KANAN LABS

HANDBOOK ON INDIA'S AI REGULATIONS

Quick Glance at Global Regulation

Region	Regulatory Style	Primary Trigger	Timeline	What It Means
EU	Hard Restrictions & Obligations	Your system falls into a listed risk tier	Before market entry (certification required)	Classify your system, complete documentation, and obtain approvals
US	Voluntary standards	Harm to consumer, or sectoral violations	Post-hoc litigation + Market Driven	Launch fast, but legal exposure increases if harm or unfair practice is proven
India	**Techno-Legal Sector-wise regulation**	**SDF Designation + Volume/Risk Threshold**	**Ex-ante technical + Audit driven**	**Governance must be built into architecture (auditability, traceability, explainability)**

SREENATH GOVINDARAJAN

India's "techno-legal" framework is embedding legal obligations directly into digital infrastructure and system architecture. For founders and product teams, this means compliance is no longer a post-launch checklist. It needs to be a design decision. Regulatory anticipation must begin at the build stage if you want to avoid structural redesigns, capital erosion, and delayed scaling.

The Three Pillars

India is not regulating AI through one single system of laws. Rather, it is constructing a distributed control architecture that rests on three structural pillars. For product developers, these pressure points determine where your product will face scrutiny, escalation, or structural redesign.

Sectoral Enforcement

India distributes authority across sectoral regulators. The RBI governs AI in finance. Health regulators influence clinical AI deployment. Telecom authorities shape network-level AI. In essence, compliance is not uniform across industries, regulation has become more contextual. Your industry determines exposure.

Infrastructure Embedded Governance

India's model embeds governance within Digital Public Infrastructure. Consent must be programmable. Data flows must be traceable and outputs need to be explainable. For companies, the essentially burden shifts from legal declarations to technical defensibility. If governance is not built into your system's architecture, it cannot be retrofitted without cost. In this model, compliance is not about filing forms. It is about how your system is built.

Implicit Risk Orientation

Regulatory intensity increases through designation mechanisms, most notably the Significant Data Fiduciary (SDF) classification under the DPDP framework. Risk is inferred from data volume, sensitivity, electoral impact, and systemic exposure. Once designated, obligations escalate, creating a dynamic compliance threshold. The same product can move from light-touch regulation to high-intensity oversight simply by scaling. Growth itself becomes a regulatory variable.

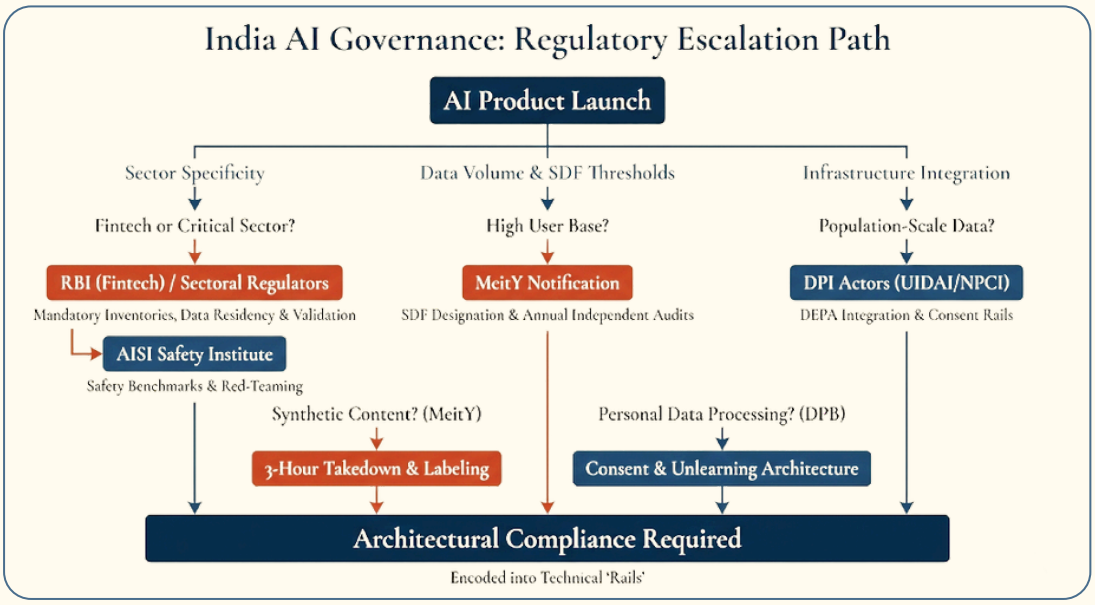

Where Regulation Enters Your Stack

AI Governance is slowly shifting towards being systems-based. In 2026, we're noticing regulatory obligations increasingly manifest as technical requirements: consent logs, audit trails, explainability layers, and takedown protocols. The practical implication is simple, compliance is moving from a legal team's burden to an engineer's. If your system cannot demonstrate traceability, consent provenance, or explain model behaviour, you are exposed.

Lifecycle Friction Mapping

Data Ingestion: Build Consent

Under the DPDP framework, personal data cannot be casually repurposed for model training. Consent must be explicit, purpose-specific, and further language-accessible. Bundled terms-of-service agreements are unlikely to satisfy scrutiny if data is used for secondary AI training.

This means that:

- You need a consent ledger.
- You need traceability of data origin.
- You need revocation pathways.
- You need clarity on training vs service-use boundaries.

Deployment & Interface: Transparency & Speed

The 2026 IT Amendments introduce operational duties at the interface level. Platforms must visibly label synthetically generated content and respond rapidly to certain forms of abuse, including non-consensual deepfakes.

For product teams, this means:

- UI-level disclosure mechanisms
- Automated content detection triggers
- Escalation dashboards
- Defined takedown workflows

Model Training:

India's evolving data regime introduces tension between the right to erasure and the persistent nature of model weights. If a user withdraws consent, the system must demonstrate reasonableness in honouring erasure requests. This is crucial since ignoring this at design stage may lead to downstream liability.

Practical mitigation strategies may include:

- Tokenisation at ingestion
- Dataset segmentation
- Machine unlearning research pipelines
- Training data isolation layers

> Long-term scalability in India requires treating regulatory anticipation as a system design phase. Success belongs to those who ensure that "Provable Compliance" is not a post-facto report, but a technical artifact of the digital foundation

BUILD – CONSIDER – DO CHEATSHEET

IF BUILDING: consumer-facing GenAI tool

CONSIDER: SGI labelling mandates

IMMEDIATE MOVE: Integrate metadata embedding and cryptographic watermarking at the inference layer.

IF BUILDING: automated recruitment/financial tool

CONSIDER: Scrutiny regarding "Black-box" opacity

IMMEDIATE MOVE: Standardise validation using Nishpaksh toolkit for fairness auditing.

IF TRAINING: models on domestic data

CONSIDER: Separate consent mandate & TDM

IMMEDIATE MOVE: Implement "Model Improvement" opt-ins in regional languages & map data flows.

IF DEPLOYING: in regulated sectorsl

CONSIDER: 48-hour incident reporting windows

IMMEDIATE MOVE: Collate semi-annual data inventory using the National AI Incident Database schema.

IF MANAGING: third-part AI vendors

CONSIDER: DPDPA timelines & retention SLAs

IMMEDIATE MOVE: Amend contracts to require 1-year long log retention & auditability access.

IF DEVELOPING: medical diagnostic software

CONSIDER: ICMR Clinical validation for imaging

IMMEDIATE MOVE: Establish a "Human-in-the-loop" protocol to manage deficiency in service liability.

IF OPERATING: a social media or content platform

CONSIDER: 3-hour takedown window for deepfakes

IMMEDIATE MOVE: Deploy the Saakshya project toolkits for automated deepfake detection.

IF AUDITING: internal governance and accountability

CONSIDER: 3rd party algorithmic audits for SDFs

IMMEDIATE MOVE: Utilise Track-LLM for governance testing to build a "provable" record of compliance.

IF IMPLEMENTING: legal or compliance automation

CONSIDER: Judicial penalty for hallucinations

IMMEDIATE MOVE: Implement verification layers to double-check citations with primary sources.

IF DESIGNING: modified content systems

CONSIDER: Graded liability model under the DIA

IMMEDIATE MOVE: Tag all AI-modified data packets at the point of origin to ensure persistent traceability.

IF TARGETING: public sector or critical infrastructures

CONSIDER: TEC 57050:2023 ratings

IMMEDIATE MOVE: Obtain technical certification from AISI for fairness and robustness.

IF PROCESSING: global data via Indian infra

CONSIDER: extra-territorial reach of EU AI Act

IMMEDIATE MOVE: Adopt interoperable "techno-legal" standards to maintain EU supply chain access.

Regulatory Foresight (2026-2027)

Probability	Development	Practical Shift	Business Impact
Highly Likely	Phased enforcement of DPDP Rules by mid-2027. AISI technical benchmarks for "Safe & Trusted AI."	Data governance and safety standards become enforceable and measurable.	Consent traceability, audit logs, and model evaluation processes must be engineered — not documented post-facto.
Emerging	Digital India Act (DIA) replacing IT Act; possible graded liability and narrowing of safe harbour. Copyright/TDM framework may formalise training rules.	Platform accountability increases. AI-modified content faces tighter scrutiny. Training economics may shift.	Content moderation, provenance tracking, and dataset strategy require forward planning.
Structural Watchpoints	Possible centralised AI Ethics/Accountability body.	Oversight becomes more uniform and potentially stricter.	Cross-sector AI deployments face standardised, less negotiable compliance.

Early Warnings to Watch Out For

Signal	Why It Matters	Who Feels the Impact?
First SDF Notification List	Defines scale thresholds for heightened compliance.	High-scale platforms, AI-enabled fintechs, health-tech firms, and data-intensive startups.
AIGG Uniform Standards	Indicates cross-sector technical alignment.	Companies building cross-sector AI tools or operating in regulated industries.
DPBI Operationalisation	Activates formal breach enforcement timelines.	Any company processing significant volumes of personal data, especially those without incident-response systems.

Final Takeaways

- Compliance is no longer a job for the General Counsel alone. It requires Product Managers who understand metadata tagging and audit logs.
- If you cannot explain the output, you cannot sell it to a regulated Indian sector. Explainability is not a feature; it is the license to operate in Finance, Health, and Government.

sreenathfromkanan@gmail.com | kananlabs.in