



Event Report

Operationalising India's New AI Labelling Rules: Readiness and Challenges

April 2026



Event Report

Operationalising India's New AI Labelling Rules: Readiness and Challenges

April 2026



Koan Advisory Group is a New Delhi-based public policy consultancy. It specializes in policy and regulatory analysis in both traditional and emergent sectors and markets. For more information, please visit: www.koanadvisory.com

Authors
Prakhar Rai and Srishti Joshi

©2026 Koan Advisory Group

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means without permission in writing from Koan Advisory Group

contactus@koanadvisory.com | www.koanadvisory.com

A

Context Overview

India recently implemented the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, marking the operationalisation of the country's first dedicated regulatory framework to address harms arising from artificial intelligence. The amendments impose obligations that mandate the labelling of synthetically generated information (SGI)¹ and significantly shortened the timelines for responding to content takedown requests.

The amendments came into force on February 20, 2026. Subsequently, the Ministry of Electronics and Information Technology (MeitY) issued an advisory on March 16, 2026, emphasising the need for intermediaries to ensure compliance with due diligence obligations.² The advisory called on digital intermediaries to enforce user terms of service and AI usage restrictions, including deterrent measures and to remove or disable access to harmful SGI within the timelines prescribed under the Rules.

B

Roundtable Overview

Koan Advisory Group, in collaboration with the Federation of Indian Micro Small and Medium Enterprises (FISME), convened a multi-stakeholder roundtable to unpack the operational readiness and implementation challenges associated with the new rules. The discussion brought together 30 stakeholders, with representatives from domestic and international digital service providers, including social media services, advertising agencies, home-grown AI startups and technology law and policy experts. It served as a platform to surface practical concerns emerging from the implementation of the rules.

The discussion focused on identifying areas of ambiguity, operational constraints and stakeholder-specific challenges in complying with the new regulatory requirements. It was structured across two thematic segments: (1) the application of SGI labelling obligations across the advertising value chain, with a focus on responsibility at each stage of content creation and dissemination; and (2) the expanding scope of intermediary obligations, particularly in light of shortened takedown timelines and evolving interpretations of intermediary status.

Event Snapshots





Overview of the New Rules

The new rules require digital intermediaries to prominently label all SGI in a clear and noticeable manner. Intermediaries must now notify users every three months (moving away from the earlier yearly obligation) about penalties for sharing unlawful content and deploy reasonable technical measures to prevent the creation and spread of harmful SGI. Intermediaries must act promptly to remove unlawful content upon awareness through monitoring, user complaints or government and court orders.

For Significant Social Media Intermediaries (social media intermediaries with over five million registered users), additional obligations include collecting and verifying user declarations before SGI is shared, clearly labelling such content, and proactively filtering serious harms such as rape and child sexual abuse material, as well as previously removed identical content. Failure to comply may result in loss of safe harbour protections.

The new rules also shorten compliance timelines. Intermediaries must act on court or government takedown orders within 3 hours (earlier 36 hours), remove certain explicit content within 2 hours (earlier 24 hours), resolve user grievances within 7 days (earlier 15 days) and resolve user requests for removal of illegal content within 36 hours (earlier 72 hours).

D

Key Recommendations

- 1 Strengthen consumer communication:** Invest in consumer education to build awareness on AI-generated content, including what it is, how it is legitimately used and what SGI labels signify. To implement this, digital service providers should develop clear disclosures, user-facing explanations and awareness campaigns to reduce misinterpretation and preserve trust in beneficial SGI content, to ultimately mitigate negative perceptions associated with labelling.

- 2 Build internal capacity and awareness:** Sensitise internal legal/policy and product teams to SGI-related obligations under the new rules and accordingly invest in training to ensure consistent and informed implementation across functions.

- 3 Simplify labelling formats:** Terms like “synthetically generated” may be difficult to understand for users because of its dense technical nature. Use intuitive visual indicators, such as a universal AI icon, to ensure labels are accessible across India’s multilingual user base.

- 4 Allow phased implementation:** Provide a defined 12-month transition period for compliance, drawing on international approaches such as South Korea’s AI Basic Act, to allow industry sufficient time to integrate provenance-based standards (such as C2PA metadata) into widely used creative tools. This would help ensure that, when enforcement begins, compliance is anchored in verifiable technical signals rather than platform-side attestation or guesswork.

- 5 Recognise advertiser attestation as due diligence for intermediaries:** Clarify that intermediaries which implement advertiser attestation frameworks to obtain disclosures on AI-generated content may be deemed to have fulfilled their due diligence obligations under the Rules, and accordingly retain safe harbour protections.

- 6 Distribute obligations across the value chain:** Avoid concentrating compliance burdens. Responsibilities should be proportionately shared between content creators, intermediaries, advertisers and publishers with obligations calibrated to both their role in the value chain and what is technically feasible at each stage.

- 7 **Clarify liability through contracts:** Based on technical feasibility, clearly allocate responsibility for SGI labelling and compliance across brands and agencies. This can be defined through detailed contractual provisions, including internal checklists and audit mechanisms to ensure adherence to labelling requirements.

- 8 **Adopt a risk-based approach:** Differentiate between categories of content and types of actors, ensuring that obligations are calibrated based on risk, scale and capacity rather than applied uniformly across the ecosystem.

- 9 **Adopt a hybrid compliance approach:** Given technological limitations and where economically possible, combine proactive measures such as SGI labelling with reactive tools like detection systems and takedown workflows to manage risks effectively.

- 10 **Economic impact:** The compliance costs imposed by these rules risk eroding AI's core value proposition as a low-cost scaling tool. This could constrain digital market access and limit growth in a sector that relies on AI-driven efficiency and scalability. Industry associations should document and share economic consequences with policymakers to ensure regulatory calibration is grounded in evidence of real-world impact.

E

Industry concerns

1 Trust and Consumer Engagement

Stakeholders noted that mandatory SGI labelling may adversely impact consumer perception and trust. In an environment where AI-generated videos, audio, and text increasingly blur the line between human and machine generated content, users often rely on labels as a primary signal of authenticity. The presence of an SGI label may lead users to assume that an entire piece of content is artificially generated, even where only limited elements are synthetic.

Labelling all SGI creates an unintended risk of misinterpretation, particularly for brands advertising real products using partially AI-generated content. Using AI for “digital mannequins” or benign AI-generated environments, such as simulated event settings, is essentially non-deceptive, yet would still trigger labelling obligations. In such cases, the application of an SGI label may lead consumers to question authenticity of a given product or service, thereby eroding trust and reducing willingness to engage with or purchase the advertised goods.

Illustratively, stakeholders pointed to scenarios where a brand uses AI to generate a professional runner (digital mannequin synthetic model) in a realistic setting to showcase a real running shoe, noting that consumers in a purchase-oriented context may conflate the label with the authenticity of the product rather than the surrounding elements. The application of an SGI label in such cases may lead consumers to question the authenticity of the product itself, thereby eroding trust and reducing willingness to engage with or purchase the advertised goods.



Here is an example of AI-generated models used in apparel marketing.

Image Credits:
[Apparel Resources](#)

Concerns were raised around the design and placement of labelling requirements. Although the rules mandate that SGI labels be “prominent and clear,” a one-size-fits-all approach is incompatible with mobile-first advertising formats and diverse user interfaces. Stakeholders also noted that persistent front-of-ad disclosures could lead to “banner blindness” and “label fatigue,” where users begin to ignore such indicators altogether due to their ubiquity.

Additionally, forcing intrusive overlays on formats such as short-form videos or display advertisements may obscure key visual elements, reduce the utility of the content and degrade the user experience for both consumers and advertisers. A potential approach to address these concerns could be through adopting the post-click transparency model, where detailed information on AI-generated elements is made available through an accessible interface (for instance, via a single click). This allows for more meaningful user engagement with disclosures, while avoiding the negative effects associated with persistent on-screen labelling.

A few advertising stakeholders noted that AI-generated content currently tends to see lower engagement and sharing, as it is often perceived as low-effort content. While others observed that this perception is beginning to shift as AI outputs become harder to distinguish from human-created ones, they were concerned that mandatory labelling could reinforce negative biases and limit the effectiveness of such content for consumer engagement.

Stakeholders also expressed concern around an implied truth paradox, where labelling only a subset of content may inadvertently increase the perceived credibility of unlabelled material, potentially making users more susceptible to misleading or deceptive content. These concerns are also reflected in emerging research, with a study finding that prominent AI labels can lead to a 31.5% drop in click-through rates (CTR), pointing to a potential modernisation penalty. At the same time, research also indicates that fully AI-generated ads can outperform human-created ones in certain contexts, achieving a 19% higher CTR.³

2 Liability Conundrum for Brands and Advertisers

Stakeholders discussed the lack of clarity regarding the attribution of responsibility in scenarios where brands outsource creative production to advertising agencies that may use AI tools in the content creation process. In such cases, ambiguity arises as to who bears the obligation to ensure compliance with labelling requirements and the persistence of such labels as content is disseminated across platforms.

While an ad agency may be directly responsible for generating the content, it is often done on the instructions of, and for the commercial benefit of, the brand, complicating the allocation of liability. To address such concerns, stakeholders expressed that contractual arrangements are beginning to evolve to incorporate AI-specific clauses. These provisions seek to define roles, responsibilities, and compliance obligations between brands and agencies, including adherence to labelling requirements under the Rules.

Stakeholders also noted the parallel development of internal compliance checklists and review processes within organisations to ensure that AI usage is identified, assessed, and appropriately labelled before dissemination. Such measures were viewed as necessary to mitigate liability risks and avoid disputes arising from unclear attribution of responsibility.

3 Liability Conundrum for Intermediaries

From the perspective of intermediaries, an “advertiser attestation” framework may offer a practical mechanism to discharge due diligence obligations under the Rules. Under such an approach, platforms would require advertisers or their agencies to formally declare whether content contains synthetically generated elements and confirm compliance with applicable labelling requirements at the point of submission.

Intermediaries that implement such systems, supported by appropriate record-keeping and risk-based checks, could be considered to have taken reasonable steps to ensure compliance with the Rules. This would enable a clearer allocation of responsibility, where liability for declaration rests with the advertiser, while platforms retain safe harbour protections on account of having operationalised due diligence processes. Such an approach may provide a scalable compromise in light of the practical limitations faced by intermediaries in independently verifying all content at scale.

4 Ambiguity Around Liability of Publishers

Questions were also raised regarding whether and how publishers would fall within the scope of the amended rules. While the amendments have been introduced under Part II of the IT Rules, which primarily governs intermediaries,⁴ stakeholders noted that the broadened interpretation of intermediaries, particularly those publishing or hosting SGI, could potentially bring traditional publishers within scope. For instance, news organisations that use AI tools to create or modify content, or that host advertisements and sponsored content incorporating AI-generated elements, may also be treated as intermediaries under the amended framework.

In the event that publishers are brought under the ambit of these rules, stakeholders identified two areas of concern. First, there are parallels with prior directions mandating self-declaration certificates for advertisements, which had posed significant operational challenges for the part of the industry which hosts such content.⁵ Similar obligations in the context of AI labelling are burdensome, particularly given that publishers are not equipped to independently verify the claims or nature of content provided by advertisers. Assigning such responsibility solely to publishers could be impractical and the consensus view in our discussion was that obligations should instead be distributed across the value chain.

Second, with respect to content created by publishers themselves, stakeholders noted that applying AI labelling requirements could have unintended consequences for audience trust. For instance, the use of AI in news production such as AI anchors or assisted content generation may require disclosure, but labelling could lead users to question the credibility of such content. Additionally, the widespread resharing of news across platforms raises concerns around downstream liability, particularly where labels may not persist. In this context, stakeholders reiterated on the need for clear regulatory guidance on applicability and responsibility.



An example of an AI-generated news anchor
Image Credits: [Business Standard](#)

5 Compliance Burden on Small Businesses

Stakeholders flagged that smaller brands and marketplace sellers may face disproportionate challenges in complying with AI labelling requirements. Unlike larger firms, these entities often lack dedicated policy teams or the technical capacity required to interpret and implement complex compliance obligations. As a result, navigating labelling requirements may require external advisory support or reliance on intermediaries such as advertising agencies, increasing both cost and operational dependency.

For smaller businesses, the compliance burden risks limiting market access by constraining their ability to rely on AI-driven tools for low-cost scaling. This is particularly significant in the Indian context,⁶ where generative AI tools are reducing the cost of content creation and digital advertising, thereby enabling MSMEs who typically lack dedicated marketing resources to undertake customer acquisition and scale digital outreach at materially lower cost.

6 Compliance Challenges for E-commerce Service Providers

Domestic E-commerce platforms hosting MSME sellers emphasised that the scale of their operations presents a challenge. With millions of sellers on such platforms, the introduction of additional obligations around SGI labelling, verification and takedown processes is difficult given the nature and diversity of the seller base, making uniform implementation and enforcement complex.

These challenges are compounded by low levels of regulatory awareness among sellers, where baseline compliances such as seller KYC registration and GST obligations are difficult to implement.⁷ This is not due to unwillingness, but limited clarity and understanding of regulatory requirements among small and first-time digital sellers. The issue is also linked to the composition of the seller base, which includes participants from government-linked initiatives and self-help groups, many of whom have limited digital literacy.

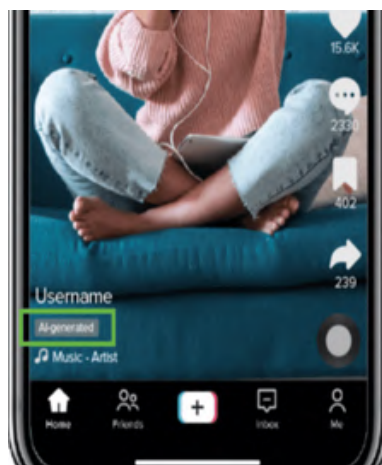
In such contexts, communicating and enforcing AI-related compliance requirements would be resource-intensive and difficult to operationalise. Additional complexities arise from the use of vernacular languages and the limitations of detection tools, which are insufficiently mature or accurate. In this backdrop, SGI labelling across all sellers is not reasonable.

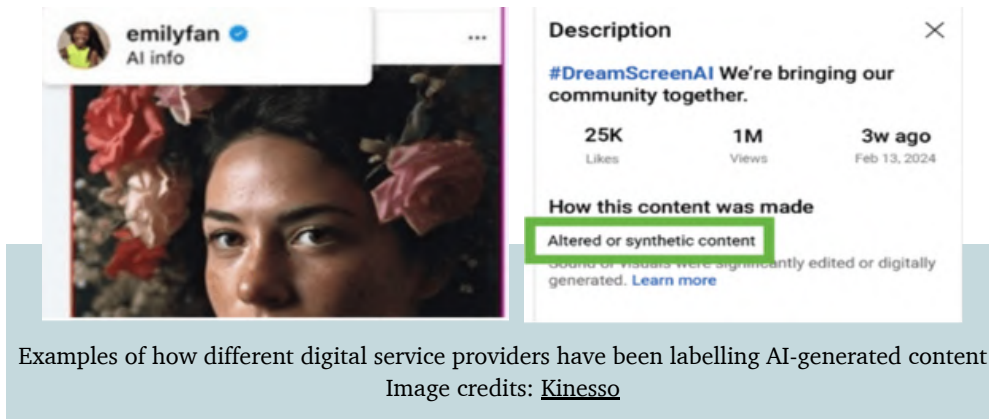
7 Technical Feasibility

Detection tools are engaged in a constant race with content generation systems, with the latter consistently outpacing detection capabilities, making risks of false positives, false negatives, and non-conformity at scale essentially unavoidable. Concerns were also raised around the durability of labelling as visual labels can be added to authentic content and omitted from SGI with equal ease and metadata frequently fails to persist as content moves across platforms.

At the same time, technology companies and platforms are actively developing solutions to address these challenges, including through collaborative initiatives such as the Coalition for Content Provenance and Authenticity (C2PA), which focuses on developing open, global technical standards for content authenticity. These efforts include provenance-based approaches such as metadata tagging at the point of content creation, interoperable frameworks, and techniques such as invisible watermarking. Such approaches are foundational in enabling content to be tracked across its lifecycle from creation to dissemination. Effective implementation of such technical solutions would require sufficient time for ecosystem-wide adoption of common standards. Drawing on international approaches such as South Korea's AI Basic Act, a defined 12-month transition period would allow industry participants to integrate provenance-based standards (such as C2PA metadata) into widely used creative tools. This would help ensure that, when enforcement begins compliance is anchored in verifiable technical signals rather than platform-side guesswork and will avoid disruption from abrupt implementations.

However, stakeholders underscored that no single technological solution can fully address the accuracy and coverage problem, and that a combination of proactive measures such as labelling and reactive tools such as detection systems would be required.





Examples of how different digital service providers have been labelling AI-generated content
Image credits: [Kinesso](#)

8 Ambiguities in Interpretation

Several key terms within the rules remain open to broad and at times inconsistent interpretation, creating uncertainty for implementation. Phrases such as “appears authentic,” “good faith,” and “reasonable and appropriate technical measures” were seen as inherently subjective, with limited clarity on how they would be applied in practice.

Uncertainty also extends to the scope of application of the rules, particularly in contexts such as one-on-one or private AI interactions. Additionally, stakeholders pointed to challenges in interpreting what constitutes permissible “basic aesthetic edits” versus content that would trigger labelling or compliance obligations, noting that such distinctions are difficult to assess without clear thresholds or contextual guidance.

In the absence of such clarity, stakeholders cautioned that platforms may adopt overly cautious approaches to preserve their safe harbour protections, including the removal or restriction of content without fully assessing context. This could have broader implications, including potential constraints on legitimate expression and a chilling effect on smaller businesses that rely on digital platforms for market access.

9 Economic Implications

The discussion also examined the broader economic implications of the rules, particularly in the context of India’s growing digital and creator-led economy. Referring to industry estimates, experts noted that India’s creator and brand-driven ecosystem could drive over \$1 trillion by 2030 in consumer spending.⁸ In this context, regulatory interventions if not carefully calibrated risk constraining this growth trajectory. They emphasised that AI tools have reduced the cost of content production, enabling brands to scale without incurring high production expenses. However, additional compliance requirements and labelling obligations are likely to increase costs, reduce return on investment, and undermine the core value proposition of AI as a low-cost scaling tool.

Cost implications are likely to disproportionately impact small businesses and individual entrepreneurs who rely on AI to overcome resource and time constraints. MSMEs highlighted that regulatory interventions affecting the use of digital and AI-driven tools could have broader

economic consequences, particularly if they increase compliance costs or limit the ability of such businesses to scale efficiently. They underscored the need to assess these rules not only from a content regulation perspective, but also through the lens of their impact on economic activity.

10 Impractical Takedown Timelines

Concerns were raised regarding the feasibility of the requirement to remove or disable access to content within three hours of receiving a court order or direction from an authorised government authority. This provision was not part of earlier drafts and was introduced without prior consultation, limiting the ability of industry to assess its operational viability. Many emphasised that such compressed timelines pose implementation challenges, particularly for entities that do not have the same level of technical and compliance infrastructure as large social media intermediaries.

E-commerce players noted that they are not similarly equipped, underscoring the need for differentiated obligations or carve-outs. This is further in focus with reports indicating the possibility of further reducing timelines to one hour, which would exacerbate existing challenges.⁹

Beyond operational constraints, stakeholders highlighted concerns around the rationale and implications of such timelines. While the objective appears to be addressing the rapid spread of harmful content, they noted the absence of corresponding deterrents targeting the source of such harms. For instance, bad actors continue to offer services to generate harmful content, including non-consensual intimate imagery, through private channels such as messaging platforms, often for payment and outside the regulatory net. At the same time, stakeholders highlighted potential free speech implications, noting that in the absence of adequate safeguards, platforms may be compelled to act hastily to meet strict timelines, leading to over-removal of content without sufficient contextual assessment.

11 Broadening the Definition of Intermediaries

The new rules depart from the IT Act's original conception of an intermediary as a passive conduit of third-party content. Participants noted that the obligations introduced under the rules which attach to any entity that creates, generates, modifies, alters, publishes, transmits, shares, or disseminates SGI collapse a critical distinction between a platform's own AI model generating content and a platform merely transmitting user-generated content that happens to be AI-made.

There was consensus that regulations should ideally distinguish between platforms that create content and those that merely host or transmit third-party content, and should base obligations on whether content is lawful or unlawful rather than whether it is AI-generated. Experts pointed out that recent FAQs had effectively expanded the scope further, bringing image-generating and content-generating tools within the definition of intermediaries, a determination that was not clearly supported by the text of the Rules themselves.¹⁰

The obligations in their current form are only meaningfully enforceable against a small number of large platforms that have inbuilt AI integration while bad actors operating through non-compliant platforms remain entirely outside the regime. Smaller domestic intermediaries, one on one interaction platforms, e-commerce marketplaces will find themselves caught in a net that was evidently designed with large global platforms in mind, with little clarity on how their specific operational contexts fit within the framework.

Taken together, these concerns pointed to a framework that is simultaneously overbroad in its definitional scope and underspecified in its practical obligations creating compliance uncertainty across the value chain while leaving key harms insufficiently addressed.

| Annexure: List of Attendees

1. Amitabh Kumar – Contrails AI
2. Anahida Bhardwaj – Data Security Council of India (DSCI)
3. Ananya Mishra – Internet and Mobile Association of India (IAMAI)
4. Anil Bhardwaj – Federation of Indian Micro and Small & Medium Enterprises (FISME)
5. Apar Gupta – Internet Freedom Foundation (IFF)
6. Aravamudhan K – JioStar
7. Archit Tripathi – Federation of Indian Micro and Small & Medium Enterprises (FISME)
8. Anushka Shrivastava – Stage OTT
9. Berges Malu - Koan Advisory
10. Dhawal Gupta – Microsoft
11. Dhruv Garg – Internet Governance and Policy Project (IGAP)
12. Jhanvi Anam – Internet Freedom Foundation (IFF)
13. Kavitha KK – OpenAI
14. Madhur Bafna – RedHoney Consulting
15. Meghna Bal – Esya Centre
16. Mekha Vijayakumar – Meesho
17. Mira Swaminathan – YouTube, Google
18. Nikhil Dahiya – Zepto
19. Nikhil Pahwa – Medianama
20. Prachi Bhuchar – Meesho
21. Rahil Chatterjee – Ikigai Law
22. Rajesh Mahapatra – Basispoint Insights
23. Rakesh Maheshwari – Policy Advisor, Cyber Regulations and Compliance
24. Rohan Sahu – Content Lens
25. Roudra Bhattacharya – HT Digital Streams
26. Sagar Deoskar – Snap Inc.
27. Saikat Dutta – Deepstrat
28. Shashwat Dwivedi – FRND
29. Shatak Sahu – USISPF
30. Vaibhav – Indian Express
31. Varun Ramdas – Microsoft
32. Vasudev Devadasan – Independent Legal Researcher
33. Vidhi Mukhi – Habibi Project

| Endnotes

- 1 Section 2(1)(wa) of the IT Act defines 'synthetically generated information' as audio, visual or audio-visual information which is artificially or algorithmically created, generated, modified or altered using a computer resource, in a manner that such information appears to be real, authentic or true and depicts or portrays any individual or event in a manner that is, or is likely to be perceived as indistinguishable from a natural person or real-world event.
- 2 MeitY Advisory on Compliance with IT Rules, 2021. (Available [here](#))
- 3 "The Impact of Visual Generative AI on Advertising Effectiveness", SSRN, (Available [here](#))
- 4 Part II of the IT Rules ("Due Diligence by Intermediaries and Grievance Redressal Mechanism") applies to intermediaries, while Part III ("Code of Ethics and Procedure and Safeguards in Relation to Digital Media") applies to publishers. The present amendments have been made to Rules 3, 4, and 5 under Part II, and therefore relate specifically to intermediary due diligence requirements.
- 5 "Supreme Court Mandates Self-Declaration by Advertisers, Ad Agencies Before Releasing Ads", Storyboard18, (Available [here](#))
- 6 "AI Ads in India Fuel Fear of Brand Dilution", LiveMint, (Available [here](#))
- 7 "India's MSME Compliance Problem is Not Resistance, It's Awareness", SMEFutures, (Available [here](#))
- 8 "India's Creator Economy Set to Drive over USD 1 Trillion in Consumer Spending by 2030", The Print, (Available [here](#))
- 9 "Govt May Cut Content Takedown Timeline to One Hour from Three", Business Standard, (Available [here](#))
- 10 MeitY, "Frequently Asked Questions (FAQs) on The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026", (Available [here](#))

KOAN