Business Software Alliance

Tuesday, November, 04, 2025

Ministry of Electronics and Information Technology (MeitY)
Government of India,
New Delhi

E-mail: itrules.consultation@meity.gov.in

Cc: Shri. Ajit Kumar, Joint Secretary, MeitY e-mail: js-akumar@meity.gov.in

BSA'S COMMENTS ON THE DRAFT AMENDMENTS TO THE INFORMATION TECHNOLOGY (INTERMEDIARY GUIDELINES AND DIGITAL MEDIA ETHICS CODE) RULES, 2021

Respected team,

Greetings! On behalf of the Business Software Alliance (**BSA**),[1] we welcome MeitY's initiative to address synthetically generated information (**SGI**) risks through the draft amendments (**Draft Amendments**)[2] to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 (**2021 Rules**), and write to provide BSA's feedback on the Draft Amendments.

BSA recognizes the serious challenges posed by the growing misuse of SGI, including deepfakes, misinformation, and content that can mislead users, cause harm, violate privacy, and threaten national integrity.[3] We support the development and deployment of reliable content authentication and provenance mechanisms that can help users identify the history and origin of SGI. These tools can help consumers tell whether content is human- or AI- generated and prevent misinformation. At the policy level, BSA supports holistic solutions that prevent harm through a coordinated, multi-stakeholder approach involving industry, government, and civil society. To this end, we provide the following recommendations for the Draft Amendments below.

Summary of Recommendations

- **Avoid Prescribing Transparency Requirements that Create a Preference for Visible Labelling:** MeitY should support flexible and appropriate technical solutions to mark AI-

[1] The Business Software Alliance (www.bsa.org) is the global trade association of the enterprise software industry, representing companies that are leaders in artificial intelligence, cybersecurity, cloud computing, quantum, and other breakthrough technologies. We work in over 20 markets in the US, Europe, and Asia, advocating for policies that build trust in technology so that every industry sector and the public can benefit from innovation.

BSA's members include: Adobe, Alteryx, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Informatica, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

[2] 2025 Draft Amendment to the Information Technology Intermediary Guidelines, 2021: https://www.meity.gov.in/static/uploads/2025/10/9de47fb06522b9e40a61e4731bc7de51.pdf

[3] MeitY, Explanatory Note, 2025 Draft Amendment to the Information Technology Intermediary Guidelines, 2021: https://www.meity.gov.in/static/uploads/2025/10/8e40cdd134cd92dd783a37556428c370.pdf

generated or adjusted outputs in machine-readable formats ***without mandating visible display requirements***.

- **Avoid a One-Size-Fits-All Approach to Addressing SGI Risks:** The Draft Amendments should specify that upstream entities in the AI value chain are not responsible for transparency obligations that should be borne by companies providing AI tools and services to downstream users.
- **Support Industry-Led Internationally Recognized Standards for Content Authenticity and Provenance:** The Draft Amendments should specify the Coalition for Content Provenance and Authenticity (**C2PA**) standard as an option to meet the Draft Amendment's labeling requirements.[4]
- **Remove the Broad Prohibition on Modifying or Removing Labels/Identifiers:** Meity should remove the prohibition on "enabling" the modification, suppression, or removal of such label, permanent unique metadata or identifier.
- **Narrow the Definition of "Synthetically Generated Information":** MeitY should narrow the definition of SGI to content intended or likely to mislead or cause harm and exclude benign or incidental modifications and content generated within enterprise business environments from the scope of SGI.

Avoid Prescribing Transparency Requirements that Create a Preference for Visible Labelling

BSA appreciates that MeitY has aimed to support multiple methods of transparency in identifying AI-created content by focusing on either (1) labelling AI-generated content, or (2) embedding AI-generated content with permanent unique metadata or identifiers. However, we have significant concerns with the requirement to visibly display either a label or metadata/identifier. This visual display requirement effectively requires visible labelling and forecloses the use of globally supported metadata-based transparency methods, because metadata by nature is embedded within content (and not visible) rather than visibly displayed atop content.

Requiring visible labelling raises a range of significant concerns. Visible watermarks or labels are easily removed through basic editing tools, undermining their effectiveness as a reliable authentication mechanism by themselves. Moreover, such visible markers diminish the enjoyment and appreciation of content, making local works less competitive relative to content from other countries that do not require visible watermarks. Content credentials, embedded metadata, and visible labelling each offer different approaches to transparency, and MeitY should allow businesses to select the most appropriate solution for their use cases without creating a regulatory preference that favors one method over another.

Recommendation: MeitY should support flexible and appropriate technical solutions to mark AI-generated or adjusted outputs in machine-readable formats without mandating visible display requirements. To achieve this, MeitY should:

- Remove the visual display requirement
- Focus on ensuring that authentication mechanisms are technically robust and accessible through appropriate tools and platforms

This approach would provide industry with flexibility to deploy evolving technical standards that offer more durable protection against manipulation while preserving content integrity.

[4] Coalition for Content Provenance and Authenticity (C2PA), https://c2pa.org/

Avoid a One-Size-Fits-All Approach to Addressing SGI Risks

MeitY has attempted to provide different responsibilities for intermediaries based on whether they offer services that generate or modify SGI (addressed in Section 4) or if they allow for display or publication of such content (addressed in Section 5).

However, the Draft Amendments lack crucial distinctions between business-to-business (**B2B**) or enterprise services and business-to-consumer (**B2C**) entities. While the *Explanatory Note* issued by MeitY outlines the impact of these obligations for all "public-facing AI-generated media", the Draft Amendment should focus on SGI that is consumer facing, as this is where the risks of AI-generated or adjusted information lie. We urge MeitY to consider the differences in the roles and functions of different companies involved in the AI supply chain when prescribing obligations related to SGI.

Key service-level, technical, functional, and user-based distinctions mean that all companies do not have the same ability to address SGI-related harms, nor do they pose the same risks. Enterprise software services, for instance, pose limited risk to user safety and public order given the size and nature of their user base and the fact that they do not provide services directly to consumers. A framework that applies uniform obligations without accounting for these fundamental differences, risks imposing disproportionate burdens on entities with limited capacity to influence end-user behavior while potentially under-regulating higher-risk consumer-facing services. In particular, the Draft Amendments should ensure that companies that enable the generation or modification of content are not responsible for transparency obligations when other companies use such AI systems to provide products or services to downstream users.

Recommendation: The Draft Amendments should specify that an intermediary that offers computer resources that may enable, permit, or facilitate the creation, generation, modification, or alteration of synthetically generated information ***is not responsible*** for transparency obligations when other companies use such AI tools to provide products or services to downstream users. Such obligations should rest with the entities that use AI tools to generate content.

Support Industry-Led Internationally Recognized Standards for Content Authenticity and Provenance

We support efforts by the Content Authenticity Initiative (**CAI**) to promote the Coalition for Content Provenance and Authenticity (**C2PA**) standard. We encourage MeitY to ensure that use of C2PA will be one option for companies to satisfy requirements imposed by the Draft Amendments.[5]

The C2PA standard, which is expected to be approved by the International Standards Organization this year, is openly available. Anyone, including governments, can use this standard to incorporate digital provenance information into their products and processes. Creators can indicate whether AI was used in their work and how it was used — the content credentials will display information about how the work was created, the date it was created, and any edits that were made along the way. This standard will help consumers decide what content is trustworthy and promote transparency around the use of AI. The CAI approach provides secure, indelible provenance. Importantly, the C2PA combines different technical tools — including invisible

[5] Content Authenticity Initiative, How it works: https://contentauthenticity.org/how-it-works

watermarking, metadata, digital fingerprinting — to provide a robust, durable, and flexible solution.[6]

Many other countries are adopting or considering adopting provenance standards, including Japan, Australia, Singapore, and the European Union. Embracing open standards like that developed by C2PA facilitates international interoperability and enhances the integrity of digital content ecosystems. On the other hand, developing local standards for content authentication could limit the ability of businesses to leverage widely accepted internationally recognized standards, such as C2PA, and could lead to increased compliance burdens for companies operating in multiple jurisdictions.

Recommendation: Specify that using the C2PA standard is one option to meet the requirements created by the Draft Amendments. This will promote global approaches to user transparency, helping India leverage transparency tools utilized worldwide.

Remove the Broad Prohibition on Modifying or Removing Labels/Identifiers

The prohibition on intermediaries from modifying, suppressing, or removing labels or identifiers raises significant practical challenges. While BSA recognizes that it is important that platforms do not strip content credentials, watermarks, or metadata, the Draft Amendment's prohibition on "enabling" modification, suppression, or removal of such information may not be practical or feasible to implement.

Given the scale and automated nature of content processing on digital platforms and businesses, it is challenging to monitor, verify, or control disclosure practices across all user-uploaded content. Prohibiting companies from "enabling" modification or removal of labels or watermarks would particularly burden model developers by forcing them to police downstream use beyond their control. This is especially challenging in B2B contexts where AI capabilities are integrated into complex, multi-layered applications and workflows. Developers could be required to contractually obligate licensees to maintain labeling integrity and monitor compliance throughout distribution chains, effectively becoming guarantors of third-party practices. Instead, the Draft Amendments should focus transparency obligations on entities that directly generate or deploy SGI, rather than imposing unenforceable policing requirements on model developers or enterprise AI developers.

Recommendation: Remove the prohibition on "enabling" the modification, suppression, or removal of such label, permanent unique metadata or identifier.

Narrow the Definition of "Synthetically Generated Information"

The proposed definition for SGI extends to all artificially altered information, irrespective of intent, degree, or risk, capturing ordinary image corrections, video compression, text summarization, and content generated in enterprise business solutions. This broad scope conflicts with risk-based regulatory principles, which require that compliance measures correspond to demonstrable harm rather than the mere use of a technology. Routine enhancements such as cropping, resizing, or denoising are fundamental to digital workflows — treating them as synthetic content would over-burden creators and intermediaries, diverting resources from genuinely harmful content. The definition also fails to distinguish between consumer-facing and business-

[6] Durable Content Credentials, 8 April 2024: https://contentauthenticity.org/blog/durable-content-credentials

to-business contexts, despite B2B applications posing limited risk to user safety and public order given their user base and intended and likely uses.

Recommendation: Narrow the definition of SGI to content intended or likely to mislead or cause harm, with express exclusions for benign or incidental modifications and content generated within enterprise business environments.

Conclusion

Thank you for allowing us to provide the enterprise software sector's perspective on this important issue. We look forward to engaging with MeitY on this policy issue.

Please do not hesitate to contact the undersigned at venkateshk@bsa.org in if you have any questions or comments.

Yours sincerely,

Venkatesh Krishnamoorthy
Country Manager, India
Business Software Alliance