligations take effect.

The table below maps all 28 provisions requiring SOPs, standard formats, assessment frameworks, protocols, or other implementation guidance.

#	Instrument / Standard / Protocol Required	Regulation(s)	Responsible Body	Indicative Timeline
1	Minimum Mandatory Standards	Reg. 24(a)	Apex Body	None
2	Technical & Ethical Impact Assessment (TEIA) - standard format	Reg. 35(2) Reg. 19(1)(b)	Appropriate Authority (Apex Body, for system-wide consistency)	6 months
3	Class-Based Certification Standard (administrative tools)	Reg. 8(3) 2nd proviso	AI Secretariat	None
4	Approval Criteria for Non-Enumerated AI Uses	Reg. 19(2)	Appropriate Authority	None
5	Prohibited-Use Violation (reporting, inquiry, remedial action)	Reg. 20 Reg. 21	AI Committee / AI Secretariat	None
6	Prior-Approval Protocol for Personal Data in Training/ Testing	Reg. 20(1)(a)	Appropriate Authority	None
7	Controlled Environment Testing (CET) Protocol	Reg. 36	AI Secretariat	None

³⁵ Dentons, Digital Omnibus on AI: Provisional compromise reshapes the EU AI Act (Part 2) (10 June 2026), <https://www.dentons.com/en/insights/articles/2026/june/10/digital-omnibus-on-ai-provisional-compromise-reshapes-eu-ai-act-part-2>.

8	AI Register - prescribed form	Reg. 37	Appropriate Authority (Apex Body)	None
9	Audit Procedure & In-House Audit Standards	Reg. 38	Appropriate Authority / AI Secretariat	None
10	AI Incident Database - structure, taxonomy, escalation matrix	Reg. 39	AI Secretariat	None
11	Legacy System Compliance Review Standard	Reg. 41	AI Secretariat	Review to happen in 1 year
12	Emergency & Fall-Back Protocol	Reg. 42	Each High Court + AI Secretariat	None
13	AI Disclosure Annexures I & II & Synthetic Data Disclosure Format	Reg. 43(3) Reg. 43(5)	Appropriate Authority	Not published
14	AI Content Verification Authority - constitution, mandate, GenAI verification protocol	Reg. 3(1)(y) Reg. 44	Appropriate Authority (likely Apex Body)	None
15	Annual Transparency Report - prescribed format	Reg. 45	Apex Body	None
16	Procurement Procedure	Reg. 46(3)	Chief Justice	None
17	Vendor Evaluation Framework (technical, legal, ethical, security, financial)	Reg. 46(2)	Appropriate Authority	None
18	Standard AI Procurement Contract - mandatory clause templates	Reg. 46(4)(a)–(l)	Chief Justice / Appropriate Authority	None

19	Data Minimisation Assessment Principles	Reg. 48(3)	Appropriate Authority	None
20	Anonymisation Standard for Court Data	Reg. 20(1)(a) Reg. 48(4)	Appropriate Authority	None
21	Cybersecurity Audit Framework	Reg. 48(5)	AI Secretariat	None
22	Access-Control Protocols (least privilege / need-to-know)	Reg. 48(6)	Appropriate Authority	Annual review specified
23	Inclusivity Impact Assessment & Digital Access Protocol	Reg. 13(2)	Appropriate Authority	None
24	Centralised Record of AI Tools and Evaluations	Reg. 32(3)(c)	CoRE-AI / Apex Body	None
25	Grievance Adjudication Procedure & Complaint Format	Reg. 52(2)	Each High Court	None

Recommendation: The Draft Regulations would benefit from a Schedule listing the relevant instruments identified, the body responsible, and a deadline from commencement sequenced by operational dependency. Instruments that are preconditions for other obligations must precede the obligations that depend on them. The Draft Regulations may also benefit from either a dedicated provision clarifying which functions fall within the remit of the Apex Body notwithstanding the general definitional split, or from refining the definition itself to distinguish between the Apex Body acting in its standard-setting capacity and the Apex Body or AI Committees acting in their approval and oversight capacities.

6. Chapter V: Oversight, Audits and Incident Management

6.1 Overall assessment

Chapter V is the operational core of the Draft Regulations, bringing together pre-deployment impact assessment, controlled testing, an AI Register, periodic audits, an incident database, fallback protocols, disclosure obligations and annual reporting. Few judicial systems have attempted a framework of comparable breadth. The closest analogue is Brazil's Sinapses platform, which provides a centralised environment for the registration, testing, audit and, where necessary, decommissioning of AI systems used by courts.³⁶

An institutional architecture of this kind requires clear parameters, defined standards and reliable records if it is to function in daily practice and sustain meaningful oversight.

6.1.1. Technical and Ethical Impact Assessment (TEIA)

The impact assessment is the central gateway and we welcome its inclusion. As presently framed, however, Regulation 35(3) identifies the matters to be evaluated without prescribing the manner of evaluation, while Regulation 35(2) defers the format of the TEIA for a period of up to six months. The practical feasibility of the TEIA is also contingent on the quality and completeness of the court data on which it draws. As of December 2025, approximately 637 crore pages had been digitised under e-Courts Phase III against a target of 3,108 crore, representing around 20.5% of the target.³⁷ Where digitisation is incomplete or the underlying data is inconsistently structured, the representativeness assessment required under Regulation 35(3) cannot be conducted in any meaningful sense. The Apex Body should be required to publish minimum data quality standards that must be satisfied before a TEIA for any proposed AI system can be treated as valid

The EU's Fundamental Rights Impact Assessment offers a useful reference. It requires identification of affected persons, specific risks of harm, human oversight measures to be implemented, and steps to be taken if risks materialise.³⁸ It operates as a living instrument, updated as circumstances change. Brazil's framework adds participation and continuity, treating impact assessments as ongoing rather than one-time exercises.

³⁶ CNJ Resolution No. 615/2025.

³⁷ Press Information Bureau, Ministry of Law & Justice, eCourts Project Phase III (Aug. 1, 2024), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2040232>; Press Information Bureau, Ministry of Law & Justice, Digital Infrastructure in District Courts (Feb. 12, 2026), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2227220>.

³⁸ EU AI Act, art 27, <https://artificialintelligenceact.eu/article/27/>.

Recommendation: The TEIA format should be published before or simultaneously with commencement, not deferred to six months after. The format should require identification of affected persons and specific harms, articulation of oversight measures, and recording of mitigation and grievance pathways. The TEIA should operate as a continuing obligation with defined triggers for review, rather than a one-time pre-deployment exercise.

6.1.2. Controlled Environment Testing

Regulation 36(1) provides that the Appropriate Authority “may” direct Controlled Environment Testing (CET) prior to full-scale deployment. As currently framed, this is entirely discretionary and the Regulations do not indicate which categories of systems should ordinarily be subject to it.

A more structured approach would be to link this requirement to the risk classification framework under Regulation 12. For systems classified as high risk, particularly those with adjudicatory-adjacent functions or those capable of influencing the judicial record, CET should be mandatory rather than discretionary. At the same time, the framework should retain flexibility for lower and intermediate risk systems. This approach would ensure that CET is applied in a manner that is both proportionate and predictable, while preserving the ability to respond to emerging use cases.

Recommendation: CET should be mandatory for high-risk systems as classified under Regulation 12, with the Appropriate Authority retaining discretion for lower-risk applications based on defined criteria.

6.1.3. Audits

Regulation 38 requires periodic audits of all court AI systems at intervals not exceeding one year. Regulation 38(2) provides that audits shall be conducted “in-house”, meaning on court premises, and that source code, algorithms, datasets and architectural information may not be shared with any third party or private entity for an audit conducted outside those premises.

The provision does not specify who is responsible for conducting the audit. It is unclear whether the audit is to be conducted by the court or the AI Secretariat, by the vendor or AI service provider on court premises, or by the vendor independently with a report submitted to the court. A meaningful technical audit of an AI system, covering model performance, bias testing, data quality, and architectural integrity, requires specialist technical expertise that court

institutions, as presently staffed, do not possess. Each of these models carries different implications for independence, technical depth and accountability.

Recommendation: Regulation 38 may benefit from clarifying who is responsible for conducting audits and the minimum qualifications or technical capacity required of those conducting them. This clarity is essential to ensuring that the audit obligation operates as a substantive accountability mechanism rather than a formal compliance exercise.

6.1.4. AI Incident Database and Interface with CERT-In

Regulation 39 requires an AI Incident Database to be maintained at every AI Secretariat, with incidents reported immediately to the Secretariat and without delay to the Appropriate Authority. Under the CERT-In Directions of 2022³⁹, issued under section 70B of the Information Technology Act, 2000,⁴⁰ government organisations, including courts, are required to report specified cyber security incidents within six hours of becoming aware of them. Several AI incidents, particularly those involving data breaches, unauthorised access or system compromise, will engage both obligations simultaneously.

The present formulation, which requires reporting immediately to the AI Secretariat and without delay to the Appropriate Authority, is less precise than the CERT-In directive and does not engage with it. This raises questions as to whether and when parallel reporting obligations are triggered, and from what point in time.

Recommendation: The Regulations should state expressly that AI incidents which also constitute cybersecurity incidents must be reported to CERT-In in accordance with the 2022 Directions, and should adopt a severity classification so that external reporting obligations are triggered only for material incidents. The interaction between Regulation 39 and the DPDP Act's breach notification obligations should also be addressed, as noted in section 8 of this submission.

³⁹ Indian Computer Emergency Response Team (CERT-In), Directions Relating to Information Security Practices, Procedure, Prevention, Response and Reporting of Cyber Incidents for Safe and Trusted Internet (28 April 2022), https://www.cert-in.org.in/PDF/CERT-In_Directions_70B_28.04.2022.pdf.

⁴⁰ Information Technology Act 2000, No 21 of 2000 (India), https://www.meity.gov.in/static/uploads/2024/03/IT-Act-Rules_2000_0.pdf.

6.1.5. Review of Existing System

The AI tools currently deployed across the Indian judiciary perform materially different functions and present different levels of risk. Systems such as SUVAS support translation, SUPACE and LegRAA assist with legal research and document analysis, while TERES and Adalat.AI are used for transcription of court proceedings.⁴¹ These systems differ in the nature of the data they process, the degree of human oversight they require and their potential impact on judicial decision-making. A uniform compliance review under Regulation 41 can therefore only be meaningful if it is conducted against clearly defined standards and assessment methodologies tailored to the type of AI system being evaluated.

Regulation 41 requires AI systems already in use to be reviewed for compliance within one year of commencement. However, the compliance standard against which they are to be assessed has no publication deadline of its own. A one-year review window with no fixed benchmark to apply from the outset risks producing inconsistent outcomes across jurisdictions.

Recommendation: The compliance standard under Regulation 41 should be required to be published within the same timeframe as the Minimum Mandatory Standards under Regulation 24, so that the one-year review period operates against a defined and uniformly applicable benchmark from commencement.

6.1.6. Emergency and Fall-Back Protocol

This is a carefully conceived provision that reflects a mature understanding of operational resilience. The requirement that each High Court establish and periodically test an emergency and fall-back protocol, ensuring continuity of essential court processes through manual or alternative means, mirrors approaches seen in the regulation of critical infrastructure and financial systems. It operates as a practical safeguard for the administration of justice, ensuring that proceedings are not disrupted by system failure. It also reflects a broader institutional commitment that the functioning of the court must not be contingent on the continuous availability of technology.

⁴¹ Press Information Bureau, Ministry of Law & Justice, Use of Artificial Intelligence in Legal Field (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202159&lang=1®=3>; High Court of Kerala, Office Memorandum on Use of Adalat.AI Speech-to-Text Transcription Tool (Oct. 27, 2025), <https://cdnbbsr.s3waas.gov.in/s3ec01dc6a7e655d7e5840e66733e9ee67/uploads/2025/10/2025102791.pdf>.

Recommendation: Two refinements may be considered. First, the frequency of testing and acceptable recovery times may be linked to the risk classification of the system under Regulation 12, with higher-risk systems subject to more frequent testing. Second, the existence of a tested fallback protocol may be made a condition for approval of higher-risk systems under Regulation 35, rather than operating as a parallel obligation.

7. Chapter VI: Procurement and Private Sector Engagement

7.1 Overall Assessment

The oversight mechanisms examined in the preceding section presuppose that the AI systems subject to them have been procured through a governance framework that introduces appropriate safeguards at the point of acquisition. Chapter VI of the Draft Regulations addresses this through a detailed vendor engagement regime.

Regulation 46 establishes a comprehensive procurement and vendor governance framework. It requires prior approval before engaging any private entity under Regulation 46(1), mandates a multi-factor evaluation of technical capability, legal compliance, ethical standards, data security and financial standing under Regulation 46(2), and prescribes detailed contractual safeguards under Regulation 46(4). Substantively, the Regulation reflects a strong protective and court-centric procurement model. However, the framework adopts a largely uniform approach to vendor governance, applying extensive safeguards across all categories of AI procurement without sufficiently differentiating between low-risk and high-risk systems. This may create unnecessary administrative burdens and reduce proportionality in implementation.

7.1.1. Evaluation and Contractual Terms

Regulation 46(2) applies the same multi-factor evaluation criteria to all AI procurements regardless of the risk profile of the system being acquired. Regulation 46(4) similarly prescribes a uniform set of mandatory contractual terms. Applying the same level of due diligence to a low-risk administrative tool, such as a scheduling assistant, as to a system processing sensitive judicial data or supporting adjudicative functions may impose disproportionate compliance costs and create unnecessary barriers, particularly for smaller vendors offering lower-risk applications.

Recommendation: Consideration may be given to calibrating both the intensity of vendor evaluation under Regulation 46(2) and the applicable contractual obligations under Regulation 46(4) to the risk profile of the system being procured, aligned with the risk classification framework under Regulation 12. This would preserve rigorous safeguards for high-risk systems while reducing administrative burden where the risks are limited.

7.1.2. Sovereign and On-Premise Deployment

Regulation 46(4)(j) requires on-premise or sovereign-cloud deployment for AI systems processing sensitive judicial data. However, the provision raises three difficulties.

First, it applies uniformly to all systems processing sensitive judicial data without regard to the degree of exposure or functional use, which may impose identical infrastructure requirements on systems with very different risk profiles.

Second, the terms "sovereign" and "on-premise" are not defined. It is unclear whether sovereignty refers to physical data localisation within India, legal control over the data, or protection from foreign jurisdictional access through vendor or cloud provider structures. Physical hosting within India does not automatically address the last of these, as a vendor's corporate structure may still expose data to foreign legal compulsion.

Third, on-premise deployment does not always represent the most secure option. Certified cloud environments can offer stronger resilience, redundancy and security controls than institutionally managed on-premise infrastructure, particularly in courts with limited technical capacity.

The practical dimension of this tension is most acute in POCSO Courts and Fast Track Special Courts that process child identity records, witness testimony, and medical examination reports under statutory confidentiality regimes that make unauthorised disclosure a criminal offence.⁴² Any AI tool used for transcription or translation in these courts will engage the Regulation 46(4)(j) requirement. These are also courts where on-premise infrastructure is typically least available and operate with contractual personnel whose continuity cannot be guaranteed. The Regulation designed to provide the strongest protection for the most sensitive data may therefore in practice exclude AI tools from precisely the courts that most need efficiency improvements, unless a realistic pathway for meeting the sovereign or on-premise requirement is specified.

Recommendation: The terms "sovereign" and "on-premise" may benefit from being defined, with clarity on whether sovereignty requires physical localisation, legal control, or both. Deployment requirements may also be aligned with the risk classification framework under Regulation 12 rather than applied uniformly. Certified and audited cloud infrastructure already operating within the Indian public digital ecosystem, such as MeitY-empanelled and STQC-audited environments under the GI Cloud

⁴² Protection of Children from Sexual Offences Act, No. 32 of 2012, §§ 23, 33(7), 37 (India); Bharatiya Nyaya Sanhita, No. 45 of 2023, § 72 (India).

framework⁴³, may be recognised as compliant alternatives where they meet the relevant security standards.

7.1.3. Intellectual Property

Regulation 46(9) seeks to ensure that courts retain appropriate rights over AI tools developed using court data or court resources. While the policy objective is sound, the provision adopts a simplified approach to AI development and intellectual property. The provision does not distinguish between different forms of contribution, such as the use of court data, court infrastructure, funding, or other public resources, nor does it indicate the threshold at which such contributions give rise to ownership or licensing rights. These are materially different legal arrangements and may require clearer delineation. It also extends the claim to the resulting tool and its outputs without clarifying the intended scope of these rights, particularly where AI systems are trained on multiple datasets or continue to evolve through subsequent fine-tuning and updates.

From an implementation perspective, the provision may also benefit from greater clarity on how the ownership or licence requirement is intended to operate alongside existing intellectual property law. Computer programs are protected as literary works under the Act, and where AI systems are procured under a contract for services, copyright will ordinarily remain with the vendor unless otherwise assigned or licensed by contract. The practical operation of Regulation 46(9) may therefore depend on the intellectual property arrangements agreed during procurement, and the relationship between the Regulation and the default ownership position under the Copyright Act.

Recommendation: Consider clarifying the scope of "court data" and "court resources", the threshold for triggering ownership or licensing rights, and the extent to which these rights apply to evolving AI systems and their outputs. Clarity is also needed on whether the court's ownership or licence claim persists where an AI system is subsequently fine-tuned or retrained on data other than court data, and if so on what basis.

⁴³ Press Information Bureau, Ministry of Electronics & Information Technology, India Building Secure, Scalable and AI-Ready Cloud Infrastructure (Dec. 12, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202897&lang=2®=3>; Ministry of Electronics & Information Technology, Guidelines for Procurement of Cloud Services 2-6 (Version 2.2, 2024), https://gsdl.org.in/rms2/doc/5.%20Guidelines_Procurement_Cloud%20Services_v2.2_1.pdf.

8. Chapter VII: Data Protection and Cyber Security

8.1 Overall Assessment

Chapter VII sits alongside two existing statutory regimes. The first is the Digital Personal Data Protection Act, 2023, which governs personal data. The second is the Information Technology Act, 2000 read with the CERT-In Directions of 2022, which govern cyber security.⁴⁴ Neither regime was designed specifically for the use of AI in the judiciary. Chapter VII therefore seeks to address risks that are either outside their scope or only partially addressed. This relationship is summarised below.

Table: The existing Law and what the Draft Regulations Add

The existing regime	What it covers, and its limit for judicial AI	The gap	The safeguard the Draft adds
DPDP Act, 2023, on personal data	Defines only personal data under section 2(t), treats all personal data alike, and exempts processing by a court in its judicial function under section 17(1)(b), keeping only the security duty in section 8(5).	Does not recognise a separate category of sensitive data. Much judicial processing is exempt.	A defined category of sensitive judicial data under Regulation 3(zh), anonymisation before training under Regulation 48(4), minimisation in system selection under Regulation 48(3), and least privilege access under Regulation 48(6).
IT Act, 2000, with the CERT-In Directions, 2022, on cyber security	Requires reporting of cyber incidents within six hours of awareness, retention of ICT logs for 180 days, and clock synchronisation. It binds the courts as government bodies.	It is not AI specific, and does not provide AI incident classification.	An AI Incident Database under Regulation 39, a fall back protocol under Regulation 42, and the security duty under Regulation 15.

⁴⁴ Digital Personal Data Protection Act, No. 22 of 2023 (India); Information Technology Act, No. 21 of 2000 (India), § 70B; Indian Computer Emergency Response Team, Directions Relating to Information Security Practices, Procedure, Prevention, Response and Reporting of Cyber Incidents for Safe and Trusted Internet ¶¶ 6-7 (Apr. 28, 2022), https://www.cert-in.org.in/PDF/CERT-In_Directions_70B_28.04.2022.pdf.

The relationship between these regimes and the Draft Regulations can be understood as follows:

First, the DPDP Act is not designed specifically for judicial AI. It regulates personal data but does not recognise a separate category of sensitive data or address AI-specific risks such as model training, inference or re-identification. Much judicial processing is also exempt under Section 17. The Draft Regulations address these gaps by introducing the concept of Sensitive Judicial Data and prescribing safeguards for its use in AI systems.

Second, Section 17(1)(b) of the DPDP Act exempts processing by courts when undertaken in the discharge of judicial, quasi-judicial, regulatory or supervisory functions.⁴⁵ The scope of this exemption is, however, limited. As discussed below, a substantial portion of the data processing contemplated under the Draft arises in administrative and technological functions, which may not fall within the exempted category.

Third, the cyber security regime under the IT Act and the CERT-In Directions is binding but general. It mandates incident reporting within six hours of awareness, retention of system logs, and baseline security practices. It does not distinguish AI incidents from other cyber events, nor does it presently map its internal incident reporting mechanisms onto this statutory framework.

Finally, the Draft Regulations introduce safeguards that are absent from both regimes, including restrictions on external transfers of judicial data, requirements relating to sovereign or on-premise deployment, in-house audits, and controls on the use of judicial data for AI training and system development.

Taken together, Chapter VII represents the first attempt to establish an AI-specific data governance framework for the Indian judiciary. The observations below are directed at ensuring that this framework remains internally coherent, operationally feasible and capable of implementation across courts with varying levels of technological capacity.

8.1.1 Applicability of the DPDP Act to Administrative and Technological Functions

The position of courts under the DPDP Act warrants closer examination in the context of administrative and technological functions under Regulation 10 of the Draft. Section 17(1)(b) of the Act exempts processing by courts when undertaken in the discharge of judicial, quasi-judicial, regulatory or supervisory functions. This exemption does not extend to administrative or technological activities.

The Draft contemplates a wide range of such functions, including digitisation of records, case management systems, e-filing infrastructure, conversational

⁴⁵ Digital Personal Data Protection Act, No. 22 of 2023, § 17(1)(b) (India).

interfaces and the deployment of AI systems in court processes, as reflected in Regulations 19 and 48.

In this domain, courts are likely to assume the position of data fiduciaries, with all corresponding statutory obligations under the DPDP Act. The use of AI systems, particularly those drawing on large volumes of historical court data, further underscores this position.

This intersection would require careful treatment. At a minimum, the categories of data processed across administrative systems, the nature and purpose of such processing, and the safeguards applicable to it would need to be identified and documented.

In particular, the deployment of AI systems in these contexts would warrant structured data protection impact assessments. These should consider the volume and sensitivity of data processed, the risks arising from automated use, and the adequacy of minimisation and anonymisation measures.

Recommendation: It is important to recognise that administrative and technological functions undertaken by courts will fall outside the exemption in Section 17(1)(b) of the DPDP Act and, in such cases, attract the corresponding obligations applicable to data fiduciaries. Before any AI system is deployed for these functions, courts should identify and document the categories of data processed, the purpose of processing, and the safeguards applicable to such processing.

For AI systems in these contexts, the Regulations should require a structured data protection impact assessment that addresses the volume and sensitivity of the data, the risks arising from automated use, and the adequacy of the minimisation and anonymisation measures adopted.

8.1.2. Data Minimisation, Anonymisation and the Missing Baseline

Draft Regulations 48(3) and 48(4) introduce requirements of data minimisation and anonymisation. Their application, however, presupposes a prior baseline, namely a clear mapping of what data is collected, processed and retained across court systems.

This baseline is particularly important in the judicial context, where court records are inherently identifying in nature. Party names, case numbers, addresses, advocate details and procedural histories are not incidental attributes but form part of the judicial record. They are essential to traceability, the development of precedent and, in many cases, the evidentiary value of the material.

Moreover, as discussed in Section 3.1.5 on the definition of Sensitive Judicial Data, different categories of judicial proceedings are subject to distinct statutory confidentiality and record-management obligations. These differences are directly relevant to the implementation of Regulations 48(3) and 48(4). Data minimisation and anonymisation cannot be applied uniformly across court records, but must instead be assessed in light of the legal requirements governing the particular category of judicial data.

Recommendation: Courts should be mandated to undertake a structured data-mapping exercise before implementing obligations relating to data minimisation and anonymisation. The exercise should identify the categories of personal data collected, retained and transmitted across different court processes. It should identify information that is essential to traceability, precedent and evidentiary value, and document the statutory confidentiality and record-management obligations applicable to different categories of judicial proceedings.

This baseline would guide the application of data minimisation and anonymisation requirements, ensuring that they are implemented in a manner that is both technically feasible and consistent with existing legal obligations.

8.1.3. Infrastructure, Control and Operational Feasibility

Three provisions appear to proceed on the assumption that courts possess, or will shortly acquire, infrastructure capable of hosting AI systems entirely within court-controlled environments. Regulation 38(2) requires audits to be conducted in-house and restricts external access to technical artefacts. Regulation 46(4)(j) mandates sovereign or on-premise deployment for systems processing sensitive judicial data. Regulation 48(1) restricts transfer of such data to external systems

This assumption of infrastructural readiness does not presently reflect the operational reality of most Indian courts, which rely on infrastructure provided by the National Informatics Centre or commercially hosted platforms.

As drafted, these provisions may inadvertently constrain the use of cloud-based or managed systems that are, in practice, the most feasible and scalable option in the near term. A more calibrated approach would preserve control over sensitive judicial data while permitting certified or audited environments that meet defined jurisdictional and security criteria.

Recommendation: The requirements of in-house audit, sovereign or on-premise deployment, and restricted external transfer should be redrafted in a manner that preserves control while accommodating the infrastructure most courts actually rely upon. It would be preferable to permit certified or independently audited environments that meet defined jurisdictional and security criteria, rather than to confine courts to wholly self-hosted arrangements, so that the framework remains workable across courts of differing technological capacity while keeping sensitive data within a controlled and accountable setting.

8.1.4. AI Processing and Incident Pathways

The Draft Regulations introduce several AI-specific safeguards that do not presently exist under either the DPDP Act or the CERT-In framework, including controls on the use of judicial data for training, testing and system development.

The Draft creates an internal reporting structure through the AI Secretariat and the AI Incident Database. The CERT-In Directions impose parallel obligations for reporting cyber incidents within six hours of awareness, along with requirements for log retention. The DPDP framework introduces a separate pathway for personal data breaches. At present, these pathways are not integrated.

Recommendation:

The Regulations would benefit from a clearly articulated incident framework that specifies how an event is classified and routed across the AI Secretariat, the CERT-In channel and any personal data breach pathway under the DPDP framework. A severity-based classification, reporting timelines aligned with existing statutory requirements, and harmonised logging obligations would allow AI-specific risks to be addressed without displacing existing duties, and would give courts a clear protocol to follow when an incident occurs.

9. Chapter VIII: Capacity Building, Training and Best Practices

9.1 Overall Assessment

Chapter VIII addresses one of the most operationally critical requirements for the Regulations to function as intended. A governance framework that relies on Human-in-the-Loop oversight, meaningful verification of AI outputs, and responsible incident reporting can only be as effective as the people it depends on to exercise those functions. The framework is appropriately structured and the minimum curriculum under Regulation 49(3) covers the right ground. The concerns that arise are not with the framework's design but with the conditions into which it will be implemented, several of which the Regulations do not acknowledge.

The global evidence on this point is stark. UNESCO's 2024 Global Judges' Initiative survey of over 500 judicial operators across 96 countries found that 44% of respondents were already using AI tools such as ChatGPT for work-related tasks, while only 9% reported that their institutions had issued guidelines or provided any AI-related training.⁴⁶ The survey documents a pattern that appears consistent across jurisdictions: AI use precedes institutional guidance, and institutional guidance precedes formal training. India is not an exception to this pattern. Official material describes SUPACE as being in experimental development and current AI-based solutions as limited to controlled pilot deployments.⁴⁷ Kerala has mandated AI transcription in district courts from November 2025.⁴⁸ The Regulations are now being framed. Structured training, at the scale and depth Regulation 49 requires, does not yet exist.

9.1.1 The Training Gap: What Currently Exists

The training architecture that Regulation 49 must build upon is limited. A review of the National Judicial Academy's 2025-26 academic calendar shows no curriculum item specifically addressing AI system literacy, hallucination detection, explainability assessment, or the HITL verification protocol that Regulations 8(3), 40 and 49(3)(b) contemplate.⁴⁹ Consequently, the training obligation envisaged under Regulation 49 cannot simply be integrated into an existing AI education framework; it would require the creation of an entirely new and specialised training programme.

⁴⁶ UNESCO, Global Judges' Initiative: Survey on the use of AI systems by judicial operators, <https://unesdoc.unesco.org/ark:/48223/pf0000389786/>.

⁴⁷ Press Information Bureau, Ministry of Law & Justice, Use of Artificial Intelligence in Legal Field (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202159&lang=1®=3>.

⁴⁸ High Court of Kerala, Office Memorandum on Use of 'Adalat.AI' Speech-to-Text Transcription Tool (Oct. 27, 2025), <https://cdnbbsr.s3waas.gov.in/s3ec01dc6a7e655d7e5840e66733e9ee67/uploads/2025/10/2025102791.pdf>.

⁴⁹ National Judicial Academy, Academic Calendar 2025-26, [https://nja.gov.in/Academic_Calendars/Academic_Calendar_2025-26\(28-09-2025\).pdf](https://nja.gov.in/Academic_Calendars/Academic_Calendar_2025-26(28-09-2025).pdf).

The existing digital capacity building initiatives within the judiciary, while significant, are not designed to meet this objective. Under the e-Courts Mission Mode Project (Phases I and II), more than 14,000 judicial officers received training on the Ubuntu Linux operating system, while approximately 3,900 court personnel were trained as District System Administrators.⁵⁰ These programmes established a baseline of digital literacy and enabled the transition towards electronic court administration. However, operational proficiency in digital systems is qualitatively different from the competencies required for the responsible use of artificial intelligence. AI literacy demands an understanding of probabilistic outputs, hallucination risks, explainability limitations, appropriate verification techniques, data governance obligations, and the circumstances in which human intervention must override automated recommendations. Regulation 49 therefore contemplates a substantially higher level of institutional capability than that created through existing digitalisation initiatives.

The implementation challenge is further compounded by existing structural constraints within the judicial system. The India Justice Report 2025 records judicial vacancy rates of approximately 21% in district courts and 33% in High Courts, while the average district court judge manages a docket of nearly 2,200 pending cases.⁵¹ It also notes that most states continue to allocate less than 1% of their budgets to the judiciary.⁵² In this context, the effective implementation of AI training obligations will depend not only on curriculum design but also on institutional capacity, time availability, and sustained financial resourcing. The issue is therefore not the desirability of training, but its operational feasibility at scale within existing constraints.

Recommendation: Regulation 49 should be operationalised through a phased implementation model, beginning with priority groups followed by wider rollout. Training should be role-specific, integrated into existing judicial education structures, and delivered through modular programmes combining foundational AI literacy with periodic refresher and advanced applied modules. Implementation should be explicitly linked to dedicated budgetary allocation and supported by institutional partnerships to ensure continuity, scalability and regular updating.

⁵⁰ Lok Sabha, Unstarred Question No. 2049, https://sansad.in/getFile/loksabhaquestions/annex/183/AU2049_gQDqY2.pdf?source=pqals.

⁵¹ India Justice Report, National Press Release (2025), https://indiajusticereport.org/files/English_National%20Press%20Release_India%20Justice%20Report%202025.docx.pdf.

⁵² India Justice Report, Budgets for Justice: A Pilot Study of Justice Budgets in 11 High-GDP States, https://indiajusticereport.org/files/IJR%20Budgets_Main%20Report_Single.pdf.

9.1.2 Training Curriculum and Accessibility

Regulation 49 establishes a sound framework for AI training within the judiciary. It appropriately requires training to be tailored to users' functions, prescribes a comprehensive minimum curriculum, and recognises the need for accessibility across district courts and India's linguistic diversity. The principal challenge is therefore one of implementation rather than design. While the same minimum topics may be relevant across the judiciary, the level of theoretical knowledge and practical competence required will differ between judges, supervisory officers and litigant-facing court staff.

The multilingual training requirement under Regulation 49(2) must be understood against the actual language environment of Indian courts. District and subordinate courts operate across a wide range of languages, Hindi across northern states, Kannada, Tamil, Telugu, Marathi, Bengali among others, at the regional level, while Fast Track Special Courts and Family Courts commonly involve code-switching between English and regional languages. The tools whose use the training must address are themselves unevenly calibrated across this landscape. The point is not theoretical: of the 31,184 Supreme Court judgments translated by SUVAS, about 70% are in Hindi, while many regional languages have few or no translated judgments.⁵³ AI tools drawing on or trained against judicial text of this distribution will perform reliably in Hindi but unevenly across regional languages. The training obligation under Regulation 49(2) should therefore address not only the language in which training materials are delivered, but the accuracy of AI tools themselves in the languages of the courts where they will be deployed.

Recommendation: A competency-based training framework prescribing role-specific learning outcomes and appropriate combinations of theoretical and practical instruction for different categories of judicial personnel is recommended. The Apex Body's implementation guidelines should also specify minimum standards for multilingual training materials, their periodic review and institutional responsibility for their development.

9.1.3 The Best Practices Repository

Regulation 50 appropriately requires the Appropriate Authority to maintain a living repository of best practices, case studies, lessons drawn from AI incidents and guidance notes. By preserving institutional knowledge despite changes in judicial personnel, the repository can become an important mechanism for continuous learning and improved AI governance. Its effectiveness will, however,

⁵³ Press Information Bureau, Using regional languages in courts (8 August 2024), <https://www.pib.gov.in/PressReleaselframePage.aspx?PRID=2042983®=48&lang=2>.

depend on the ability to systematically collect, structure and retrieve information across multiple courts and jurisdictions in a consistent and usable form.

Recommendation: The Appropriate Authority should prescribe standardised templates for recording AI incidents, case studies, best practices and guidance notes, together with uniform reporting and periodic review requirements. This will ensure that information is consistently structured, easier to consolidate across jurisdictions, and more effectively curated into a coherent national repository accessible to all courts.

10. Chapter IX: Grievance Redressal

10.1 Overall Assessment

Chapter IX sets out a dual structure for addressing AI-related harm in court processes. Regulation 52 provides a grievance mechanism for harms arising from prohibited AI use under Regulation 20, while Regulation 53 preserves recourse to existing legal remedies for other AI-related harms.

This creates an uneven remedial framework. Regulation 52 offers a defined court-based mechanism for a narrow category of cases, whereas harms arising from permissible AI use, such as transcription errors, inaccurate legal suggestions, or flawed procedural assistance, are left to general legal processes under Regulation 53. Given the expected breadth of AI deployment under Regulation 19, this distinction may leave a significant category of AI-related harms without a tailored or accessible redressal pathway within the court system.

In addition, even within Regulation 52, key procedural elements such as timelines, standards of review, and procedural safeguards are not specified, which may affect its practical operation.

10.1.1. Grievance Redressal Mechanism

Regulation 52 raises concerns regarding institutional independence, the trigger for filing, and procedural clarity.

On independence, grievances relating to prohibited AI use are to be adjudicated by the same court whose processes are implicated. While Indian courts do exercise inherent and review jurisdiction to correct their own procedural errors, Regulation 52 appears to extend beyond that limited framework. Review under Order 47 CPC and Article 137 is confined to errors apparent on the face of the record. By contrast, AI-related errors are unlikely to be evident on the record and would typically require access to system logs, audit trails, or technical analysis. In the absence of any supporting investigative or disclosure mechanism, Regulation 52 may not sufficiently bridge this gap.

On the trigger for filing, Regulation 52 requires grievances to be filed at the “earliest opportunity.” This creates uncertainty as to limitation and may operate harshly where delay is not attributable to the applicant. More fundamentally, prohibited AI use may not be visible to the affected party. Without disclosure of AI deployment or access to relevant system information, a litigant may be unable to know that a grievance exists, making the “earliest opportunity” standard difficult to operationalise in practice. The Regulations also do not require disclosure or notification of AI use to affected parties, which may limit the effectiveness of the remedy in practice.

Recommendation: While Regulation 52(2) enables High Courts to prescribe procedures and formats, consideration may be given to setting baseline requirements on disclosure of AI use in proceedings and access to relevant information necessary to assess and file a grievance. This would support meaningful application of the “earliest opportunity” standard and improve the practical effectiveness of the mechanism.

10.1.2. Other Remedies (Regulation 53)

Regulation 53 preserves the right of any person to seek remedies under any law in force for harm arising from the use of AI in court processes.

The provision is appropriate in principle, but its practical effectiveness may be limited. The remedies it preserves, including review under Order 47 CPC, writ jurisdiction under Articles 226 and 227, and statutory appeals, are generally designed to address legal, procedural, or jurisdictional errors apparent on the record. They are less suited to addressing harm arising from AI system outputs embedded within court processes, such as transcription errors, inaccurate legal references, or incorrect procedural guidance, which may not be visible on the face of the record.

Recommendation: Consider clarifying the information available to parties seeking remedies under Regulation 53, including access to relevant AI-related audit records or system logs where necessary. Consideration may also be given to whether additional structured procedures are required for addressing harms arising from permissible AI use, with appropriate safeguards and differentiation in standards of review.

10.1.3. Escalation and Procedural Timelines

Regulation 52 provides no escalation pathway beyond the initial court of filing and no timelines for acknowledgment or disposal. There is no provision for referral to the Apex Body where a grievance raises systemic concerns, despite the Apex Body's function under Regulation 23 of "guiding actions in respect of grievances against AI systems" suggesting that such a role was contemplated.

Recommendation: An escalation pathway linking Regulation 52 to the Apex Body's function under Regulation 23 should be established for grievances raising systemic concerns. Minimum procedural timelines for acknowledgment and disposal should be specified, either within the

Regulation or through standards issued by the Apex Body under Regulation 24.

11. Chapter X: Miscellaneous Provisions

Regulations 54 to 57 are standard miscellaneous provisions. While they are broadly structured in conventional terms, each contains ambiguities that may affect implementation when read alongside the substantive framework.

Regulation 54(1) provides that the Regulations operate in addition to the DPDP Act and IT Act, while Regulation 54(2) states that in case of inconsistency with any other law, the other law will prevail. Where the Regulations introduce higher or more specific safeguards, particularly in relation to Sensitive Judicial Data, it is unclear whether such provisions would operate as supplementary standards or be treated as inconsistent with existing legislation and therefore displaced. Regulation 54(4), which gives precedence over administrative instructions, does not address this inter-legislative interaction.

Regulation 55 empowers the Chief Justice or Institutional Head to issue administrative directions for implementation. However, the provision does not require publication, dissemination, or any structured mechanism for oversight or alignment with the Apex Body. Given that such directions will materially shape compliance in practice, the absence of transparency requirements is notable.

Regulation 56 provides a power to relax or modify provisions, subject to compliance with Chapter II principles and exclusion of Regulation 20 prohibitions. While these are important safeguards, the Regulation does not set out criteria guiding the exercise of this discretion, nor does it provide for any review or oversight role for the Apex Body beyond receipt of intimation.

Recommendation: Consider clarifying the interaction between Regulation 54(2) and enhanced safeguards under the Regulations, particularly for Sensitive Judicial Data; basic transparency requirements for directions issued under Regulation 55; criteria and oversight mechanisms governing the relaxation power under Regulation 56; and a defined periodic review cycle under Regulation 57.

12. Operational Readiness and Implementation Sequencing

The Draft Regulations establish a governance framework of significant ambition. Their effectiveness in practice will depend not only on the quality of the principles they articulate but on whether the courts, tribunals and institutions to which they apply are in a position to implement them.

Readiness, in this context, requires an assessment of whether existing digital systems can support AI deployment, whether underlying data is of sufficient quality and consistency to sustain meaningful impact assessment and audit, whether hardware and connectivity infrastructure can bear the operational demands of the tools being introduced, and whether institutional and human resource capacity exists to discharge the oversight obligations the framework creates. This requires a careful evaluation of technical, operational and organisational factors across a judicial system that is not uniform in its current state of preparedness.

Digitisation and Data Quality

The pace of progress within the e-Courts project provides a material backdrop against which the implementation timeline of these Regulations must be assessed. Phase III, sanctioned at ₹7,210 crore and running from 2023 to 2027, set a target of digitising approximately 3,108 crore pages of legacy court records. As of 31 December 2025, approximately 637 crore pages of court records had been digitised against a target of 3,108 crore, representing around 20.5% of the target two years into a four-year phase.⁵⁴ At the current pace the 2027 deadline is unlikely to be met. Progress is also unevenly distributed: Allahabad's district courts account for over 168 crore pages digitised while Nagaland district courts record zero.⁵⁵ The practical point is not to criticise the pace of digitisation, but to ensure that AI deployment is sequenced against the real condition of the data environment in which it will operate.

DAKSH's State of Tribunals Report in India, 2025 records that more than ₹ 24.7 lakh crore, or 7.48 percent of India's GDP, is locked in over 3.5 lakh pending cases before just 350 tribunal members. It also identifies chronic vacancies, overreliance on contractual staff, jurisdictional overlaps, procedural inconsistencies and weak transparency as recurring constraints, raising concerns about the documented lines of accountability required by Regulation 8(2) and

⁵⁴ Lok Sabha, Unstarred Question No. 764, https://images.assettype.com/barandbench/2026-02-05/8j5fjad9/e_courts_.pdf; Bar and Bench, Over 637 crore pages of court records digitised under e-Courts Project: Law Ministry (5 February 2026), <https://www.barandbench.com/news/law-policy/over-637-crore-pages-of-court-records-digitised-under-e-courts-project-law-ministry>.

⁵⁵ Press Information Bureau, Digitization of Courts (12 February 2026), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2227226®=3&lang=2>.

the institutional continuity needed to discharge supervising-officer obligations under Regulation 40.⁵⁶

This uneven digital foundation has direct implications for AI deployment. Many of the tools permitted under Regulation 19, including legal research and case management systems, depend on comprehensive and standardised digital records. The baseline readiness assessment should therefore evaluate the state of digitisation within each jurisdiction before authorising AI systems that rely on digital case records.

Digitisation alone, however, is insufficient. The National Institute of Public Finance and Policy has identified persistent data-quality issues in e-Courts data, including incorrectly populated case-type fields, inconsistent statute and section fields, missing values and the absence of systematic institutional data-quality controls.⁵⁷

Poor-quality data will undermine Technology Enabled Impact Assessments under Regulation 35 by reproducing existing inaccuracies rather than identifying bias. Establishing minimum data quality standards before TEIAs are undertaken would strengthen the reliability of the assessment process.

Hardware, Connectivity and Infrastructure

Existing court infrastructure was designed primarily for digital case management rather than resource-intensive AI applications. Official material records SUPACE as being in experimental development and describes current AI-based solutions as controlled pilot deployments.⁵⁸ The Supreme Court has similarly identified hardware limitations as the main barrier to the wider deployment of SUPACE, a constraint that remains unresolved.⁵⁹ Phase II has achieved 99.5 percent Wide Area Network coverage with bandwidth ranging from 10 Mbps to 100 Mbps.⁶⁰ This is a significant achievement, but some AI applications, including real-time transcription, translation and secure audio

⁵⁶ Daksh, The State of Tribunals Report in India 2025, <https://www.dakshindia.org/the-state-of-tribunals-report-in-india-2025/>.

⁵⁷ National Institute of Public Finance and Policy (NIPFP), Problems with the e-Courts data (2020), https://www.nipfp.org.in/media/medialibrary/2020/07/WP_314__2020.pdf.

⁵⁸ Dr. Shyama Prasad Mukherjee Research Foundation, Artificial Intelligence in the Indian Judiciary: SUPACE, SUVAS and the Limits of Assistive Automation (2026), <https://spmrf.org/artificial-intelligence-in-the-indian-judiciary-supace-suvas-and-the-limits-of-assistive-automation/>.

⁵⁹ Dr. Shyama Prasad Mukherjee Research Foundation, Artificial Intelligence in the Indian Judiciary: SUPACE, SUVAS and the Limits of Assistive Automation (2026), <https://spmrf.org/artificial-intelligence-in-the-indian-judiciary-supace-suvas-and-the-limits-of-assistive-automation/>.

⁶⁰ Press Information Bureau, Computerization and Wide Area Network Connectivity under eCourt project (23 August 2023), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=1951374®=48&lang=2>; Press Information Bureau, Digital Transformation of Judicial System (11 December 2025), <https://www.pib.gov.in/Pressreleaseshare.aspx?PRID=2202160®=3&lang=2>.

processing, may require higher compute, storage, audio infrastructure and reliable connectivity depending on the use case.

Institutional Capacity and Staffing

Technical readiness must be matched by institutional capacity. The India Justice Report 2025 records vacancy rates of 21% in district courts and 33% in High Courts, with judges managing substantial caseloads. Tribunals face similar challenges. The DAKSH *State of Tribunals 2025* report highlights inconsistent infrastructure, limited operational data and heavy reliance on contractual staff.⁶¹ The additional responsibilities created by the Regulations, including human oversight, supervision and training, will place further demands on already stretched institutions. Staffing capacity should therefore form part of the Apex Body's readiness assessment before implementation.

Court Specific Constraints

Beyond the systemic constraints addressed above, several categories of court present specific compliance challenges that arise from pre-existing statutory frameworks and institutional conditions the Draft Regulations do not address directly.

The deployment of AI tools for transcription under Regulation 19(1)(b), drafting and summarisation under Regulation 19(1)(d), and conversational assistance under Regulation 19(1)(f) in **Family Court** proceedings engages obligations that arise independently of the Regulations. In *Santhini v. Vijaya Venketesh*, the Supreme Court held that technology-based hearing arrangements cannot be imposed on a party in Family Court proceedings without consent.⁶² This principle extends by analogy to AI systems that process or transmit the content of such proceedings. The Regulations do not currently address this constraint, and the risk classification framework under Regulation 12 could usefully incorporate specific protocols for AI deployment in proceedings subject to statutory consent and confidentiality requirements of this kind.

POCSO Courts and Fast Track Special Courts present particular implementation challenges. These courts frequently operate with limited physical infrastructure and contractual staffing, making compliance with the oversight, audit and supervisory obligations under Regulations 35, 38 and 40 more difficult than in better-resourced courts. Implementation planning should

⁶¹ India Justice Report, National Press Release: India Justice Report 2025, https://indiajusticereport.org/files/English_National%20Press%20Release_India%20Justice%20Report%202025.docx.pdf; DAKSH, The State of Tribunals Report in India 2025, <https://www.dakshindia.org/the-state-of-tribunals-report-in-india-2025/>.

⁶² *Santhini v. Vijaya Venketesh*, (2018) 1 SCC 1; See also SCC Online Blog, No video conferencing for reconciliation in matrimonial disputes; Allowed only after settlement fails: Andhra Pradesh High Court reiterates (14 February 2025), <https://www.sconline.com/blog/post/2025/02/14/no-video-conferencing-reconciliation-matrimonial-disputes-andhra-pradesh-high-court-legal-news/>.

therefore account for these institutional constraints rather than assume uniform technological and administrative capacity across all courts. Where AI systems are introduced, deployment should be accompanied by adequate infrastructure, staffing and technical support.

Juvenile Justice Boards, if brought within the scope of the Regulations, raise a distinct implementation issue. While the Regulations require data minimisation and anonymisation under Regulations 10 and 48, they do not address the mandatory destruction of records required under Section 24(2) of the Juvenile Justice Act, 2015.⁶³ Any AI system deployed in JJB proceedings should therefore be capable of permanent, deadline-triggered deletion of records in accordance with the statutory retention framework. This technical requirement should form part of the approval and assessment process for AI systems intended for use by JJBs.

The Case for Phased Implementation

These implementation challenges do not weaken the case for the Draft Regulations. Rather, they demonstrate the importance of aligning implementation with institutional readiness. The experience of SUPACE and SUVAS illustrates that AI can be successfully integrated into the judiciary when deployment is gradual and supported by appropriate infrastructure and institutional capacity.⁶⁴

The Apex Body should therefore undertake a baseline readiness assessment before commencement, evaluating digitisation, data quality, infrastructure, staffing and training across different categories of courts. Institutions that satisfy defined readiness criteria could implement the full framework immediately, while others should follow a phased pathway linked to measurable improvements in capacity. Such an approach is consistent with the Regulations' principles of proportionality and responsible adoption and would improve the prospects of effective implementation across the judicial system.

⁶³ Juvenile Justice (Care and Protection of Children) Act, No. 2 of 2016, § 24(2) (India).

⁶⁴ Press Information Bureau, Ministry of Law & Justice, Use of Artificial Intelligence in Legal Field (Dec. 11, 2025), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2202159&lang=1®=3>; Press Information Bureau, Ministry of Law & Justice, Promotion of Hindi in Higher Courts (Feb. 2, 2024), <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2001863>.

ANNEXURE - I

Mapping of Key Definitions: Draft Regulations for Use of AI in Courts against Comparative Frameworks

The table below maps the definitions in the Indian draft against equivalent provisions in the EU AI Act, Brazil CNJ Resolution 615/2025, ISO/IEC 42001:2023, the OECD AI Incident framework, the GDPR, Singapore's Registrar's Circular No. 1 of 2024, and other relevant instruments. The Significance column identifies the analytical import of each definition, including gaps, departures, and implications for enforcement and interpretation.

No.	Term	Meaning & Context	Significance
1	AI Audit	Regulation 3(1)(c) defines 'AI Audit' by reference to Regulation 38. Regulation 38 provides further operative detail - all AI Systems and AI Tools used by Courts must undergo technical, legal and ethical audits at intervals not exceeding one year; such audits must be conducted in-house; and the resulting reports must be submitted to the Appropriate Authority and recorded in the AI Register. The EU AI Act relies on conformity assessment and post-market oversight rather than employing the same defined term (Arts. 43 and 72–73). ISO/IEC 42001:2023 contains management-system requirements relating to internal audit and corrective action. Brazil's CNJ Resolution No. 615/2025 provides for audit methodology, qualified technical groups and periods for corrective action (Arts. 39–41).	The Draft framework establishes a periodic audit requirement and certain procedural safeguards. It does not, however, prescribe minimum audit criteria, requirements concerning auditor independence and qualifications, sampling and testing methodologies, a risk-based scope, timeframes for corrective action, or specified consequences where adverse findings remain unresolved. The requirement that audits be conducted in-house may also limit the scope for independent scrutiny. A binding Audit Protocol could address these matters and distinguish routine internal review from independent audit.

2	AI Committee	Regulation 3(1)(d) describes the AI Committee as a 'Committee of Judges', whereas Regulation 33(2) provides for a broader composition comprising at least three judges, a Senior Member of the AI Secretariat and such other members of the Secretariat as the Committee may invite. The Draft also permits external stakeholders to submit suggestions, but does not provide for their membership or voting participation. In the EU framework, the AI Office is a function of the European Commission, while the European Artificial Intelligence Board is a separate body composed of representatives of the Member States (Arts. 64–66). Brazil's CNIAJ has a multi-stakeholder composition under Article 15. Members in categories I–VI have voice and vote, while representatives of the OAB, the Public Prosecution Service, the Public Defender's Office and civil society in categories VII–X have voice only (Art. 15 §3). Resolution No. 674/2026 separately expanded CNJ Councillor representation from two to three and included the Councillor responsible for personal-data protection.	Regulation 3(1)(d) and Regulation 33 may benefit from closer alignment so that the definition of the AI Committee fully reflects its operative composition. The judge-led structure is supported by technical inputs through the AI Secretariat, the Apex Body's Technical Committee and CoRE-AI. Institutional clarity could be further strengthened by expressly setting out the Committee's membership, voting arrangements, recusal requirements and the circumstances in which external technical expertise may be invited or required.
3	AI Incident	Regulation 3(1)(e) covers an event in which the use, operation, failure, erroneous output or malfunction of an AI System causes, or creates a substantial risk of, harm, infringement of rights, disruption of a Court process, or breach of data security or	The Draft adopts a precautionary approach by bringing realised and potential harm within a single definition. It also includes 'disruption of a Court process' as an independent category of harm, thereby recognising a judicially specific concern not expressly reflected in the corresponding EU

		confidentiality. The OECD distinguishes between an AI incident involving realised harm and an AI hazard involving a potentially harmful event. The EU AI Act defines a narrower category of ‘serious incident’ and requires providers of high-risk AI systems to report specified incidents to the relevant market-surveillance authorities (Arts. 3(49) and 73).	and OECD formulations. Greater clarity on the meaning of ‘substantial risk’ would support consistent reporting and remediation across Courts. The Regulations may also benefit from specifying the relevant factors for assessment, including likelihood, severity, reversibility and the rights affected, and from distinguishing near-misses from realised incidents through proportionate reporting and escalation requirements.
4	AI Register	Regulation 37 requires each Court to maintain a register of approved AI Systems, including their approved purposes and conditions, service providers, Technical and Ethical Impact Assessments, audits and AI Incidents (Reg. 37). Public dissemination is subject to applicable data-protection, confidentiality and cybersecurity requirements. Under the EU AI Act, obligations concerning the registration of specified high-risk AI systems are principally set out in Article 49, while Article 71 establishes the EU database; the extent of public access varies across categories (Arts. 49 and 71). Brazil’s CNJ Resolution No. 615/2025 separately provides for annual public reporting and for registration and cataloguing through the Sinapses platform (Art. 18; Arts. 23–25). Sensitive, confidential or intellectual-property information may be omitted from the publicly available impact summary (Art. 24 §2).	The Draft adopts a decentralised, Court-level register, in contrast to Brazil’s central Sinapses catalogue and the EU database. Regulation 37 provides for the recording of systems following approval, but does not expressly make completed registration a condition precedent to deployment. Consistency and transparency could be strengthened by requiring registration before use, prescribing a common national data schema, identifying the fields that must be publicly accessible, and providing for central aggregation through the Apex Body or CoRE-AI.

5	AI Secretariat	Regulation 34 establishes a dedicated AI Secretariat for each AI Committee, headed by an officer in the cadre of District Judge and comprising officers and experts in judicial administration, technology, data science and law (Reg. 34). In addition to its administrative and technical functions, the Secretariat may grant expedited approval for specified low-risk administrative AI Tools and must maintain a register of tools approved through that process (Regs. 34(4)–(5) and 46(7)–(8)). The EU AI Office operates within the European Commission and exercises implementation, coordination and enforcement functions, particularly in relation to general-purpose AI (EU AI Act, Arts. 64–69). Brazil's CNJ Resolution No. 615/2025 establishes the CNIAJ, but does not provide for an equivalent court-level secretariat structure.	The Indian Secretariat performs a wider role than a purely administrative body, including a limited approval function. The principal issue is therefore the allocation of authority among the Secretariat, the AI Committee and the Apex Body. Institutional clarity may be strengthened by specifying the review or appeal of expedited approvals, conflict-of-interest safeguards, minimum staffing and expertise requirements, arrangements for cross-Court coordination, and the circumstances in which Secretariat experts exercise advisory or decision-making functions.
6	AI Service Provider	Regulation 3(1)(h) defines an 'AI Service Provider' as any person or organisation that, by virtue of a contract, designs, develops, trains, fine-tunes or integrates an AI System for deployment in a Court process, and remains accountable for its performance, accuracy, fairness, transparency and security throughout its lifecycle (Reg. 3(1)(h)). The EU AI Act defines a 'provider' by reference to the development and placing on the market, or putting into service, of an AI system under the provider's own name or trademark, and imposes both pre-market and	The Draft adopts a contract-linked service-provider category, while other actors within the AI supply chain are not separately defined with comparable precision. Greater clarity may therefore be warranted regarding responsibility where a Court fine-tunes or materially modifies a system, or where subcontractors, upstream model providers or API suppliers contribute to its operation. The framework may also benefit from a provision analogous to Article 25 of the EU AI Act, under which an entity that substantially modifies an AI system assumes the corresponding provider obligations.

		post-market obligations (Art. 3(3); Arts. 16, 20 and 72–73). It also reallocates provider obligations where another actor substantially modifies an AI system (Art. 25). Brazil's CNJ Resolution No. 615/2025 separately defines an AI-system developer and regulates externally contracted solutions (Art. 4(IV); Arts. 34–35).	
7	AI System / AI Tool	Regulation 3(1)(i) defines an 'AI System' or 'AI Tool' as any software, platform, application, device or process that employs Artificial Intelligence in connection with a Court process (Reg. 3(1)(i)). The exclusion for general-purpose software or digital tools appears in the separate definition of 'Artificial Intelligence' and therefore applies to an AI System when the two definitions are read together (Reg. 3(1)(m)). The EU AI Act adopts a capability-based definition requiring a machine-based system that infers from inputs how to generate outputs capable of influencing physical or virtual environments (Art. 3(1)). Brazil's definition similarly refers to a machine-based system operating with varying degrees of autonomy and generating probable and coherent outputs, although it does not reproduce every element of the EU definition (Art. 4(I)).	The Draft framework combines a purpose- and context-based definition of an AI System with a separate definition of Artificial Intelligence. Since the two definitions overlap but are not identical, some uncertainty may arise regarding their application to rule-based automation and routine software containing embedded AI functionality. Greater clarity could be achieved by expressly setting out the relationship between the two definitions and introducing a materiality threshold, so that the regulatory obligations apply where AI functionality materially performs or shapes a Court process, while excluding conventional deterministic automation and merely incidental AI-enabled features.
8	Algorithmic Decision-Making (ADM)	The Draft defines 'Algorithmic Decision-Making' to include the use of algorithmic outputs, whether in whole or	The Indian definition appropriately extends to hybrid human–algorithm decision-making, recognising that AI may materially influence an

		in part, to inform, recommend or arrive at a decision affecting any person or process (Reg. 3(1)(j)). Article 22 of the GDPR establishes a specific right in relation to decisions based solely on automated processing that produce legal or similarly significant effects; hybrid decisions remain subject to the GDPR's other requirements concerning lawfulness, fairness, transparency and data-subject rights (GDPR, Art. 22). The Draft separately prohibits judicial or quasi-judicial outcomes based solely on Algorithmic Decision-Making or AI-generated information (Reg. 20(1)(b)).	outcome without replacing human decision-making altogether. Its breadth may, however, also encompass routine administrative applications affecting a 'process', including scheduling and case allocation. The Regulations may therefore benefit from distinguishing adjudicatory, rights-affecting and purely administrative uses, and from calibrating impact-assessment, explanation and review requirements according to the material effect of the algorithmic involvement.
9	Anonymisation	Regulation 3(1)(k) defines 'Anonymisation' as the irreversible alteration or removal of personal information from a dataset so that the individual cannot be identified, directly or indirectly. The GDPR excludes genuinely anonymous information from its scope, having regard to all means reasonably likely to be used for identification, while European data-protection guidance treats anonymisation as a process of irreversible de-identification (GDPR, Recital 26; Article 29 Working Party, Opinion 05/2014). The GDPR separately defines 'pseudonymisation', under which information remains personal data because it may still be attributed to an individual through the use of additional information (GDPR, Art. 4(5)).	The Draft expressly incorporates irreversibility into the definition of anonymisation, consistently with the European understanding of genuine anonymisation. It does not, however, replicate the GDPR's distinction between anonymisation and pseudonymisation or prescribe continuing safeguards for pseudonymised data. This creates gaps in relation to the regulatory status of pseudonymised data, re-identification testing, residual-risk assessment, linkage attacks and continuing controls. In the judicial context, where case records may be pseudonymised for research, analysis or publication, greater clarity on the treatment of data that fall short of true anonymisation would be valuable.

10	Artificial Intelligence	Regulation 3(1)(m) defines 'Artificial Intelligence' as a machine-based system that infers, learns and generates decisions, predictions or recommendations from data with varying degrees of autonomy, including algorithms, computational processes and software deployed for Court processes. General-purpose software or digital tools are excluded unless specifically embedded with, augmented by or functionally dependent upon AI. The EU AI Act defines an 'AI system', rather than Artificial Intelligence as a standalone concept, by reference to autonomy, possible adaptiveness, inference from inputs and the generation of outputs capable of influencing physical or virtual environments (Art. 3(1)). Recital 12 indicates that systems based solely on rules defined by natural persons should fall outside that definition.	The Draft defines Artificial Intelligence separately from an AI System, although the two definitions substantially overlap. The use of 'infers, learns and generates' may be read as requiring each of these characteristics, potentially excluding systems that infer and generate outputs but do not themselves continue to learn. The definition also omits the EU requirement that generated outputs be capable of influencing physical or virtual environments. The exclusion for general-purpose software is not directly equivalent to the approach in EU Recital 12 and may become difficult to apply as AI functionality is increasingly incorporated into ordinary digital tools. Greater clarity on the relationship between Regulations 3(1)(i) and 3(1)(m) would support more consistent interpretation.
11	Black Box	Regulation 3(1)(n) defines a 'Black Box' as an AI System, typically employing deep learning, whose internal processes and decision-making logic are not transparent, comprehensible or capable of explanation by reference to identifiable rules or factors accessible to the user or affected person. The EU AI Act does not expressly define the term; instead, it addresses opacity through transparency, documentation and human-oversight requirements, particularly for high-risk AI systems (Recital 72; Arts. 13–14).	The Draft gives express regulatory recognition to opacity, while the EU AI Act addresses it through functional transparency and oversight obligations. Black Box character attracts heightened scrutiny, cannot be invoked to avoid accountability, and is restricted where an undisclosed, opaque or unexplainable system may materially affect personal liberty or lawful rights (Regs. 7, 8 and 20(1)(e)). The definition may, however, be more durable if framed by reference to the inability to provide an adequate explanation of an output, rather than by reference to a particular technical architecture.

12	Controlled Environment Testing	Regulation 3(1)(na) defines 'Controlled Environment Testing' as testing conducted in a secure and isolated environment without affecting the integrity, security or functioning of a Court's primary operational systems. Before full-scale deployment, the Appropriate Authority may require time-limited testing under the supervision of the AI Secretariat and against documented evaluation parameters; outputs generated during testing may not be used in actual adjudicatory or administrative decisions (Reg. 36). The EU AI Act establishes regulatory sandboxes as supervised frameworks for the development, training, testing and validation of AI systems before their placement on the market or putting into service (Art. 57).	The Draft mechanism is tailored to deployment readiness and the protection of Court operations, while the EU sandbox also serves broader regulatory-learning and innovation objectives. The Draft incorporates important safeguards, including the protection of live operational systems and the prohibition on using test outputs in actual Court decisions. However, testing remains discretionary, and the Regulations do not prescribe mandatory triggers, common testing standards or arrangements for recognising and sharing results across Courts. A common testing protocol could promote greater consistency while preserving institutional flexibility.
13	Data Minimisation	Regulation 3(1)(t) defines 'Data Minimisation' as the principle that an AI System shall collect, process and retain only such personal data as are strictly necessary for its specified purpose, and expressly applies this principle at the stages of system design, procurement, deployment and audit. The GDPR requires personal data to be adequate, relevant and limited to what is necessary, and incorporates data minimisation within the requirements of data protection by design and by default (GDPR, Arts. 5(1)(c) and 25(2)). Argentina's Law No. 25.326 similarly requires personal data not to be excessive in relation to their purpose and provides for their	The Draft provides useful operational specificity by expressly identifying the stages of system design, procurement, deployment and audit at which data minimisation must be considered. The GDPR also treats data minimisation as a continuing obligation rather than as a single assessment at the point of collection. The Indian formulation is nevertheless valuable in the judicial context because it directly incorporates the principle into AI procurement and audit processes. Greater clarity on necessity assessments, retention periods, deletion verification, system logs and backups, and the interaction with statutory judicial-record retention requirements would support consistent implementation.

		destruction when they are no longer necessary or relevant (Art. 4).	
14	Ethical Impact Assessment	Regulation 3(1)(v) defines an 'Ethical Impact Assessment' as a structured pre- and post-deployment evaluation of the design and deployment of an AI System against the principles set out in Chapter II. Regulation 35 separately requires a Technical and Ethical Impact Assessment before approval and specifies the minimum matters to be examined (Reg. 35). The EU AI Act requires specified deployers to conduct a Fundamental Rights Impact Assessment before first deployment and to update it where relevant elements have changed or are no longer current (Art. 27).	The Draft's assessment is benchmarked against the principles set out in Chapter II, rather than expressly against the full range of fundamental rights considered under the EU framework, and may therefore be narrower in scope. Although the definition contemplates both pre- and post-deployment evaluation, Regulation 35 principally establishes a pre-approval process and does not prescribe the trigger, timing, responsible body or procedure for subsequent reassessment. The Draft would benefit from requiring reassessment following material modification, retraining, a serious incident, a change in intended use, or at prescribed intervals, and from connecting such reassessment to audit findings and continued-deployment decisions.
15	Explainability	Regulation 3(1)(w) defines 'Explainability' as the capacity of an AI System, upon request, to provide a comprehensible account of the reasoning, factors, weightings and process by which it arrived at a particular output or decision, in terms understandable by the judicial officer, Court staff member or litigant concerned, without requiring specialist technical knowledge. The EU AI Act addresses related concerns through transparency and information requirements for deployers, effective human oversight, and a qualified right to a clear and meaningful explanation for certain persons affected by decisions	The EU AI Act addresses explainability through transparency and information requirements for deployers (Art. 13), human oversight by appropriately competent persons (Art. 14), and a qualified right to a clear and meaningful explanation in specified cases (Art. 86). The Draft's definition is more directly Court-user facing: upon request, it requires a comprehensible account that can be understood by judicial officers, Court staff and litigants without specialist technical knowledge. This imposes a more operationally demanding standard. Its application should, however, be calibrated to what is technically feasible and sufficient to support meaningful review and challenge, rather than

		based on the outputs of specified high-risk AI systems (Arts. 13, 14 and 86).	requiring reconstruction of a model's internal reasoning.
16	Generative Artificial Intelligence	Regulation 3(1)(y) defines 'Generative Artificial Intelligence' as a class of AI Systems that generates new content, including text, images, audio, video and computer code, by learning statistical patterns from large volumes of existing data. It further provides that AI-generated content may not be filed or produced before a Court without disclosure of its origin and verification through a designated centralised mechanism. Regulations 43 and 44 separately provide for declaration and verification mechanisms, including a declaration or certificate for AI-assisted submissions and an AI Content Verification Authority (Regs. 43(3) and 44). Singapore's Registrar's Circular No. 1 of 2024 defines GenAI, requires Court users to verify AI-generated material, provides for disclosure where the Court so directs, and prohibits the use of GenAI to fabricate or manipulate evidence.	Singapore places primary responsibility for verification on the Court user and does not generally require advance disclosure unless directed by the Court. The Draft adopts a dual mechanism comprising user-level disclosure under Regulation 43(3) and institutional verification through a centralised mechanism, an approach not reflected in Singapore's 2024 Guide. The relationship between the verification requirement in Regulation 3(1)(y), the declaration requirement in Regulation 43(3) and the AI Content Verification Authority under Regulation 44 would benefit from clarification. The substantive disclosure and verification obligations may also be more appropriately located in the operative provisions rather than within the definition itself.
17	Hallucination	Regulation 3(1)(z) defines 'Hallucination' as the generation by an AI System of outputs that appear plausible or coherent but are factually incorrect, fabricated, misleading or unsupported by verifiable data or source material, including the fabrication or misstatement of legal facts, evidence, precedents, statutory provisions, rules or legal	Neither Singapore nor South Korea treats hallucination as a formally defined legal category; both address it principally through verification requirements and institutional guidance. The Indian Draft gives the concept express legal recognition tailored to the judicial context. This creates a regulatory reference point for accountability, while the operative consequences arise through provisions that prevent

		principles. Singapore's 2024 Guide expressly uses the term 'hallucinating' in relation to invented cases, statutory provisions and facts, and requires users to verify AI-generated material independently. South Korea's National Court Administration published the <i>AI Guidebook for Judges</i> in February 2026, addressing risks including hallucination, bias, security and privacy.	hallucination from being invoked to avoid accountability and permit the Court to act where AI-generated material is fabricated, false, misleading or inaccurate (Regs. 8 and 43(6)).
18	Human-in-the-Loop	Regulation 3(1)(zb) defines 'Human-in-the-Loop' as a governance process in which the outputs, recommendations or results generated by an AI System are subject to mandatory human review, supervision and verification, with final decision-making authority, responsibility and accountability resting at all times with a human. The EU AI Act requires high-risk AI systems to enable effective human oversight by competent persons who can understand the system's limitations, remain alert to automation bias, interpret its outputs, disregard or override them, and interrupt the operation of the system (Art. 14). Argentina's San Juan Protocol requires human control before and after the use of generative AI and strictly prohibits the delegation of decision-making to such systems (Acuerdo General No. 102/2024, 'Valoración de Resultados').	The EU framework is principally design and operation oriented, requiring systems to enable effective human oversight. The San Juan Protocol similarly requires human control and prohibits the delegation of decision-making. The Indian formulation is more expressly outcome-oriented, requiring final decision-making authority, responsibility and accountability to remain with a human. This constitutes a strong articulation of human primacy, but does not by itself impose personal legal liability. The Draft may also benefit from greater consistency in its use of 'Human-in-the-Loop', 'human oversight' and related concepts, particularly where different levels of human involvement are intended.

19	Risk Scoring	Regulation 3(1)(zg) defines 'Risk Scoring' as the use of an AI System to assign a numerical or categorical score to an individual estimating the probability of specified future behaviour, including the commission of an offence, recidivism or failure to appear before a Court. Regulation 20(1)(d) prohibits Risk Scoring for any purpose in Court processes, including the assessment of flight risk, recidivism, bail or credibility. The EU AI Act prohibits specified criminal-offence risk assessments based solely on profiling or the assessment of personality traits and characteristics (Art. 5(1)(d)), while treating other specified law-enforcement systems used to assess the risk of offending or reoffending as high-risk (Annex III, point 6(d)). Systems assisting judicial authorities are addressed separately in Annex III, point 8(a).	The Draft is notable in providing a standalone definition of Risk Scoring, a term that is not separately defined in the EU AI Act. This creates a clear interpretive trigger for Regulation 20(1)(d), which prohibits Risk Scoring for any purpose in Court processes. The Indian approach is therefore broader and stricter than the EU framework: Article 5(1)(d) prohibits specified criminal-offence risk assessments based solely on profiling or the assessment of personality traits and characteristics, while certain other law-enforcement risk-assessment systems are classified as high-risk under Annex III, point 6(d).
----	--------------	--	--

Sources: Argentina, Personal Data Protection Law No. 25.326 (2000); Article 29 Data Protection Working Party, *Opinion 05/2014 on Anonymisation Techniques* (2014); Regulation (EU) 2016/679 (General Data Protection Regulation) (2016); ISO/IEC 42001:2023; Regulation (EU) 2024/1689 (EU Artificial Intelligence Act) (2024); OECD, *Defining AI Incidents and Related Terms* (2024); Supreme Court of Singapore, Registrar's Circular No. 1 of 2024, *Guide on the Use of Generative Artificial Intelligence Tools by Court Users* (2024); Court of Justice of San Juan, Acuerdo General No. 102/2024, *Protocol for the Acceptable Use of Generative Artificial Intelligence* (2024); Brazil CNJ Resolution No. 615/2025, as amended by Resolution No. 674/2026; Draft Regulations for Use of Artificial Intelligence in Courts, 2026 (India); South Korea National Court Administration, *AI Guidebook for Judges* (2026).

॥ न्यायस्तत्र प्रमाणं स्यात् ॥



nludelhi.ac.in | info@nludelhi.ac.in | igap.in | relations@igap.in