Business Software Alliance

Friday, March 27, 2026

Shri. Ajit Kumar
Joint Secretary
Ministry of Electronics and Information Technology (**MeitY**)
Government of India,
New Delhi.

E-mail: js-akumar@meity.gov.in

Cc: Shri. S. Krishnan, Secretary

BSA PRIMER ON CONTENT AUTHENTICITY AND SYNTHETICALLY GENERATED INFORMATION

Dear Shri. Ajit Kumarji,

Greetings on behalf of the Business Software Alliance (**BSA**)[1].

MeitY notified the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026 (**SGI Amendment**)[2] on February 10, 2026.

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] To this end, we provided our inputs to MeitY in support of the development and deployment of reliable content provenance and authenticity mechanisms that can help users.[4]

Content provenance and authenticity are rapidly developing fields. A range of technologies now exist to identify how content is created or modified, including machine-readable watermarks, digital fingerprints, and secure metadata. These tools can work together as layers of verification and their capabilities continue to advance. As SGI-developing technologies are rapidly evolving as well, it is important that regulatory frameworks remain flexible to accommodate such developments and support the adoption of mature, internationally recognised technical methods. The global regulatory landscape for AI-generated content and content authenticity continues to

[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, Amadeus, Amazon Web Services, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cloudflare, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, 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] The Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Amendment Rules, 2026, February 10, 2026 at https://www.meity.gov.in/static/uploads/2026/02/f55fe52418b03f58b0669f6a8bc03b6d.pdf

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

[4] BSA's Comments on the draft amendments to the Information Technology (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021at https://www.bsa.org/files/policy-filings/11042025bsaindcmtsit.pdf

evolve. Jurisdictions around the world, such as the EU and Singapore, among others, continue to review regulatory approaches to address the challenges posed by SGI, deepfakes, and AI enabled misinformation.

Considering these developments, we suggest that MeitY conduct a review of the SGI amendments within 3 to 5 months. Drawing on BSA's learnings on content authenticity globally, described in the attached Primer,[5] we respectfully offer the following recommendations for MeitY's consideration in its review process of the SGI Amendment framework:

- **Establish Role-Appropriate Obligations:**
    - Content provenance and authentication obligations should focus on consumer-facing audio, visual, and audio-visual content. They should not apply to AI-generated text, which raises distinct issues, or in business-to-business settings, where companies are better positioned to create specific methods to address content authenticity based on their particular uses.
    - Obligations to apply content provenance information, like watermarks, digital fingerprints, or secure metadata, should fall on companies that develop AI systems. These obligations should not fall on the company that develops the underlying AI model, since an AI model is not used to generate content until it is integrated into a specific AI system. As a result, these measures are best applied at the AI system level, by the developer of the AI system.
- **Avoid mandating visible watermarks** that can be easily removed and create a false sense of security. Instead, policies should encourage the use of machine-readable, invisible watermarks alongside digital fingerprints and secure metadata as complementary verification layers.

BSA will be grateful for the opportunity to support MeitY in this important policy process. Our primer on content authenticity is enclosed. Please feel free 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

[5] BSA's Primer on Content Authenticity, 2026: https://www.bsa.org/policy-filings/primer-on-content-authenticity

Business Software Alliance

Primer on Content Authenticity

Transparency about AI-generated content is an important part of promoting responsible AI.

The Business Software Alliance supports the development and deployment of reliable content authentication and provenance mechanisms that can help users identify the history and origin of AI-generated content. This can help consumers know when content is human- or AI-generated and can help to address misinformation and disinformation.

Across the globe, policymakers are confronting important questions about how consumers can tell if the pictures and videos they see online are real. To do that, consumers need to know about tools that already exist today, like watermarks, digital fingerprints, and secure metadata. These tools can show who made a picture or video—and whether digital tools like AI have modified it. Policymakers should support the use of these and other tools to help consumers understand if content was made by humans or by AI.

BSA's primer on content authenticity:

- Defines content authenticity and content provenance, which are foundational to AI transparency policies.
- Describes existing tools that can identify who created an image or video and whether that image or video has been altered. These include machine-readable watermarks, digital fingerprints, and secure metadata.
- Describes the different actions that different types of companies can take to help support the use of content authenticity and provenance tools.
- Explores the policy issues raised by deepfakes.
- Explains how different AI policies can help consumers tell fact from fiction online.

What Are Content Provenance and Content Authentication?

Content Provenance

Content provenance information describes where content came from and how it changed. It tracks the origin, ownership, and modification history of a digital file, like an image, video, or audio clip. Example: A photograph includes cryptographically signed metadata that shows the date and location in which a photo was taken, the type of camera used to take the photo, and any software editing programs used to edit the file.

Content Authentication

Content authentication helps indicate whether the content and its provenance information are trustworthy and whether the metadata, watermarks, or credentials for a piece of content have been tampered with. Example: a content authentication tool can check a video's embedded signature to confirm it matches the issuer's public key and that the file hasn't been modified since publication.

What Tools Can Help Identify AI-Created Content?

A range of technologies can identify how content was created or modified, including content made through AI tools. Policymakers should not require use of a single technique but should ensure companies can use a range of technologies to mark content as AI-generated or not.

Machine-Readable Watermarks

Machine-readable watermarks are added directly into a file, by embedding a small amount of information that can be detected by a specialized tool. A watermark can be embedded in a file's pixels or data stream, so the origin of the file can be verified even if it is later stripped of metadata. Example: an invisible watermark inside an image can tie that image back to a creator or authenticity platform.

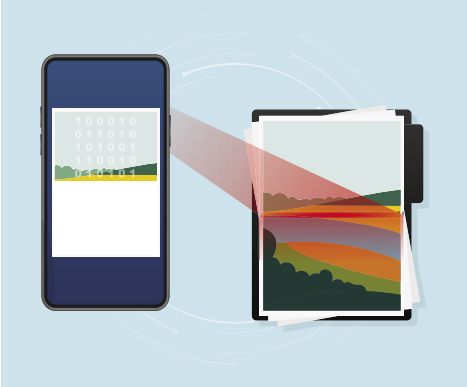

Digital Fingerprints

Digital fingerprints generate a unique identifier (or "fingerprint") from the content's core features, such as its color patterns, structure, or encoding. A digital fingerprint does not modify the file. Instead, it is a unique mathematical signature, such as a hash, that can be stored separately from the file. A file's digital fingerprint can be checked against a database of digital fingerprints, to confirm if a file matches an original or has been altered. Example: a video can be broken into frames to create a fingerprint from visual and audio cues; that signature can be checked against a database of digital signatures to confirm if the video was modified or altered.

Secure Metadata

Secure metadata can be used to store information about who created a file, what tools were used to create it, and how the file was edited. That information establishes the file's content provenance. Example: an image can include metadata that records the creator's identity, editing history, and software used; this can be digitally signed and embedded in the file so that viewers can verify its authenticity and origin.

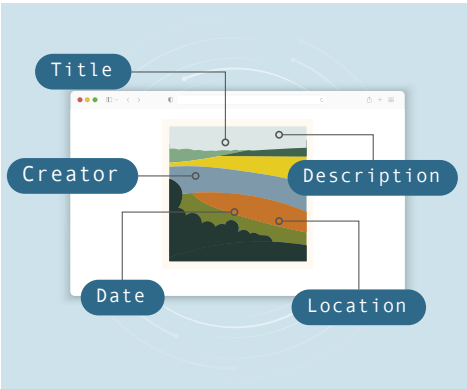

A Robust Approach: Combining Multiple Tools

These technologies can work together to support a robust approach to content authenticity, giving individuals the tools to know how content was created and whether it has been altered. All three processes act as layers of verification.

WATERMARKING
embeds authenticity

FINGERPRINTING
identifies content

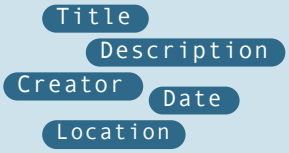

METADATA
can document context and authorship

While implementing all three processes may be ideal in many scenarios, it is not always necessary to use all three techniques in every instance. Rather, policies should encourage the use of technical methods that are mature enough for adoption today while remaining flexible enough to accommodate a range of technical and operational uses.

A Bad Fit for AI: Visible Watermarks

Visible watermarks are a low-tech method used for decades to label content, including to clearly identify documents as 'Draft' or 'Confidential.' In the AI context, requiring visible watermarks is not practical, since they can be easily removed. That undermines their effectiveness and creates a false sense of security.

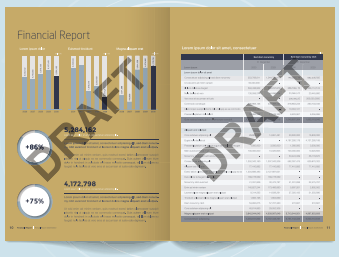

IN FOCUS

AI GENERATED TEXT

There are a range of tools to label images and audio-visual content, to help consumers know if that content is authentic or AI-generated. In contrast, any requirements to label AI-generated text raise a host of practical concerns. To ensure consumers know about text-based AI generated content, we recommend focusing on requirements that help individuals know when they interact with an AI system.

IN FOCUS

COALITION FOR CONTENT PROVENANCE AND AUTHENTICITY (C2PA) STANDARD

The C2PA standard can be used by anyone to incorporate digital provenance information into products and processes. It combines several technologies and is expected to be approved by the International Standards Organization as a global standard.

HOW DOES IT WORK?

1. **CREATING A FILE.** When a person creates a new file, she can use a tool that generates information about the file's provenance—including how it was created and any edits.
2. **CONTENT CREDENTIALS.** The provenance information is encoded into a content credential.
3. **CRYPTOGRAPHIC SIGNING.** The content credential is signed with the private key of the software or hardware used to create it, ensuring its authenticity. A public key is made available for authentication.
4. **EMBEDDING AND/OR WATERMARKING.** The content credential can be stored in several places, including embedding it within the file or storing it in a separate database where it is associated with an invisible watermark or digital signature.
5. **VERIFICATION.** To check the authenticity of a file, an individual can use a tool that checks content credentials to confirm if the file is authentic or has been modified.
6. **PRESENTATION.** Content that is verified as authentic can be marked with a clear indicator, like a badge or icon.

What Are Content Credentials?

The latest C2PA standard relies on content credentials, which combine multiple techniques.

- Content credentials contain metadata about a file, including the date and time it was created and the technology used to create it.
- Watermark identifiers and digital signatures can be added into the content credential, which already includes the metadata.
- The package of credentials is digitally signed and uniquely connected to that file.
- The content credentials are both embedded into the content and stored separately, such as in a database of content credentials.

The AI Supply Chain: Different Roles and Responsibilities

The AI supply chain is evolving and includes a variety of companies. Lawmakers must recognize these different types of companies as they create policies on content authenticity—since different types of companies will have access to different information and be in a position to take different actions to protect consumers.

For example, one company may develop an AI model, a second company may integrate that AI model into an application, and yet another company may use that application to create new content.

Model Developers

These companies create AI models, which can be used in a wide variety of different applications. For example, a company may build a foundation model that can be adapted for many different use cases. The same foundation model can be used to power search engines, chatbots, spam detection software, tools that summarize long text documents, and a variety of other applications. The model developer will have information about how that foundation model is developed—but generally lacks information about how other companies deploy the model for specific uses.

Integrators

These companies may integrate an AI model into a particular application, for use by other companies. Some integrators may simply connect the AI model to a specific service or application, while other integrators may fine tune or modify the AI model before including it in a service or application. These companies will generally have information about any changes they made to the AI model, but they usually do not have direct insight into the model's initial development or the particular circumstances in which other companies use the resulting AI application. For example, a company may act as an integrator when it develops an AI application that incorporates one or more AI models.

Deployers

Companies that use an AI tool for a specific purpose are often called deployers. These companies decide when and how to use a particular AI technology—so they will have insight into the facts of a particular use case. But deployers often obtain AI tools from other companies, so they typically lack direct insight into the initial training of the AI tool.

Any policies on content authenticity need to reflect these different roles.

If policies apply one-size-fits-all requirements to these very different types of companies, it will not effectively help consumers. For example, developers of AI applications that generate content are in the best position to apply content provenance information to content generated by that AI application. In contrast, the developer of a foundation model would lack the technical ability to apply watermarks or other digital markings to content that is generated through an AI application developed and used by other entities.

Ensuring laws remain fit over time. What constitutes state of the art in ensuring solutions for content provenance will evolve over time; policymakers should ensure that any legislative framework accommodates such developments. Embracing open standards like C2PA can help in this effort.

Combatting Deepfakes

Tools like invisible watermarks, digital fingerprints, and secure metadata can help to identify authentic content—so that individuals can know when the images, videos, and audio clips they interact with are authentic.

Together, this helps combat the problem of deepfakes. While no bad actor is likely to label their deepfakes as fake, it is important to give individuals tools to understand when content is authentic. For this to work, these tools need to be adopted across devices and platforms. People also need to be educated about these tools, so they know to watch for information about the provenance of an image or a video, while maintaining a healthy skepticism about information they encounter online.

Bad actors will continue to find novel ways to exploit technologies, including AI, for deceptive purposes. But tools like secure metadata and the C2PA standard can be crucial in helping good actors prove the authenticity of their content—helping all of us tell fact from fiction online.

As policymakers consider the most effective way to address deepfakes, they should also consider the different roles and functions of different types of companies. This is crucial since different types of companies will be positioned to address different issues, due to key service-level, technical, functional, and user-based distinctions.

Policymakers should also recognize the different risks involved in distinct types of services. For example, business-to-business software services pose limited risk to user safety and public order, given the size of their user base and the fact that they do not provide services directly to consumers, which may create fewer deepfake-related concerns.

How AI Policies Can Help Consumers

When policymakers want to ensure consumers can tell fact from fiction online, it is important to rely on existing tools that can identify if content was created by AI or not.

POLICY GOAL	SOLUTION
 **Tell consumers when they interact with AI systems.**	When companies provide AI systems that interact directly with consumers they should tell consumers they are interacting with an AI system, unless it is obvious.
 **Help consumers know if content is AI-generated.**	Focus obligations to use content provenance and authentication tools on consumer-facing audio, visual, or video content. These requirements should not apply to text or in business-to-business settings. This is important because using content authenticity mechanisms doesn't make sense for text—like when a consumer uses AI tools to edit the text of an email or translate a document. In business-to-business settings, companies can create separate methods to address content authenticity, based on their specific uses.
 **Support the use of leading global technologies to recognize AI-created content.**	Ensure requirements to identify AI-created content can be satisfied by leading global tools, including open standards like the Coalition for Content Provenance and Authenticity (C2PA). A range of countries are considering regulations on content provenance; adopting country-specific standards can limit consumers' ability to use widely accepted standards that are updated over time.
 **Facilitate transparency about AI-generated content.**	Obligations to apply content provenance information to AI-generated content are usually best placed on developers of AI applications—since AI applications are used to generate content. Companies can use tools like machine-readable watermarks, digital fingerprints, and secure metadata to identify the origin of that content.
 **Make sure content provenance information remains available to users.**	Prohibit stripping content provenance information such as machine-readable watermarks and secure metadata, absent security concerns.