KNOWLEDGE PAPER

TRUST, TRACEABILITY, AND TRANSPARENCY:

India's 2026 IT Amendment Rules for AI-Generated and Synthetic Content

An analysis of the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026

FEBRUARY 2026

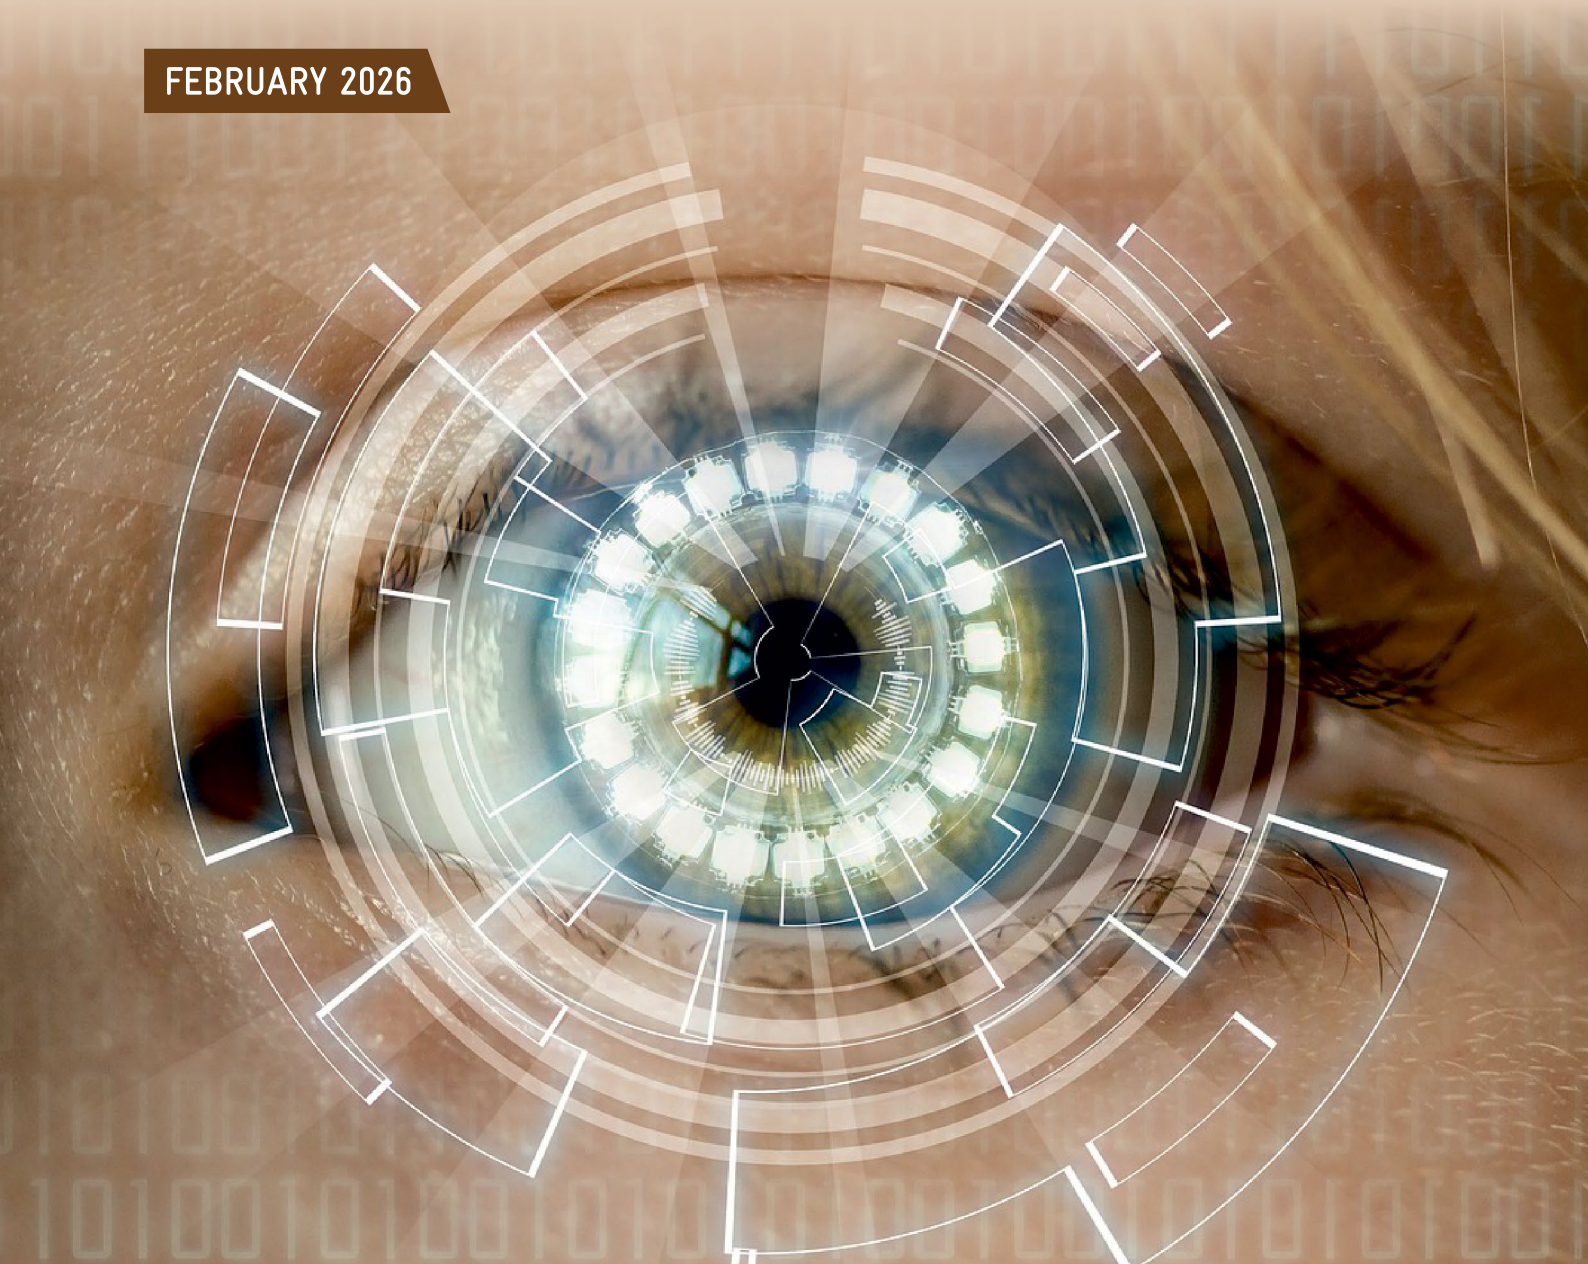

Foreword

Ms. Ranjana Khanna

Director General CEO
AMCHAM India

The Ministry of Electronics and Information Technology (MeitY) notification dated 10th February 2026 has responded with commendable amendments through the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, which came into force on 20th February 2026. These SGI Amendments represent a landmark step in establishing India as a global leader in AI content governance and set a template that other democracies are watching closely particularly as India aspires to build a Trillion USD digital economy built significantly on AI-led growth.

The three pillars of the SGI framework Trust, Traceability, and Transparency align closely with the principles that our member companies have long advocated and embedded into their global operations. Mandatory labelling of synthetic content, embedded metadata for provenance, and clear platform accountability are precisely the kind of enforceable, innovation-compatible standards that can make India's digital ecosystem safer and more resilient.

This knowledge paper details several provisions of the amendments, highlighting key industry challenges.

- The scope of intermediary due diligence.
- The preservation of safe harbour protections.
- Compliance & Reporting Timelines
- The treatment of pre-existing synthetic content on platforms are areas requiring clear and workable guidance.

Furthermore, it is suggested for effective implementation be carried out in a manner that least disrupts the existing digital ecosystem.

The U.S.–India strategic technology partnership, already deepened by shared imperatives around cybersecurity, critical infrastructure protection, and responsible AI, provides a powerful foundation for collaborative action. Our member companies bring deep expertise in threat detection, AI governance frameworks, content moderation at scale, and regulatory compliance across multiple jurisdictions. This collective wisdom is available to support India's efforts to implement the SGI Amendments in a manner that is effective, equitable, and commercially sustainable.

Representing over 400 American companies, AMCHAM India drives investment and technology across the Indian market through 20 sectoral committees. Its influence spans critical industries including IT and cybersecurity, defence and aerospace, and sustainable energy infrastructure. The chamber also plays a vital role in advancing healthcare, agriculture, and manufacturing, further strengthening the strategic and economic partnership between the United States and India. In this regard, US industries operating in India extend their full support to the Government's initiatives for regulatory framework in this current scenario.

The American Chamber of Commerce in India

AMCHAM India is the leading apex chamber of U.S. industry in India. Established in 1992, AMCHAM has over 400 U.S. companies as members and plays a pivotal role in fostering strong ties between the U.S. and India. The incumbent U.S. Ambassador to India is the Honorary President of AMCHAM. The chamber enjoys a close relationship with the U.S. Embassy and complete support in fulfilling its objectives.

Country Heads of leading U.S. companies constitute the elected national executive board. The chamber's mission is to assist member companies to succeed in India through advocacy, information, networking and business support services. Headquartered in New Delhi, AMCHAM extends its influence through regional chapters in Bengaluru, Chennai, Hyderabad, Kolkata, Mumbai and Pune.

AMCHAM India is a member of the U.S. Chamber of Commerce in Washington DC and the AmCham's of Asia Pacific.

Executive Summary

India's Ministry of Electronics and Information Technology (MeitY) issued the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendments Rules, 2026. This is a well-intentioned approach to addressing public harms at large arising from for AI-generated and synthetic content[1].

These amendments arrive at a defining moment for India's digital economy and India's aspiration for a Trillion USD digital economy[2] now rests significantly on AI-led growth. Ensuring that this AI-powered future is safe, trustworthy.

The SGI Amendments establish India as a global leader in AI content governance creating enforceable standards that balance user protection with innovation and setting a template that other democracies are watching closely.

Key Provisions[3]

- Introduces a comprehensive definition of 'Synthetically Generated Information' (SGI) covering AI-generated audio, visual, and audio-visual content that appears indistinguishable from reality
- Mandates labelling requirements for SGI, with prominent disclosure and embedded metadata
- Calls for enhanced due diligence obligations for intermediaries, especially Significant Social Media Intermediaries (SSMIs)
- Prohibits the creation, publication, or sharing 'Prohibited SGI' including deepfakes depicting real persons, non-consensual intimate imagery, and forged documents
- Requires compressed compliance timelines: 3 hours for government takedown orders, 2 hours for SGI-related grievances
- New victim rights including identity disclosure of SGI perpetrators through legal process

The Three Pillars

01 TRUST	02 TRACEABILITY	03 TRANSPARENCY
Protecting users and the public from deceptive AI - generated content through clear prohibitions, defined liability, and enforceable platform accountability.	Mandating embedded meta-data, unique identifiers, and provenance mechanisms so all synthetic content can be identified, attributed, and acted upon.	Requiring mandatory dis-closure, prominent labelling, and regular user notifications so audiences can identify and contextualize synthetic content.

1. https://timesofindia.indiatimes.com/technology/tech-news/not-just-rules-techno-legal-approach-needed-to-curb-ai-generated-harmful-content-it-minister-ashwini-vaishnaw/articleshow/128458119.cms
2. https://www.pib.gov.in/PressReleaseIframePage.aspx?PRID=1565669&reg=3&lang=2
3. Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026 – G.S.R. 120E dated 10th February 2026

Implications and Recommendations at a Glance

The SGI Amendments open a vital platform for ongoing stakeholder dialogue on AI governance. While their intent is sound, several provisions raise operational and legal questions that require regulatory clarity to be workable in practice. The table[4] below captures the five cross-cutting issues that industry and policymakers must address collaboratively.

No.	Provision	Implications	Request for Consideration
1.	**Notice to Users**	Users interact with platforms in multiple, often ad-hoc ways making it operationally complex for platforms to track notification frequency and confirm receipt across anonymous, multi-device, and guest users.	MeitY may clarify acceptable notification delivery methods, the definition of a 'user' for notification purposes, and what constitutes sufficient proof of notification.
2.	**Platform / Intermediary Due Diligence & Grievance Redressal**	Shifting SGI detection responsibility to intermediaries raises serious concerns: intermediaries may lack the subject matter and legal expertise to assess violations across all applicable laws; most lack the institutional or technical capacity to verify user-posted information; mandatory identity disclosure risks conflicting with established IT Act grievance principles; and compliance costs will rise significantly.	Clarity is needed on the scope of due diligence, minimum detection standards, and a safe harbour for good-faith compliance efforts. A tiered model based on intermediary size and risk is essential.
3.	**Compliance & Reporting Timelines**	The 3-hour (government orders) and 2-hour (user grievances) SGI timelines are operationally extremely stringent therefore raising serious concerns of technical feasibility. Sustained 24/7 compliance at this speed demands dedicated infrastructure that most intermediaries especially smaller platforms do not currently possess.	Structured dialogue on workable, tiered timelines differentiated by intermediary size, content risk level, and technical capacity is urgently needed before active enforcement begins.
4.	**Maintaining Safe Harbour**	Safe harbour is the cornerstone of India's internet economy. Compliance obligations that are unclear or technically unfeasible risk inadvertently eroding this protection generating unintended liabilities contrary to established legal principles and decades of internet jurisprudence.	Good-faith compliance must be explicitly and robustly linked to safe harbour preservation. Liability may follow wilful non-compliance, not good-faith technical limitations.
5.	**Retrospective Impact**	While platform and user liability is prospective, pre-existing SGI content on platforms may nonetheless require retrospective action creating unforeseen consequences for both platforms and users regarding historic liabilities and the scope of user rights.	Regulatory clarity on the treatment of pre-notification SGI content is needed urgently to give platforms and users certainty on historic obligations and any retrospective sweep.

4. https://www.business-standard.com/industry/news/sgi-labelling-norms-unimplementable-social-media-meity-125111302103_1.html

SECTION 1

Background and Problem Statement

The Challenge of Synthetic Media

The proliferation of generative AI technologies has fundamentally altered the information landscape. Tools capable of creating hyper-realistic synthetic audio, images, and videos are now widely accessible no longer confined to research laboratories or well-resourced actors. In 2023, 500,000 deepfake files circulated online. By 2025, that number had surged to over 8,000,000 a 16x increase in under two years. Over the same period, fraud attempts fuelled by synthetic media spiked 3,000%, with voice cloning emerging as the leading attack vector[5].

Key challenges addressed by the 2026 amendments include:

- **Deepfakes and Impersonation:** AI-generated content falsely depicting real individuals in compromising, misleading, or fabricated scenarios enabling reputational harm, fraud, and violation of personal dignity
- **Non-Consensual Intimate Imagery (NCII):** Synthetic intimate content created without the subject's consent, frequently weaponized for harassment, blackmail, and exploitation
- **Electoral Manipulation:** Synthetic media deployed to spread political misinformation or manipulate voter perception threatening the integrity of democratic processes
- **Forged Documents:** AI-generated official records, identity documents, and financial instruments used for fraud
- **Platform Accountability Vacuum:** The absence of a clear liability framework for intermediaries hosting or enabling harmful synthetic content

SGI Amendments address concerns and specifically synthetically generated content. While general prohibitions on unlawful content existed, they were not tailored to the unique challenges posed by AI-generated media particularly the speed of creation, ease of dissemination, and the technical difficulty of detection and attribution. The regulatory gap left victims without clear remedies, platforms without defined obligations, and law enforcement without jurisdictional clarity. A targeted amendment to the existing rules offered the fastest and most operationally certain path to addressing this gap.[6]

India's amendments arrive with urgency as the harms they address are already widespread and accelerating. Their reliance on the existing IT rules architecture means regulated entities can adapt within a familiar compliance framework. However, several obligations assume technical capabilities in detection, verification, and metadata embedding that are still maturing. The framework's long-term success depends on implementation guidance that keeps pace with the technology, not just the regulation.[7]

5. https://www.dhs.gov/sites/default/files/publications/increasing_threats_of_deepfake_identities_0.pdf
6. https://www.meity.gov.in/static/uploads/2025/10/065b6deb585441b5ccdf8be42502a49c.pdf
7. https://buffett.northwestern.edu/documents/buffett-brief_the-rise-of-ai-and-deepfake-technology.pdf

Regulatory Gap: The Global Context

Globally, there is no uniform regulatory approach to AI[8]. The U.S. has adopted ad-hoc, state-level approaches to specific harms such as NCII and election deepfakes and recently enacted a comprehensive federal legislation[9]. The EU AI Act addresses AI transparency requirements but does not constitute a dedicated synthetic media regime. Most major democracies are still in consultative phases.[10]

While India continues to deliberate on a separate comprehensive AI framework, the IT Act and Rules form the bedrock of digital regulation in the country. The SGI Amendments utilize this existing legislative infrastructure rather than waiting for standalone AI legislation to deliver rapid, targeted protections. This approach provides regulatory certainty for intermediaries already operating within the legal framework, minimizes transition costs, and demonstrates that agile, evidence-based AI governance is achievable without landmark new legislation.

India's first-mover status creates both influence and responsibility. Other jurisdictions will watch its implementation experience what works, what requires recalibration, how industry responds as they develop their own frameworks. This positions India to shape global AI content governance norms, provided its framework remains technically grounded, practically workable, and open to iterative refinement through stakeholder engagement.

8. https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-united-states
9. https://www.citizen.org/article/tracker-intimate-deepfakes-state-legislation/
10. https://spicyip.com/2025/11/deepfake-regulation-same-problem-different-approaches-yet-none-is-an-error-free-resolution.html

SECTION 2

Regulatory Context

Evolution of the IT Rules

The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 were enacted under Section 87 of the IT Act, 2000. They established the foundational, principle-based framework for intermediary liability, user safety, and content moderation in India. The 2026 SGI Amendments exercise powers under Section 87(1) and Section 87(2)(z) and (zg) of the IT Act which authorize the Central Government to make rules for intermediaries and for matters ancillary to technology use.

The amendments build on the 2021 framework through targeted modifications. Rather than creating a separate statute which could introduce fragmentation and legal uncertainty MeitY grafted a comprehensive SGI architecture onto the existing, well-understood rules structure, introducing new definitions rules 2(1) (wa), 2(1)(ca)), two new sub-rules (2(1A), 2(1B)), and new due diligence obligations within existing rule numbers.

The Consultation Process[11]

In October 2025, MeitY released draft amendments for public consultation inviting comment from technology companies, civil society organisations, legal experts, and academic institutions. The final amendments, notified February 10, 2026, reflect meaningful stakeholder input.

The consultation process reflects MeitY's recognition that SGI regulation touches complex technical, legal, and commercial interests that require genuine multi-stakeholder input to be practically workable. The changes from draft to final rules show this process was substantive, not performative.

Four significant changes from draft to final rules carry important compliance implications:

- **Narrowed SGI definition:** Pure text content is explicitly excluded from scope protecting AI writing tools, chatbots, and document generation platforms from inadvertent capture. However, the boundary between audio-visual SGI and text with embedded AI-generated media will require further interpretive guidance as multimodal AI tools proliferate.
- **Removed fixed 10% label size:** Replaced with a flexible 'prominence' standard giving intermediaries design freedom while maintaining the consumer-facing transparency objective. Platforms must document their labelling design rationale to demonstrate that prominence standards are met.
- **Broader application:** Intermediaries that enable publication or transmission of SGI (not just creation) are now captured significantly expanding the compliance universe beyond generative AI platforms to include social networks, messaging apps, and cloud storage services.
- **Enhanced safe harbour:** Compliant takedowns are expressly protected under Section 79 of the IT Act removing the chilling effect that might otherwise cause platforms to delay action for fear of content removal liability.

11. https://www.thehindu.com/sci-tech/technology/what-has-government-laid-down-on-ai-labelling-explained/article70633146.ece

SECTION 3

Core Definitions and Scope

Synthetically Generated Information — Rule 2(1)(wa)

> **"Synthetically generated information means 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."**

The definition is structured to capture harm through effect rather than technique. The phrase 'likely to be perceived as indistinguishable' adopts an audience-centred, objective standard focusing on how a reasonable viewer experiences the content, not on how sophisticated the underlying AI model is. By limiting scope to audio, visual, and audio-visual content (excluding pure text), the drafters focused the framework on the highest-harm format types where deception is most visceral and most difficult to detect. The definition is designed to remain durable as both generation and detection technologies evolve.

Four key definitional elements carry direct compliance significance: the content must be created algorithmically or artificially using a computer resource; it must appear real, authentic, or true; it must depict individuals or events; and it must be indistinguishable or likely to be perceived as indistinguishable from a natural person or real event. The 'likely to be perceived' standard introduces interpretive uncertainty at the margins: what appears convincing to some audiences may not to others. MeitY guidance or judicial interpretation establishing consistent benchmarks will be essential particularly for imperfect but still potentially deceptive content.

Critical Carve-Outs: What is NOT SGI

Three provisos explicitly exclude legitimate AI use from the SGI framework:

- **Proviso A — Routine Technical Enhancements:** Colour correction, noise reduction, transcription, compression, and similar edits that do not materially alter content substance, context, or meaning. (Examples: brightness adjustment, audio noise reduction, video compression for streaming)
- **Proviso B — Good-Faith Document Creation:** AI-generated presentations, PDFs, educational materials, and template-based content provided no false document, or false electronic record is created. (Examples: AI-generated business slides, educational diagrams, report templates)
- **Proviso C — Accessibility and Quality Improvements:** Subtitles, captions, audio descriptions, translation, searchability improvements where the material substance of the underlying information is not altered. (Examples: AI subtitles, audio descriptions for visually impaired users, language translation)

These carve-outs reflect the drafters' recognition that AI-assisted content modification is pervasive in modern digital workflows. Without them, routine use of AI in content production could inadvertently trigger labelling and due diligence obligations that serve no protective purpose. The provisos operationalize the principle that the SGI framework targets deceptive synthetic content specifically not AI-assisted content creation in general.

The carve-outs provide important safe spaces for product teams building AI tools. However, their application requires careful case-by-case analysis. The line between 'material alteration' (within scope) and 'technical enhancement' (excluded) is not always self-evident particularly for tools that combine routine enhancement with generative synthesis, such as AI-powered video up-scalers or style-transfer applications. Organizations may maintain documented classification rationale for their products to support any future compliance defense.

Definitional Extensions — Rules 2(1A) and 2(1B)

Rule 2(1A) extends the meaning of 'information' throughout the IT Rules' unlawful content provisions to include SGI ensuring the full weight of existing content moderation obligations applies to synthetic content. Rule 2(1B) provides express safe harbour protection: removing or disabling access to SGI in compliance with the rules shall NOT constitute a violation of Section 79(2)(a) or (b) of the IT Act.

Rule 2(1B) is practically significant for intermediaries: it means platforms that act in good faith on SGI takedown orders even where the 'prohibited' determination involves judgement calls are protected from downstream liability. This may encourage proactive enforcement rather than defensive inaction. The corollary is equally important: intermediaries that fail to act despite actual knowledge, or that lack the systems to develop such knowledge, cannot claim safe harbor protection.

SECTION 4

Obligations for Intermediaries

The SGI Amendments create a three-tier framework of obligations calibrated to the type and scale of intermediary imposing proportionate requirements based on both the risk created by the intermediary's activities and its capacity to implement them.

4.1 All Intermediaries — Universal Baseline

Mandatory User Notifications — Rule 3(1)(c)

Every intermediary must notify users at least once every three months in a simple and effective manner, in English or any language specified in the Eighth Schedule to the Constitution covering: (i) the consequences of creating, publishing, or sharing prohibited SGI; (ii) a clear definition of SGI and examples of prohibited content; and (iii) the grounds for account suspension or termination for non-compliance.

Regular, mandated user notification closes the information gap between platform obligations and user awareness. Most users are unaware of the specific legal and platform-level consequences of creating or sharing synthetic content. Quarterly notifications create a documented mechanism for establishing adequate notice both a compliance objective in its own right and a foundation for enforcement action when violations occur. The Eighth Schedule language requirement ensures accessibility across India's linguistic diversity.

The quarterly cadence is operationally demanding at scale. Platforms with hundreds of millions of users must develop infrastructure to track notification delivery across diverse engagement patterns including anonymous browsing, logged-out sessions, and multi-device users. The definition of 'user' for notification purposes, and what constitutes sufficient proof of delivery, will require regulatory clarification. Platforms may build notification confirmation mechanisms into their UX flows and maintain audit logs as compliance evidence.

4.2 Intermediaries Enabling SGI — Heightened Due Diligence

Prohibition on Prohibited SGI — Rule 3(3)(a)(i)

Intermediaries enabling the creation, generation, modification, publication, transmission, or dissemination of SGI must ensure users do not create, publish, or share prohibited SGI content containing CSEAM, NCII, forged documents, weapons-related content, or content falsely depicting real persons or events. Intermediaries must take expeditious action upon actual knowledge and deploy automated tools to prevent generation or sharing of such content.

The prohibited category targets the highest harm use cases for synthetic content those where harm is severe, consent of those depicted is absent, and the public interest in suppression is unambiguous. The 'actual knowledge' trigger preserves the intermediary's safe harbour without creating strict liability for all content on a platform. The requirement for automated detection tools acknowledges a fundamental operational reality: at the scale of modern AI-generated content, manual moderation alone cannot be effective.

The automated detection obligation is the most technically demanding in the framework. No current detection system achieves the accuracy required to act on its outputs without significant false positive rates. Platforms must invest in or licence state-of-the-art detection capabilities, build robust appeals mechanisms to address false positives, and document their detection methodology to demonstrate 'reasonable and appropriate' measures. Smaller intermediaries face

disproportionate compliance costs hence the case for a tiered standard based on platform size and risk profile is compelling and may be addressed in MeitY guidance.

Mandatory Labelling of Lawful SGI — Rule 3(3)(a)(ii)

For lawful SGI, intermediaries must ensure prominent on-screen labelling for visual content; prominently prefixed audio disclosures; embedded permanent metadata or unique identifiers enabling identification of the intermediary's computer resource (where technically feasible); and prohibition on enabling users to remove, alter, or suppress labels or markers.

Labelling serves the transparency pillar directly: it ensures that even lawful synthetic content — AI-generated art, entertainment, educational simulations is visible to audiences as such. Embedded metadata serves the traceability pillar: unique identifiers linked to the platform's infrastructure create an audit trail for retrospective attribution even after content has been shared across platforms. The anti-tampering obligation ensures bad actors cannot defeat the system by stripping metadata or overlaying labels.

The metadata requirement is technically forward-looking but creates near-term implementation challenges. Standards for AI content provenance metadata are still being developed by bodies such as the Coalition for Content Provenance and Authenticity (C2PA). Intermediaries may engage actively with these standards processes to ensure India's requirements are met by emerging interoperable protocols. The 'technically feasible' qualifier provides important flexibility: platforms that cannot currently embed metadata without disrupting existing workflows are not in breach, provided they demonstrate good-faith efforts to implement as technology matures.

4.3 Significant Social Media Intermediaries (SSMIs)

SSMIs — intermediaries whose registered user base in India exceeds the threshold under Rule 2(1)(w) — bear maximum obligations under Rule 4(1A): (i) a mandatory user SGI declaration system at upload; (ii) technical verification of user declarations using metadata, signals, and appropriate tools; and (iii) clear and prominent labelling of all confirmed SGI content.

SSMIs are singled out for maximum obligations because of their combination of scale, technical capacity, and social impact. A platform with tens of millions of users has both the greatest potential for harm — synthetic content spreads further and faster — and the greatest resources to implement sophisticated detection and verification. The two-layer system (user declaration + platform verification) makes SGI identification more robust than either measure alone and places primary responsibility where technical capacity is highest.

The verification obligation requires SSMIs to build or license AI detection systems capable of cross-checking user declarations against content itself. SSMIs that already deploy AI content moderation infrastructure will need to extend those systems to specifically flag and label SGI. The user declaration UX design matters considerably: a frictionless, clearly worded declaration toggle will achieve higher self-reporting accuracy than a buried or confusing prompt.[12]

12. https://www.khaitanco.com/thought-leadership/MeitY-notifies-the-IT-Amendment-Rules-2026

SECTION 5

Compliance Timelines and Enforcement

The SGI Amendments compress response timelines for SGI-related content moderation and introduce a suite of enforcement mechanisms for violations.

Action Required	Timeline	Applicable Rule
Government takedown orders (SGI content)	**3 hours**	Rule 3(1)(d)
User grievances — SGI-related	**2 hours**	Rule 3(2)(b)
General government orders	36 hours	Rule 3(1)(d)
General user grievances (acknowledgment)	24 hours	Rule 3(2)
General user grievances (resolution)	15 days	Rule 3(2)
Periodic user notifications	Once every 3 months	Rule 3(1)(c)

The 3-hour and 2-hour timelines for SGI-specific actions reflect a fundamental asymmetry in the harm dynamic of synthetic content: a deepfake can cause irreversible reputational, personal, or electoral damage within hours of publication. Standard 36-hour or 15-day windows appropriate for many content moderation scenarios are inadequate to prevent the most acute harms from viral synthetic content. The compressed timelines signal that the government views SGI harm as categorically urgent, and that the regulatory framework must reflect this.

These are the most operationally challenging aspects of the SGI framework. They require 24/7 monitoring capability, pre-established escalation protocols, dedicated SGI review queues, and the organizational infrastructure to respond authoritatively to government orders outside business hours. For large SSMIs with existing trust and safety operations, this is an extension of current capability. For smaller intermediaries, it may require significant new investment or the use of third-party managed compliance services. Industry stakeholders have rightly highlighted that implementation of these timelines in their current form is extremely stringent tiered, and capacity-sensitive standards are needed.

Enforcement Mechanisms

- **Account Suspension / Termination:** For creation or sharing of Prohibited SGI upon confirmed violations
- **Identity Disclosure:** Victims may obtain perpetrator identity through legal process enabling civil and criminal remedies against bad actors operating anonymously
- **Mandatory Reporting:** To law enforcement under applicable provisions of the Bharatiya Nagarik Suraksha Sanhita (BNSS) for grievous offences
- **Criminal Penalties:** Under relevant provisions of the IT Act and other applicable laws

The enforcement architecture creates meaningful deterrence without imposing strict liability that would chill platform investment in India. Account-level sanctions are proportionate and reversible. Criminal liability is reserved for the most serious violations, channelled through established legal processes. Safe harbour preservation under rule 2(1B) ensures platforms that invest in compliance are protected from the liability exposure that would otherwise make operating in India commercially unattractive.

Safe harbor is the linchpin of the framework's commercial viability. Compliance obligations that are unclear, technically unfeasible, or disproportionate risk inadvertently eroding this protection generating unintended liabilities contrary to established legal principles and decades of internet jurisprudence. Maintaining a clear, workable link between good-faith compliance and safe harbor preservation is not optional; it is essential to the framework achieving its intended outcomes. Platforms may document all compliance efforts meticulously as evidence of good-faith engagement.

SECTION 6

Strategic Implications and Opportunities

Global Leadership in AI Governance

India's SGI Amendments represent an engagement opportunity for open global dialogue and discourse on frameworks for AI Generated Content. They balance several competing imperatives: user protection (clear prohibitions on harmful synthetic content with enforceable penalties); innovation preservation (carve-outs for legitimate AI use in education, accessibility, and creativity); platform accountability (tiered obligations matching platform capabilities and scale); and technical feasibility (flexible implementation standards that adapt to evolving AI detection technology).

India's unique position as both a major AI consumer and a growing AI producer, with a 1.4-billion-person digital population, a world-class technology sector, and significant AI research capacity gives its regulatory choices disproportionate global significance. The SGI Amendments reflect a governance philosophy grounded in the New Delhi Declaration's principles: that technological advancement must be matched with meaningful user protection, and that developing nations can and may shape global AI governance norms rather than merely adopt them.

India's first-mover status creates both influence and risk. The influence is real: other jurisdictions will use India's implementation experience as a reference point as they develop their own frameworks. The risk is that provisions that prove unworkable in practice become entrenched before course corrections can be made. A feedback loop between implementation experience and regulatory iteration through structured stakeholder dialogue, compliance data collection, and regular review is essential to converting first-mover advantage into durable regulatory leadership.[13]

Alignment with International Frameworks: BTA and TRUST[14]

The SGI Amendments were notified in the same month as the U.S.-India Bilateral Trade Agreement (BTA) and the TRUST (Transforming the Relationship Utilizing Strategic Technology)[15] initiative, a bilateral AI cooperation framework between India and the U.S., already exists. This timing creates structured opportunities for regulatory harmonization before divergence becomes entrenched.

Regulatory fragmentation between India and the U.S. different SGI definitions, incompatible metadata standards, misaligned enforcement mechanisms would impose significant compliance costs on globally operating platforms and reduce the effectiveness of cross-border enforcement. The BTA and TRUST initiative provide the political and institutional infrastructure for avoiding this fragmentation proactively, rather than attempting costly harmonization after the fact.[16]

Four specific collaboration opportunities with direct commercial and governance implications:

- **Joint Standards Development**: India and the U.S. can co-author interoperable SGI and deepfake definitions, forming the basis of bilateral and multilateral AI content standards reducing compliance costs for global platforms operating in both markets

13. https://sflc.in/press-coverage/the-hindu-businessline-quoted-sflc-ins-analysis-on-why-meitys-draft-sgi-rules-risk-overstepping-the-it-act-and-curbing-online-free-expression/
14. https://www.mea.gov.in/bilateral-documents.htm?dtl/39066#:~:text=18.,technologies%20and%20reduce%20regulatory%20barriers
15. Ibid
16. https://www.mea.gov.in/bilateral-documents.htm?dtl/39066#:~:text=18.,technologies%20and%20reduce%20regulatory%20barriers

- **Technical Protocol Harmonisation**: Joint development of AI content provenance protocols building on C2PA/Content Credentials standards, enabling platforms to comply with both Indian and U.S. requirements through a single technical implementation[17]
- **Cross-Border Enforcement**: A bilateral mutual legal assistance framework for SGI enforcement, enabling victims to pursue perpetrators across jurisdictions and closing the regulatory arbitrage gap that bad actors currently exploit
- **Democratic AI Governance Model**: Together, the IT Rules 2026, BTA, and TRUST initiative can establish a template for AI content governance rooted in transparency and accountability distinguishable from surveillance-led approaches and potentially adoptable at G20 or UN level

Commercial Opportunities

The SGI framework creates substantial commercial opportunities across several categories:

- **AI Detection Technology:** Growing market for compliant detection, verification, and labelling tools capable of operating at platform scale with high accuracy that India's compliance requirements will help accelerate
- **Content Authentication Services:** Third-party certification and metadata embedding services analogous to SSL certificate authorities in the web security space
- **Compliance Consulting:** Legal and technical advisory for intermediaries navigating the new framework, designing defensible compliance programs, and responding to enforcement actions
- **Infrastructure Investment:** U.S. and global investment in India's AI data center infrastructure, enabled by the regulatory clarity the SGI Amendments provide

17. https://cyberpeace.org/resources/blogs/c2pa-an-international-standard-for-media-provenance-in-india

SECTION 7

Conclusion and Recommendations

The SGI amendments open a new chapter on AI governance globally. This dialogue must be expanded to incorporate views of the industry to ensure seamless compliance and concrete enforcement. Addressing disinformation, deepfakes and unlawful synthetically generated content requires coordinated and concerted efforts to create safety and trust. The framework's success, however, will depend on technically achievable obligations, robust implementation guidance, and ongoing stakeholder dialogue.

India's SGI Amendments 2026, the BTA, and the TRUST initiative together give India and the U.S. a shared starting point and the tools to build the world's first democratic framework for honest, trustworthy AI content.

Key Takeaways[18]

- The rules balance user protection with innovation prohibiting harmful synthetic content while preserving legitimate AI use through carefully crafted definitional carve-outs
- Compliance obligations are tiered and proportionate, with SSMIs bearing the greatest responsibility for detection, verification, and labelling infrastructure
- Compressed timelines demand significant technical and operational investment tiered, capacity-sensitive standards are needed for the framework to be workable across intermediary types
- Safe harbor must remain robustly maintained, the link between good-faith compliance and liability protection must be explicit and achievable
- The rules create meaningful opportunities for U.S.-India collaboration on democratic, transparency-first AI governance standards with global reach

Recommendations for Intermediaries

- Invest in automated detection and verification infrastructure and treat this as a strategic technology investment, not a pure compliance cost
- Engage with C2PA and other international content provenance standards bodies to ensure technical implementations are interoperable and future-proofed
- Document all compliance efforts meticulously good-faith, evidence-backed compliance is the foundation of safe harbor protection

Recommendations for Policymakers[19]

- Issue comprehensive implementation guidance covering detection standards, notification delivery methods, treatment of pre-existing SGI, and safe harbor thresholds

18. https://www.nishithdesai.com/research-and-articles/hotline/technology-law-analysis/ai-generated-content-and-combating-deepfakes-what-indias-new-rules-mean-for-global-platforms-15532

19. https://hasgeek.com/fifthelephant/2025-winter/sub/shaping-indias-ai-regulations-a-dialogue-on-meitys-RAgmbrakqZFnec6FXhzzT7

- Develop a tiered compliance model differentiating obligations by intermediary size and technical capacity smaller platforms must not be forced into non-compliance by operationally unfeasible requirements
- Establish a structured stakeholder dialogue mechanism with regular technical roundtables with industry, civil society, and academic experts to collect implementation data and enable timely framework calibration
- • Engage U.S. counterparts through BTA and TRUST frameworks to develop interoperable SGI definitions, metadata standards, and cross-border enforcement arrangements before divergence becomes entrenched

Recommendations for Technology Providers[20]

- Develop India-compliant detection, labelling, and metadata solutions as the Indian market's scale makes compliance investment commercially significant
- Pursue TRUST-certified AI content certification to serve both Indian and U.S. markets simultaneously, leveraging regulatory alignment as a compliance efficiency opportunity
- Explore infrastructure investment opportunities in India's rapidly growing AI ecosystem, where the SGI Amendments contribute to the regulatory clarity that long-term capital deployment requires

20. https://community.nasscom.in/communities/public-policy/information-technology-intermediary-guidelines-and-digital-media-ethics

Appendix A: Defined Terms

Synthetically Generated Information (SGI): Audio, visual, or audio-visual information artificially or algorithmically created, generated, modified, or altered using a computer resource, appearing to be real, authentic, or true and depicting individuals or events in a manner indistinguishable or likely to be perceived as indistinguishable from reality. (Rule 2(1)(wa))

Prohibited SGI: SGI containing CSEAM, NCII, sexually explicit content, forged documents, false electronic records, weapons-related content, or content falsely depicting natural persons or real-world events through deceptive misrepresentation.

Intermediary: Any person who on behalf of another person receives, stores, or transmits electronic records, or provides any service with respect to such records. (Section 2(1)(w), IT Act 2000)

Significant Social Media Intermediary (SSMI): A social media intermediary with registered users in India exceeding the threshold notified by the Central Government under Rule 2(1)(w).

Safe Harbor: Legal immunity under Section 79 of the IT Act, 2000 shielding compliant intermediaries from liability for user-generated content. Rule 2(1B) expressly preserves this protection for SGI-compliant takedown actions.

CSEAM: Child sexual exploitative and abuse material.

NCII: Non-consensual intimate imagery intimate or sexually explicit content created or shared without the subject's consent.

MeitY: Ministry of Electronics and Information Technology, Government of India the nodal ministry responsible for IT policy and regulation.

BTA: U.S.-India Bilateral Trade Agreement a comprehensive trade framework establishing the basis for expanded economic cooperation between the United States and India.

TRUST Initiative: 'Transforming the Relationship Utilizing Strategic Technology) a bilateral U.S.-India AI cooperation framework promoting secure, trustworthy, and interoperable AI systems.

C2PA / Content Credentials: Coalition for Content Provenance and Authenticity an open standards body developing specifications for AI content provenance and metadata embedding to enable authentication of digital content.

BNSS: Bharatiya Nagarik Suraksha Sanhita India's updated code of criminal procedure, under which mandatory reporting obligations for grievous SGI offences apply.

JAM Trinity: Jan Dhan – Aadhaar – Mobile: India's foundational digital public infrastructure enabling digital payments and public service delivery at scale.

Eighth Schedule: The schedule to the Constitution of India listing the 22 officially recognized languages of India. User notifications under Rule 3(1)(c) must be made in English or any Eighth Schedule language.

AMERICAN CHAMBER OF COMMERCE IN INDIA

PHD House, 4th Floor, 4/2, Siri Institutional Area, August Kranti Marg, New Delhi - 110 016

Tel: 91-11-3548 0630, 4650 9413 | **Fax:** 91-11-3548 0631

Email: amcham@amchamindia.com | **Website:** www.amchamindia.com