AI Transparency and Disclosure Framework

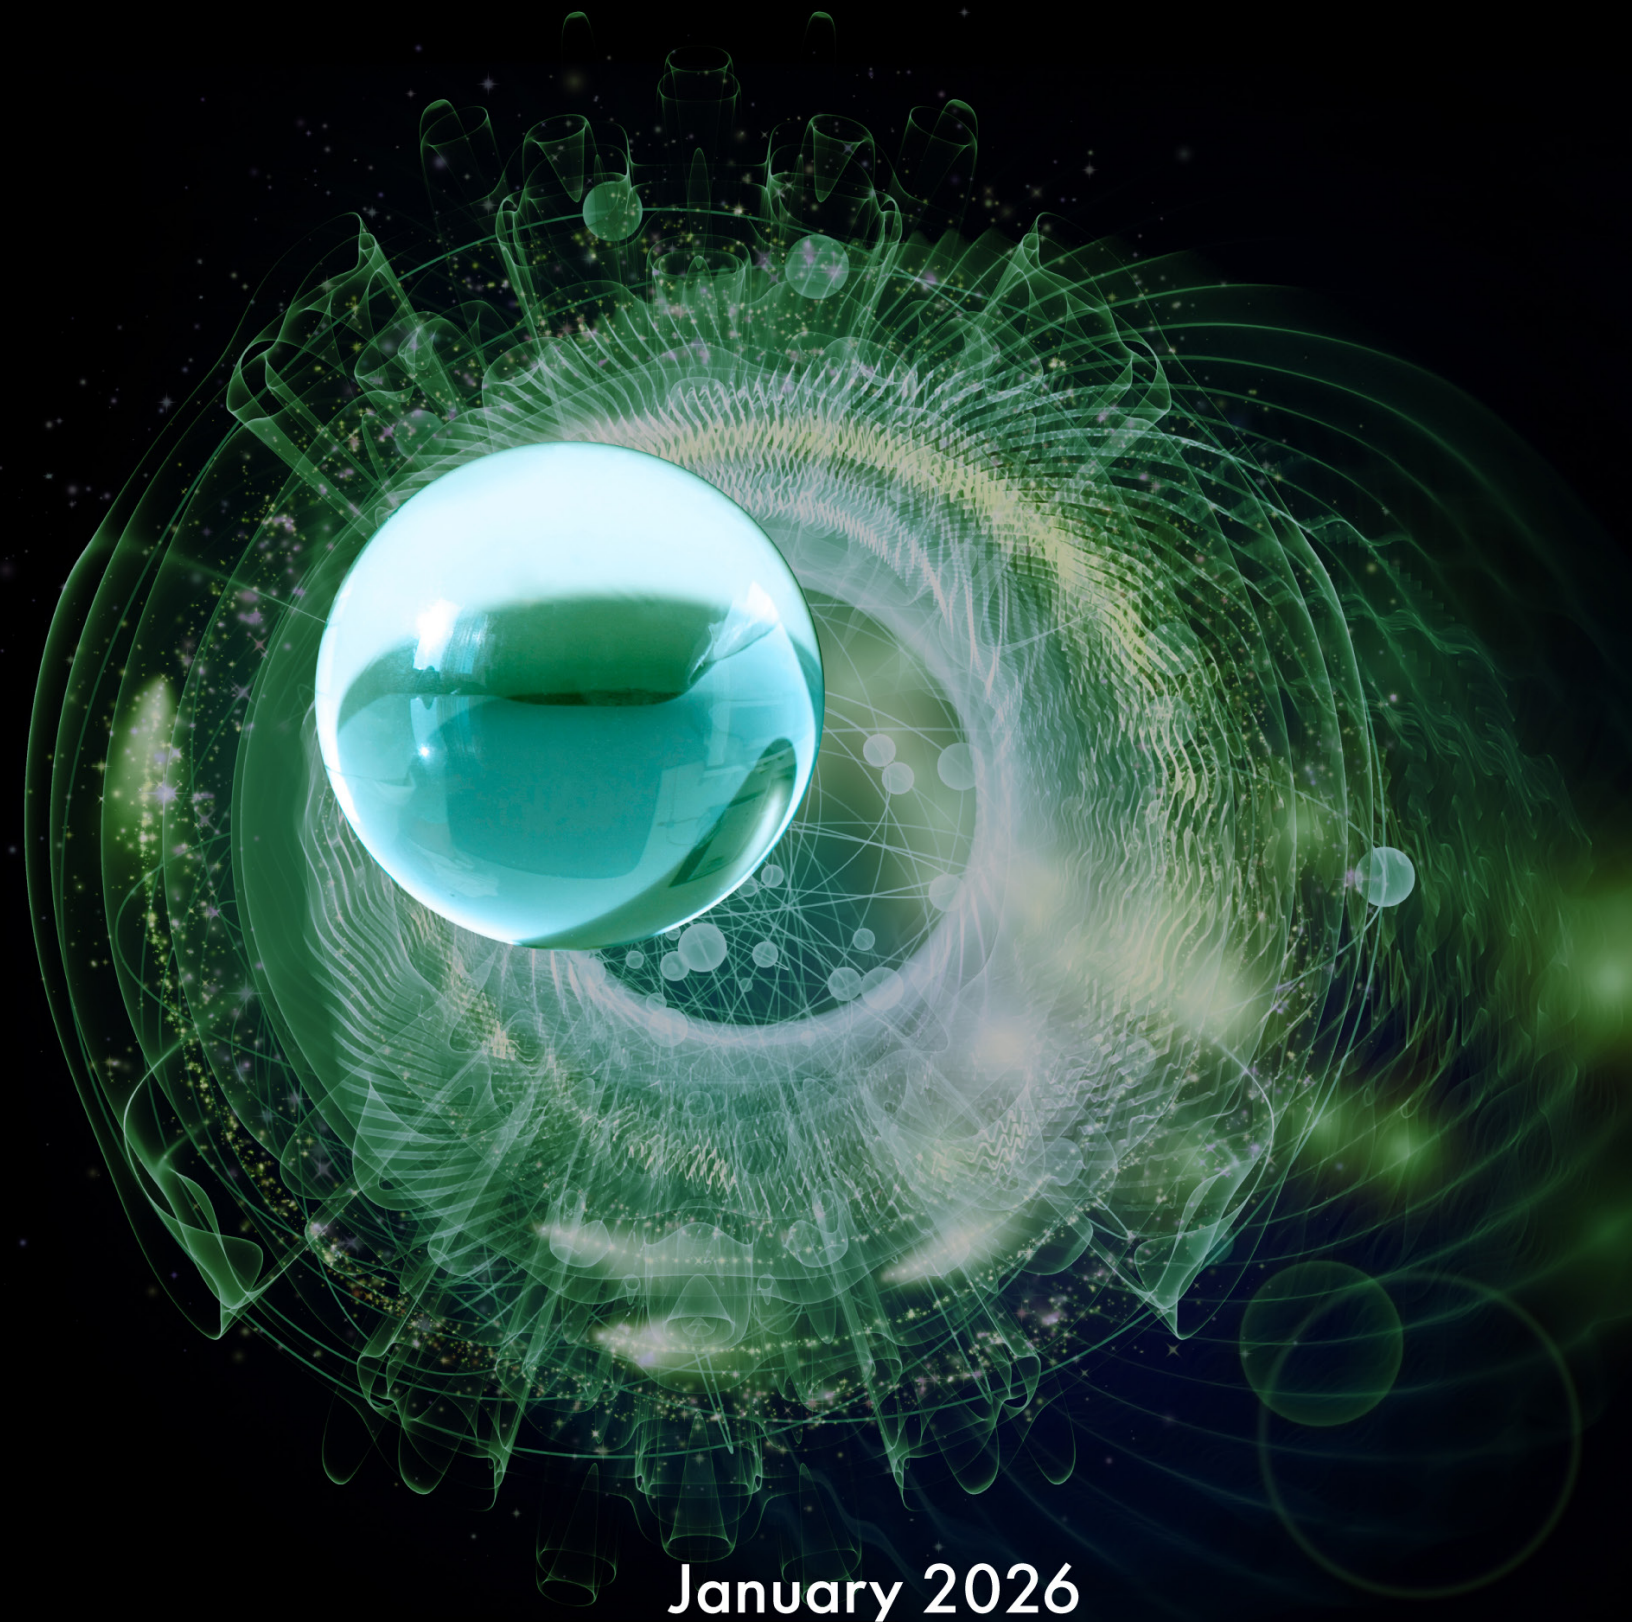

January 2026

Table of Contents

Maximize this Playbook

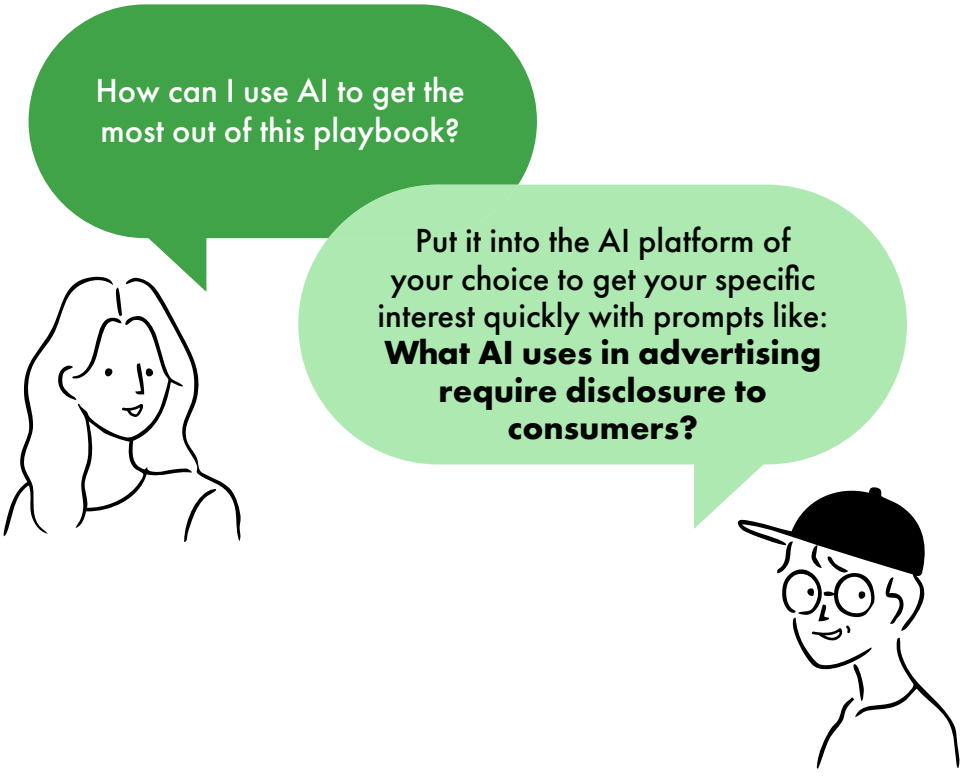

Executive Summary

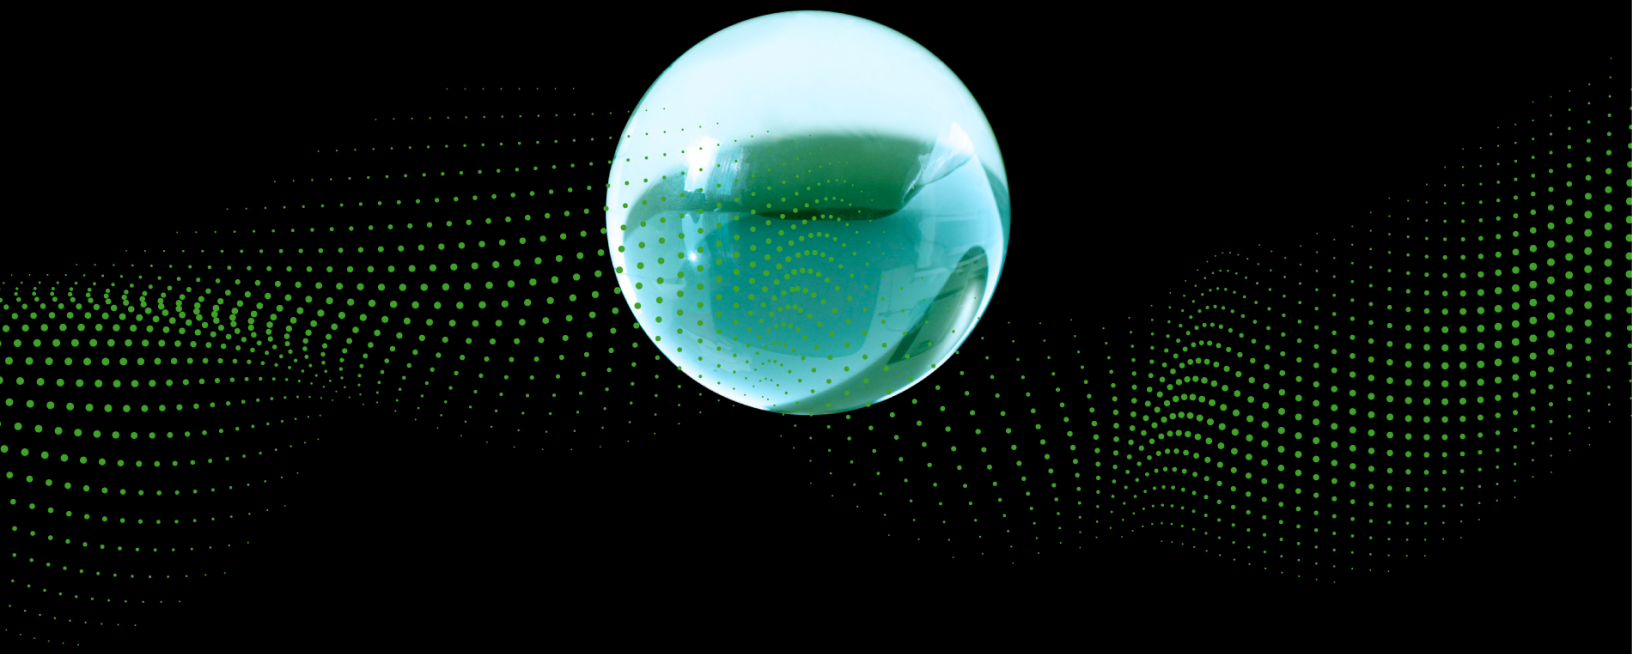

iab.

AI Transparency and Disclosure Framework

Executive Summary

Artificial intelligence has transformed advertising production at unprecedented speed, with AI-driven tools moving from experimental pilots to widespread adoption across creative workflows, enabling marketers to generate product imagery, synthetic voiceovers, video avatars, and large-scale personalized content. This framework applies specifically to advertising and marketing communications, including paid, sponsored, or brand-controlled content. It does not extend to organic social media, editorial journalism, or entertainment content unless those assets are used or repurposed as advertisements. This rapid integration creates both opportunity and risk; opportunity to reduce costs and expand creative output, and risk of consumer mistrust, operational inconsistency, and regulatory exposure.

This Framework looks to balance both that opportunity and risk by providing clear, actionable guidance to answer five essential questions:

1. Do we disclose AI use in advertising creative?
2. When?
3. How?
4. To whom?
5. Who is responsible?

Why Disclosure Matters Now

Inconsistent disclosure practices in GenAI advertising creative are fragmenting trust. Some campaigns label even minor AI involvement, while others disclose nothing unless legally mandated, creating confusion and "label fatigue" where audiences ignore disclosures entirely. Consumer expectations compound the challenge. Research shows broad majorities want transparency when AI is used, yet disclosure practices remain inconsistent across the industry. An IAB study found a significant perception gap, with 38% perceiving GenAI ads as creative and 20% as inauthentic. This disconnect highlights why transparent and consistent disclosure is crucial, not to discourage AI use, but to foster the trust necessary for responsible adoption. Meanwhile, regulatory pressure intensifies as the EU AI Act mandates baseline transparency requirements, U.S. states enact targeted legislation, and major platforms have taken early steps in implementing their own AI labeling approaches.

Key Findings and Recommendations

This Framework establishes a **risk-based, two-layer disclosure model** that balances transparency with operational feasibility.

IAB establishes these disclosure thresholds as **industry best practices** for responsible AI transparency. While adoption is voluntary, organizations that implement these standards position themselves as industry leaders, reduce regulatory risk, and build stronger consumer trust. This Framework does not carry enforcement authority; rather, it provides standardized guidance to help the industry self-regulate effectively. While transparency about AI use encourages consumer trust, disclosure does not exempt advertisers from compliance with laws prohibiting deceptive, unfair, or harmful advertising practices.

Consumer-Facing Disclosures are required only when AI materially shapes content in ways that could mislead a reasonable consumer about authenticity, identity, or representation. Disclosure triggers include synthetic humans, images, and videos, which are created entirely with AI and AI voice cloning of deceased individuals. This primarily applies to:

- Images or videos generated by prompt (text-to-image or image-to-image generation), where human input is limited to refinement, editing, or compositing
- AI-generated voices of deceased persons creating new statements they never actually made, even with estate authorization
- AI-generated voices of living persons making statements about specific events, actions, commitments, or circumstances that never occurred, as distinct from scripted commercial endorsements or brand messaging
- Photorealistic synthetic humans (fully AI-generated entities) in primary roles
- Digital twins of deceased individuals in any capacity, even with estate authorization
- Digital twins of living individuals depicted in specific events, scenarios, or locations that never occurred, as distinct from standard product endorsements or brand representation
- AI chatbots or conversational agents in ads that simulate human interaction in customer-facing contexts

Disclosure is not required for standard production techniques including, but not limited to, routine post-production, obviously stylized characters, or uses where AI serves as a tool, much as prior traditional tools would, without altering consumer perception of authenticity.

Machine-Readable Metadata should accompany AI-involved assets, using Coalition for Content Provenance and Authenticity (C2PA) content credentials, an emerging open standard for embedding provenance information. This includes both standard provenance fields and IAB-specific fields that record whether disclosure is required and whether a label was applied. This metadata enables verification, audit, and platform enforcement to determine whether consumer-facing disclosure is required. More detail on these fields and implementation is provided in the *Operational Framework*.

Accountability rests with advertisers (brand and/or agency), who bear ultimate responsibility for disclosure assessments and labelling. Platforms enforce compliance within their ecosystems but do not assume advertiser accountability.

Framework Overview

The Framework provides format-specific guidance across images, video, audio, text, and synthetic influencers, with clear thresholds for when disclosure is required. It notes regulatory requirements from the FTC, EU AI Act, and state-level legislation while positioning industry self-regulation as the bridge between legal gaps and consumer-first transparency.

Implementation follows a phased roadmap: establish governance foundations and assign an AI Disclosure Lead within 60 days, pilot disclosures and integrate compliance tooling within six months, and scale automated systems with ongoing monitoring within 12 months. Minimum viable practices, including pre-launch checklists, standardized labels, and C2PA metadata capture can be implemented immediately with minimal infrastructure.

Industry Impact and Urgency

By embedding disclosure into creative, operational, and technical workflows now, advertisers, agencies, publishers, and platforms can establish credible, scalable practices that sustain both trust and growth.

Introduction & Current Landscape

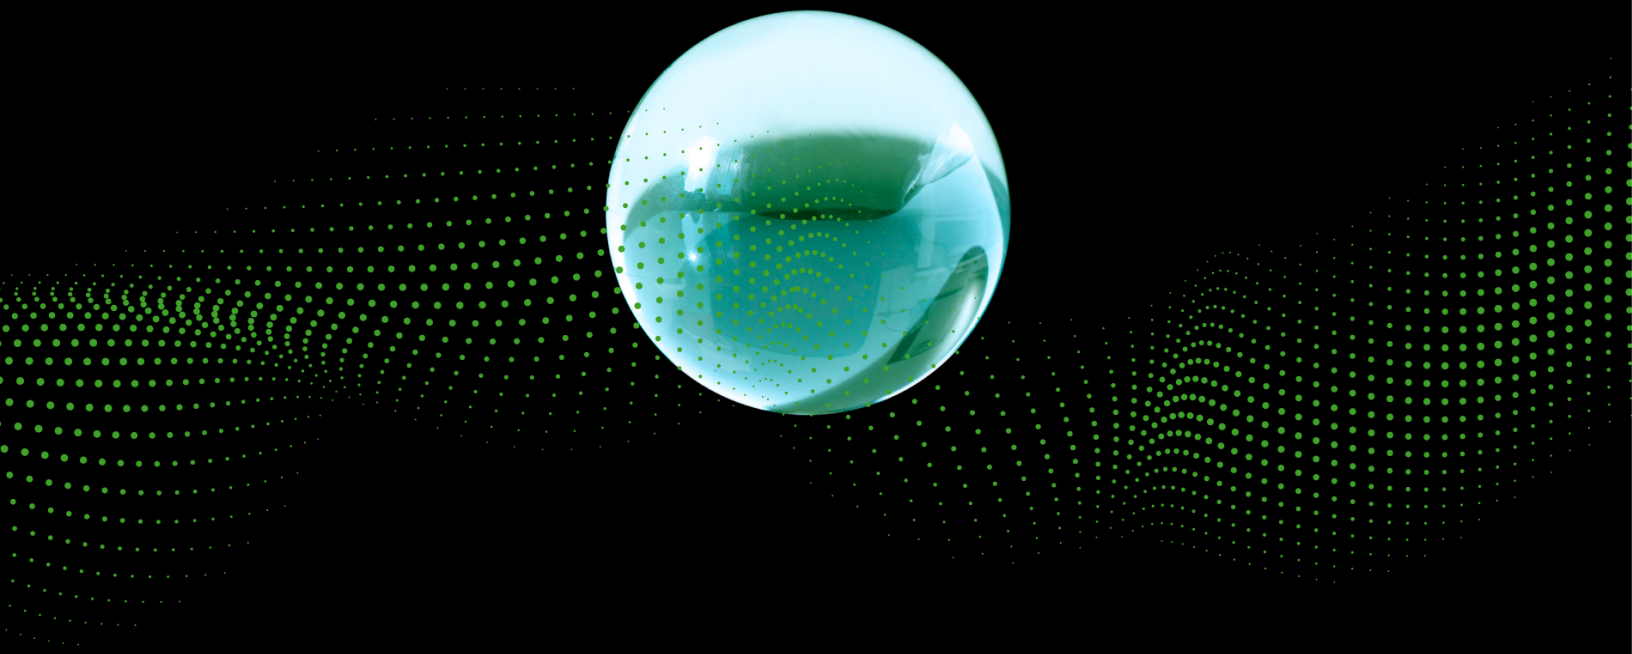

iab.

AI Transparency and Disclosure Framework

Introduction & Current Landscape

The rapid rise of generative AI has transformed advertising, publishing, and branded content creation. In less than three years, AI-driven tools have gone from experimental pilots to widespread adoption across creative, production, and media workflows. Marketers now use AI for product imagery, synthetic voiceovers, video avatars, and large-scale copy generation. Publishers and platforms are also adopting AI to accelerate content production, localization, and personalization. This pace of adoption creates both opportunity and risk: opportunity to reduce costs and expand creative output, and risk of consumer mistrust, inconsistent practices, and fragmented regulation.

This deliverable aims to answer five practical questions for the industry:

1. **Do we disclose AI use in advertising creative?**
2. **When?**
3. **How do we disclose?**
4. **To whom?**
5. **Who is responsible for disclosing?**

For the purposes of this framework, "advertising" follows the ICC Marketing Communications Code definition: any form of marketing communications carried by the media, typically in return for payment or other valuable consideration. Non-commercial editorial, entertainment, or user-generated content are out of scope unless they are distributed as advertising. This framework builds on insights from global self-regulatory bodies such as the International Chamber of Commerce (ICC), the European Advertising Standards Alliance (EASA), and the International Council for Ad Self-Regulation (ICAS), platform policies, and advertiser and publisher experiences to provide actionable guidance that is feasible, credible, and aligned with current regulation.

This Framework provides **voluntary industry best practices** for responsible AI disclosure in advertising, offering guidance to help organizations build consumer trust, reduce regulatory risk, and position themselves as industry leaders in responsible AI use.

The Rise of AI-Generated Content

Generative AI now touches nearly every format in advertising:

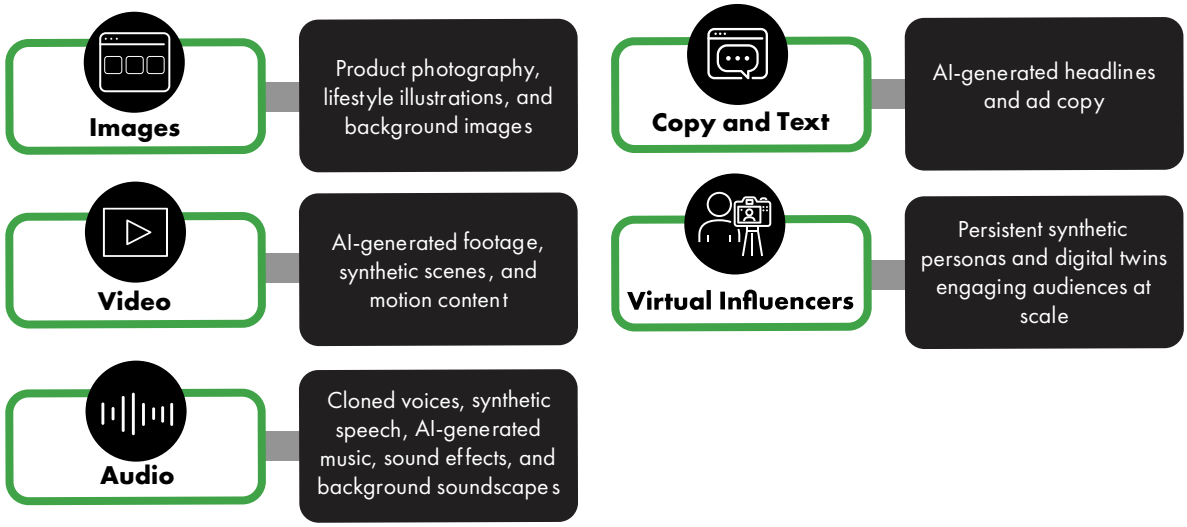

Industry adoption is exponentially on the rise because AI reduces the time and cost of production and enables personalization at a level not previously possible. However, this scale amplifies the risk that consumers may feel misled if they cannot distinguish between human and AI-generated content.

Industry Pain Points

Inconsistency. Brands, agencies, publishers, and platforms apply disclosure rules unevenly. Some campaigns label even minor AI involvement, while others label nothing unless legally mandated. This creates confusion and risks "label fatigue," where audiences ignore disclosures entirely.

Mixed consumer expectations. Research shows consumers broadly want to know when AI is used (Pew Research), with 76% of U.S. adults saying that they would like to know when GenAI is used in creation; however labels are interpreted unevenly. An IAB study of Gen Z and Millennial audiences found that younger consumers have some split perceptions about advertisements with GenAI elements. While 38% would regard AI-advertising brands as "creative," 20% would consider those brands "inauthentic."

Regulatory uncertainty. To date, few jurisdictions have implemented advertising-specific AI disclosure laws, though several have finalized legislation set to take effect in coming years. Existing consumer protection rules still apply, making advertisers liable for deceptive and unfair practices regardless of whether AI was involved. But emerging regulations (e.g., the EU AI Act, U.S. state laws, and South Korea's advertising labeling mandates) are moving toward greater transparency in AI advertising creative, though the scope and strictness of those requirements vary widely, from the EU's comprehensive, precautionary model to South Korea's blanket labeling approach to more decentralized, state-by-state approaches in the United States.

Operational burden. Overly broad disclosure requirements can increase costs and cause workflow disruptions without providing meaningful transparency. Industry groups, such as EASA and ICC, caution against blanket labeling and stress disclosure should depend on deception risk and consumer impact.

Stakeholder Impact

- **Advertisers** benefit from cost and speed AI efficiencies, but face heightened reputational risk if disclosures are absent or inconsistent.
- **Publishers** must balance trust in their content with advertiser demands, often acting as enforcers of labeling rules.
- **Agencies** are under pressure to integrate AI responsibly, flag usage to clients, and share accountability when creating content on behalf of brands.
- **Platforms** (e.g., Meta, YouTube, TikTok) have taken early steps to implement AI-labeling approaches, though practices vary and continue to evolve. For example, YouTube encourages disclosure of realistic synthetic media in user-generated content, while Meta applies "AI info" labels to certain ads based on advertiser declarations.
- **Consumers** want clarity but not overload. Poorly designed disclosures risk undermining trust rather than building it.

Research: Disclosure Effects on Trust and Credibility

While consumers generally want to know when AI is used, it comes with trade-offs. Labels can create unintended effects as audiences may assume unlabeled content is "real" by default (the "implied-truth effect"), or distrust creators who label responsibly. The implied-truth effect refers to a behavioral bias in which audiences perceive unlabeled or unflagged content as more credible or authentic simply because other content carries a disclosure label. This tension highlights the need for carefully designed and consistent disclosure standards.

Defining the Challenge

AI is not "just another tool." Unlike past technologies, AI outputs are probabilistic, opaque, and capable of fabricating convincing but false representations. The challenge is to develop disclosure standards that:

Protects **advertiser credibility** by preventing accusations of deception.

Builds **consumer trust** by signaling when AI has materially shaped content.

Ensures **regulatory alignment** while avoiding unnecessary burden.

This requires distinguishing between low-risk uses (AI as assistive background editing) that may not require disclosure, and high-risk uses (synthetic humans, altered claims, deepfake-style manipulations) where disclosure is essential.

Definition of Materiality

Within this framework, "materiality" refers to AI use that could mislead a reasonable consumer about what is authentic, factual, or human-created in an advertisement. The threshold is consumer impact, not the volume of AI tools used.

This approach aligns with longstanding advertising principles: advertisers remain responsible for truthful representation regardless of production methods, and disclosure is required when omission would deceive consumers about material aspects of the advertising content.

Core Principles of Transparency

To address these challenges, disclosure standards must be guided by four principles:

1. **Transparency:** clear, conspicuous signals that AI is involved
2. **Proportionality:** requirements scaled to risk and impact
3. **Consistency:** alignment across formats and geographies
4. **Clarity:** consumer-first best practices designed to inform audiences, not just meet legal minimalism

Takeaway

The current landscape is defined by high adoption, inconsistent disclosure, and mounting regulatory scrutiny. The industry's credibility depends on proactive, standardized practices that balance transparency with feasibility. The following sections move from landscape to action, starting with **Disclosure by Content Type**, to provide clarity on where disclosure is required and how it should be implemented in practice.

Disclosure by Content Type

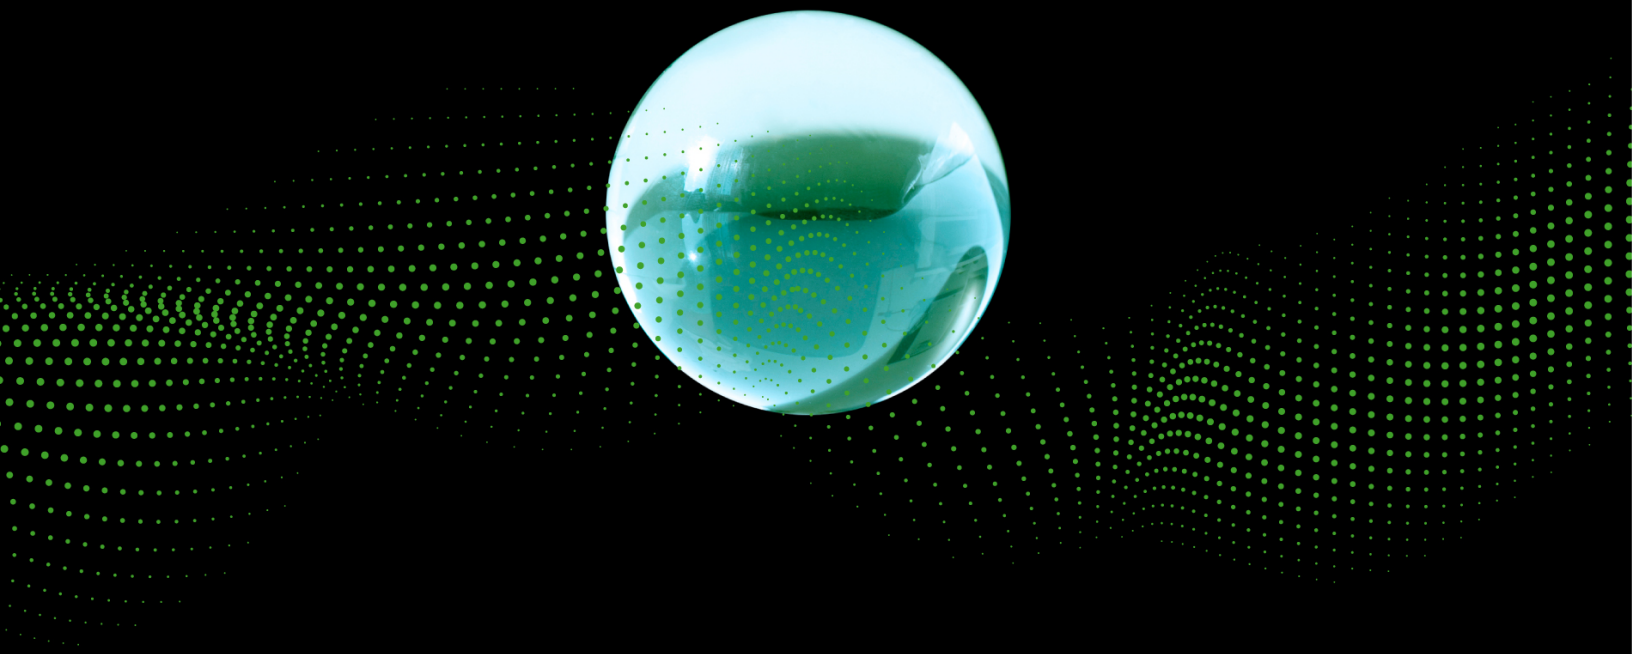

iab.

AI Transparency and Disclosure Framework

Disclosure by Content Type

This section builds directly on the core principles of transparency and consumer trust outlined earlier. The following guidance is designed to balance operational feasibility with consumer clarity. The threshold for disclosure centers on a fundamental question:

Does the AI involvement create a material risk that consumers will be misled about what they're seeing, hearing, or interacting with?

The answer depends not on the sophistication of the technology, but on whether its use could reasonably alter consumer perception or decision-making.

Establishing the Threshold

Before addressing specific content types, it's critical to understand what triggers disclosure requirements. The framework rests on three evaluative criteria:

Deception Potential:

Could a reasonable consumer be misled about the nature, origin, or authenticity of what they're viewing?

Material Impact:

Does AI involvement affect product representation, performance claims, or social proof in ways that influence consumer decisions or conduct, like purchase decisions?

Expectation Alignment:

Does the content fall outside the range of what consumers would reasonably expect from standard creative production?

These criteria deliberately focus on consumer impact rather than technical process. Consumers don't need to be aware of every AI tool in the production chain. They do, however, need to know when AI involvement could mislead them about what they're evaluating or being asked to buy.

Note on Platform Labeling:

These disclosure thresholds guide advertiser decision-making under this IAB framework. However, major social platforms may automatically apply "AI-generated" or similar labels based on their internal detection or upload rules, even when the use does not meet IAB's designated threshold. Advertisers should align internal policies with both this framework and applicable platform disclosure requirements.

AI-Generated Images

Included: Still photography, product shots, lifestyle imagery, and illustrations used in commercial contexts

Disclosure Required For:

Synthetic creation: Images generated from AI prompts (text-to-image or image-to-image generation), regardless of subsequent human refinement, editing, or compositing, excluding obviously non-realistic content.

Disclosure Not Required For:

General principle: Using AI to assist with standard editing tasks (retouching, color correction, compositing) does not trigger disclosure if the same edits would not require disclosure when performed through more traditional methods. The threshold is based on the nature and impact of the change regarding whether there is a material risk that the consumer would be misled, not whether AI tools were involved in the editing process.

- **Routine post-production** editing analogous to established editorial practices (e.g., color correction, lighting adjustments, background cleanup, dust/blemish removal, contrast optimization) unless this creates a material
- **Clearly stylized or fantastical imagery** where no reasonable consumer would perceive photographic realism (illustrated characters, clear CGI effects, artistic interpretations)
- **AI-assisted technical improvements** that don't alter content meaning: upscaling resolution, de-noising, format conversion, background alteration that doesn't affect product representation

Example Applications:

- An e-commerce retailer uses AI to generate images from prompts of their furniture products in various room settings without photographing either the product or the rooms: **disclosure required** (AI-generated from a prompt)
- A CPG retailer uses AI to remove a distracting shadow from a product photo taken in their studio: **no disclosure required**
- A supplement company uses AI tools to generate lifestyle backgrounds for product photos, placing bottles in kitchens, gyms, and outdoor settings: **no disclosure required** (real product photo with AI-generated backgrounds - covered under background alteration)

AI-Generated Images

Included: Synthetic footage, digitally generated scenes, and motion content

Disclosure Required For:

- **Synthetic video:** Video generated from AI prompts (text-to-video, image-to-video, or video-to-video generation), regardless of subsequent human refinement, editing, or compositing, excluding obviously stylized or fantastical content.

Disclosure Not Required For:

General principle: If a video editing task could be performed with traditional tools (Adobe Premiere, After Effects, VFX software) without requiring disclosure, using AI to perform the same task does not trigger disclosure, unless there is a material risk that the consumer would be misled.

- **Stylized or obviously fictional content:** Clearly animated characters, fantasy sequences, or content with evident non-realistic elements (cartoon avatars, mascots with exaggerated features)
- **Standard post-production:** Color grading, frame rate adjustments, background replacement that doesn't alter core content meaning, stabilization, routine visual effects

Example Applications:

- Financial exchange Kalshi creates an AI-generated ad from a text prompt: **disclosure required** (video created from text prompt)
- An automotive ad uses AI to seamlessly remove safety rigging from a stunt sequence: **no disclosure required** (standard VFX practice)
- A hotel brand uses AI to enhance sunset colors and remove tourists from footage of their property: **no disclosure required** (standard post-production editing)

AI-Generated Audio

Included: Synthetic speech, voice cloning, AI-generated music, sound design, and audio effects used in commercial contexts

Disclosure Required For:

- **Synthetic voice of deceased persons:** AI-generated voice content that creates new speech or statements from deceased individuals (words they never actually spoke while alive), even with estate authorization. This excludes the use of archived, actual recordings where the person genuinely spoke those words during their lifetime.
- **Synthetic voice of living persons making statements they never made:** AI-generated voice content of living individuals (with authorization) making statements that convey or imply real-world facts, experiences, personal actions, or documented events that the individual did not actually do, experience, or state, as distinct from scripted commercial endorsements or brand messaging that are clearly promotional in nature and explicitly authorized by the individual.

Disclosure Not Required For:

- **General principle:** If audio production or editing could be performed with traditional tools (Pro Tools, Logic Pro, audio mixing software) without requiring disclosure, using AI to perform the same task does not trigger disclosure. The threshold is whether the AI creates a false impression about who is speaking or what was said, not merely how the audio was produced.
- **Authorized synthetic voice of a living person for scripted commercial content:** When a living individual has explicitly authorized the use of AI to replicate their voice for standard advertising purposes (product endorsements, brand messaging, scripted narration)
- **Generic synthetic voices:** AI-generated voiceovers that don't impersonate specific individuals and are used for standard narration, where the identity of the speaker is immaterial to the message
- **Background music and soundscapes:** AI-generated music, ambient sound, or sound effects where no reasonable consumer would interpret the content as being created by or associated with a specific artist or composer
- **Standard audio enhancement:** Noise reduction, equalization, mixing, volume normalization, or other technical improvements that don't alter the substance of what was recorded
- **Minor voice corrections:** Breath removal, timing adjustments, pitch correction, or de-clicking that fall within standard audio engineering practices

Example Applications:

- A watch brand creates an ad using an AI-generated voice that sounds like a deceased celebrity endorsing their product (with estate approval): **disclosure required** (voice clone of deceased identifiable person)
- A meditation app generates calming background music using AI composition tools: **no disclosure required** (background music where production method is immaterial)
- A podcast ad uses AI to clean up background noise and normalize volume levels: **no disclosure required** (standard audio post-production)
- A software company uses a generic AI voiceover to explain their product's technical features in a demo video: **no disclosure required** (standard narration where speaker identity is immaterial)
- An automotive company uses an authorized AI voice clone of a racing champion saying "I set the track record at Daytona using this exact model" when they never drove that car or set that record: **disclosure required** (synthetic voice depicting fabricated statements about real-world events)

Synthetic Influencers and Virtual Personalities

Included: Digital twins (AI replica of real individuals), synthetic influencers (entirely fictional AI-generated personas (e.g., Lil Miquela)

Disclosure Required For:

- **Digital twins of real individuals (deceased):** AI replicas of real people who are deceased, used in advertising in any capacity, even with estate authorization
- **Digital twins depicting fabricated events or scenarios:** AI replicas of living individuals (even with authorization) shown participating in specific events, locations, interactions, or circumstances that never occurred, where consumers might believe they are viewing authentic documentation
- **AI chatbots or conversational agents in ads:** When AI-powered personas engage directly with consumers in ways that simulate human interaction

Disclosure Not Required For:

General principle: If a character's synthetic nature is immediately apparent through visual style, context, or presentation, additional disclosure may be redundant. The threshold is whether a reasonable consumer might believe they're interacting with or viewing a real human being.

- **Digital twins in authorized standard advertising contexts:** AI replicas of living individuals with authorization, used in typical commercial advertising formats (product endorsements, brand representation, scripted messages) where no specific real-world event or scenario is being fabricated
- **Obviously non-human or stylized characters:** Cartoon mascots, animated brand characters, fantasy creatures, or virtual personas with clearly non-realistic visual styles (anime-style, cartoon-like, fantastical features) where no reasonable consumer would perceive them as real people
- **Incidental background figures:** AI-generated personas, such as synthetic influencers or generic AI-generated individuals in non-focal roles such as crowds, background scenes, or ambient elements where consumer attention is directed elsewhere and individual authenticity is not material to the ad's message

Example Applications:

- A fashion brand creates "Mia," a photorealistic AI influencer with an Instagram presence, who posts regularly and responds to comments: **disclosure required** (synthetic persona that could be mistaken for a real person)
- A beverage company showcases a digital twin of a celebrity as a spokesperson with celebrity authorization: no **disclosure required** (authorized digital twin of a real individual with brand endorsement)
- An insurance company uses an animated gecko mascot in their commercials: **no disclosure required** (obviously non-human character)
- A game developer features stylized anime-style virtual characters in promotional materials: **no disclosure required** (clearly fictional, non-realistic visual style)
- A political campaign creates an ad showing a digital twin of the candidate visiting a disaster site and comforting residents, using footage and scenarios that never occurred: **disclosure required** (digital twin depicting fabricated events, even with authorization)

Text and Copy

Included: Headlines, ad copy, product descriptions, email content, and any written material used in commercial contexts

Disclosure Required For:

For text and copy, disclosure is generally not required. Advertisers remain responsible for the accuracy and substantiation of all claims regardless of authorship method. AI-generated false claims, unverified data, or expert advice without proper human review are prohibited under existing advertising and consumer protection laws, and disclosure does not cure these violations.

Disclosure Not Required For:

- **Standard marketing copy:** Headlines, slogans, product descriptions, email subject lines, social media posts, and other promotional text where AI serves as a drafting tool
- **General creative writing:** Blog posts, brand storytelling, and editorial content where AI assists but human editors validate accuracy and tone
- **Translations and localizations:** AI-powered translation or adaptation of approved content into multiple languages

Example Applications:

- A clothing brand uses AI to draft hundreds of product descriptions that are reviewed and approved by merchandising staff: **no disclosure required**
- A travel company uses AI to translate its English website content into 15 languages: **no disclosure required**

Section Takeaway:

Disclosure requirements vary by content type, but the underlying logic remains consistent: transparency is required when AI involvement creates a material risk of consumer deception which falls outside current consumer protection laws. The framework intentionally avoids mandating disclosure for every AI tool in the production pipeline, focusing instead on consumer-facing impact. The following section addresses how to implement these disclosures effectively.

Operational Framework

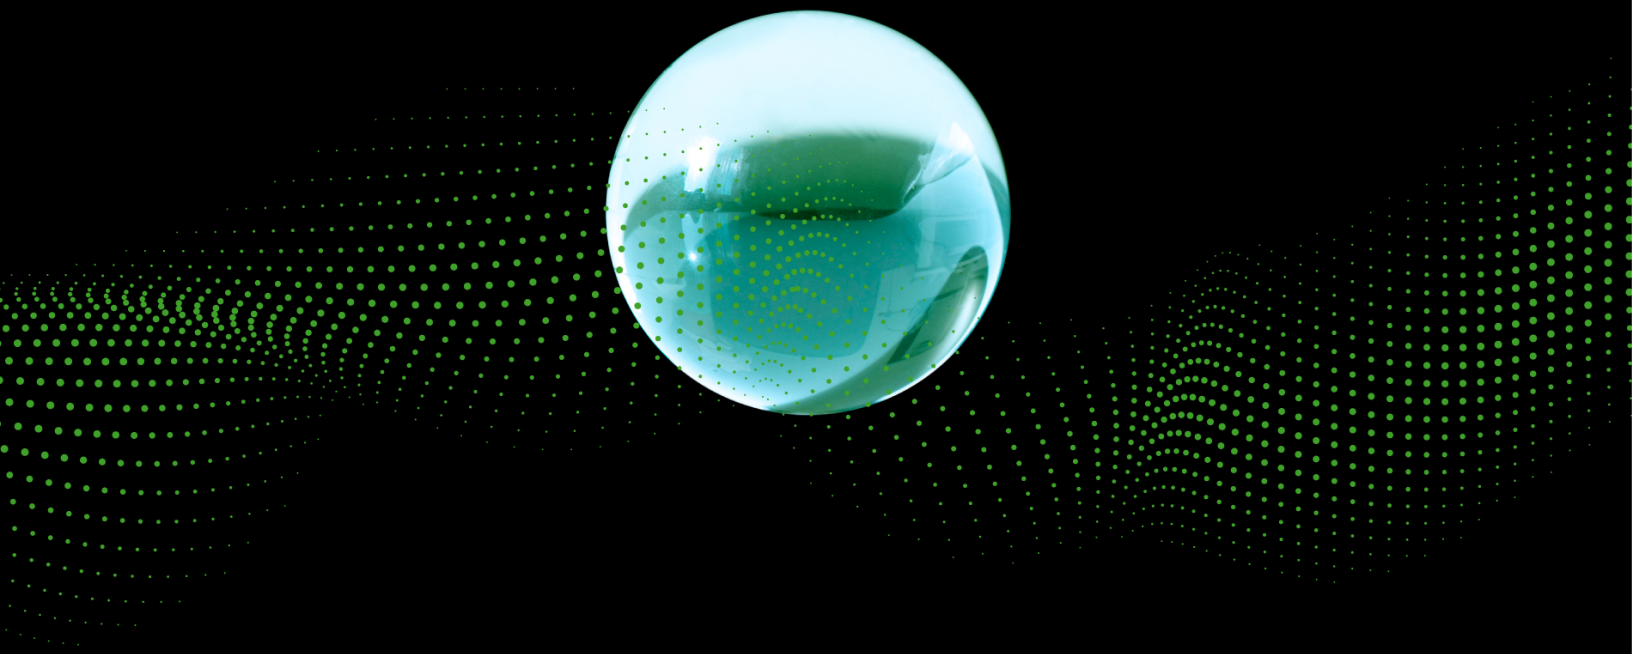

iab.

AI Transparency and Disclosure Framework

Operational Framework

The preceding section established *what* advertising content requires disclosure based on materiality thresholds. This section addresses how to operationalize those decisions: disclosure mechanics, audiences, timing, and accountability.

The Two-Layer Model

Effective AI disclosure operates across two interconnected layers:

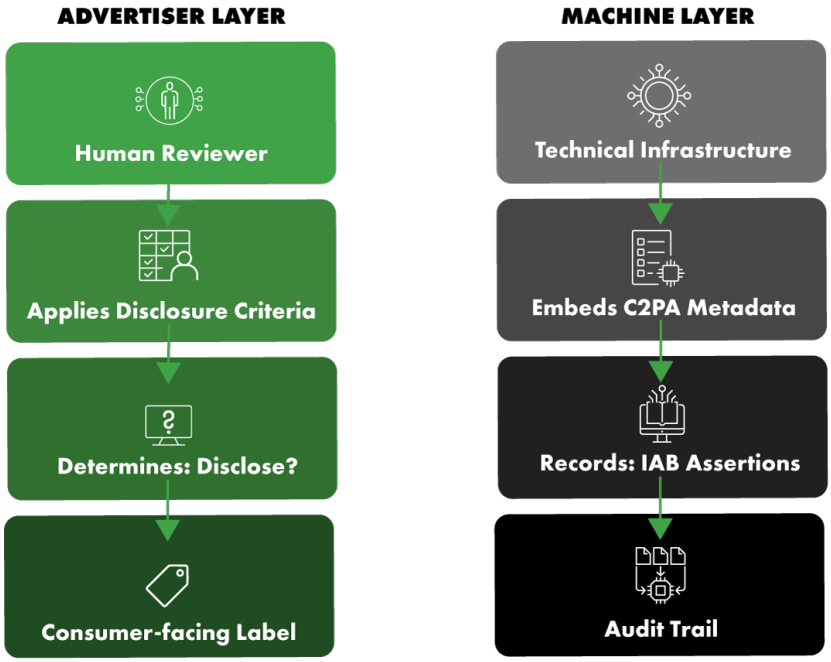

Machine Layer (Provenance Infrastructure): Using the Coalition for Content Provenance and Authenticity (C2PA) standard, an open technical standard that records how digital content was created, edited, and published by embedding metadata cryptographically into the file itself, this layer provides an auditable record of content creation and modification, capturing what AI tools were used, when changes occurred, and the advertiser's threshold determination. While C2PA adoption is still maturing and cannot guarantee completeness or prevent all forms of tampering, it enables platforms, regulators, and auditors a method to verify declared disclosure decisions, providing transparency infrastructure that will strengthen as ecosystem adoption scales.

Advertiser Layer (Disclosure and Accountability): Human reviewers apply IAB disclosure thresholds to determine whether AI involvement poses a risk of consumer deception. When materiality triggers disclosure requirements, advertisers implement consumer-facing labels or explanations.

The two-layer model enables both verification (through metadata) and transparency (through consumer-facing labels). Simply labeling content as 'AI-generated' provides limited value; disclosure is meaningful only when AI involvement affects consumer perception of authenticity, origin, or representation.

1. How to Disclose (Signal)

Disclosure operates through two parallel mechanisms: visible consumer-facing signals and embedded provenance metadata.

Consumer-Facing Signals

Visual Methods: Clear, persistent text indicators are the primary disclosure method, providing the most reliable

disclosure for consumer comprehension and accessibility. Effective text labels appear near content, use plain language ("AI-generated spokesperson" not "synthetically produced asset"), remain visible throughout exposure, and provide sufficient color contrast to ensure readability.

Alternative disclosure methods may be used when text labels would materially compromise advertising effectiveness:

- **Visual indicators** (watermarks, badges, or standardized icons) when creative space is constrained
- **Interactive elements** (tap/hover information icons) in digital environments
- **Adjacent placement** (disclosure placed next to rather than on the creative asset)

Alternative methods must meet the same visibility, clarity, and accessibility standards as text labels. Organizations using visual indicators should align with emerging platform conventions (i.e., Meta, TikTok, YouTube, etc.) and C2PA provenance standards to support cross-platform consistency.

Audio Methods: Voice-only contexts require verbal disclosures preceding or following AI-generated segments, distinct audio cues, or companion visual disclosures in available interfaces.

Accessibility: Methods must include alt text for screen readers, audio descriptions, sufficient contrast ratios, and use both visual and audio disclosure methods when possible. For audio-only content where no visual interface exists (e.g., radio, podcasts, smart speakers), disclosures must be clearly spoken in the same language as the advertisement and repeated at least once if the ad exceeds 60 seconds. The disclosure should be best positioned so that it is heard before or immediately after the AI-generated segment.

The Guiding Principle: The disclosure must be perceptible and accurate, clear enough to inform, but unobtrusive enough to avoid 'label fatigue.'

Required Disclosure by Content Type

Content Type	Threshold Trigger	Disclosure Method	Placement/Duration
Synthetic Image	Images generated from AI prompts (text-to-image or image-to-image), excluding obviously non-realistic content	Primary text Label: "AI-generated image" OR Alternative: standardized visual indicator	Placement near image when feasible
Synthetic Video	Video generated from AI prompts (text-to-video or image-to-video), excluding obviously stylized or fantastical content	Primary Text Label: "AI-generated video" OR Alternative: standardized visual indicator	Must appear on first frame/ introduction and remain visible throughout video
Synthetic Voice of Deceased Person	AI-generated voice of deceased person, even with estate authorization	Verbal disclosure: "AI-generated voice"	Verbal disclosure: "AI-generated voice"

Content Type	Threshold Trigger	Disclosure Method	Placement/Duration
 **Synthetic Voice (Living) - Fabricated Statements**	AI-generated voice of living person (even with authorization) making statements about specific events, actions, commitments, or circumstances that never occurred, as distinct from scripted commercial endorsements or brand messaging	Verbal disclosure: "AI-generated voice"	Audio: verbal disclosure pre- or post-content. Video: text overlay
 **Synthetic Avatars**	Fully AI-generated avatar which could be mistaken for a real human	Primary Text Label: "AI-generated person" OR Alternative: standardized visual indicator	Must appear on first frame/ introduction and remain visible throughout synthetic avatar use
 **Digital Twin (Deceased)**	AI re-creation of real person's likeness (deceased) with estate authorization	Primary Text Label: "AI-generated likeness" OR Alternative: standardized visual indicator	Must appear on first frame/ introduction and remain visible throughout use
 **Digital Twin (Living) - Fabricated Events**	AI replica of living person (even with authorization) depicted in specific events, scenarios, or locations that never occurred, as distinct from standard product endorsements or brand representation	Primary Text Label: "AI-generated likeness" OR Alternative: standardized visual indicator	Must appear on first frame/ introduction and remain visible throughout use
 **AI Chatbots/ Conversational Agents in Ads**	Interactive AI-powered personas embedded in advertising that engage directly with consumers in ways that simulate human interaction	Text disclosure identifying AI nature of interaction	At initiation of interaction and readily accessible throughout ad experience

Disclosure Method Standards

Text Labels (Primary Method)

Text labels are the recommended standard for consumer clarity and accessibility. Visual indicators (watermarks, badges, icons, or symbols) may be used as an alternative when text labels would materially compromise advertising effectiveness.

Visual Indicators (Alternative Method)

Recommended Visual Indicators:

1. **Sparkle/Star Icon** (✨): A sparkle or star icon provides a subtle, recognizable indicator that aligns with emerging platform conventions. Several major platforms use variations of sparkle iconography to denote AI involvement. A sparkle or star icon reflects an emerging cross-platform design convention for denoting AI-assisted or synthetic content.
2. **C2PA Content Credentials Icon (CR):** The Content Credentials icon indicates that provenance metadata is available. When paired with IAB-compliant C2PA custom assertions (com.iab.threshold), the CR icon provides both consumer-facing disclosure and technical verification. Users can tap/click the icon to view full disclosure details.
3. **Platform-Applied Indicators:** When distributing on platforms that automatically apply AI disclosure indicators based on C2PA metadata (such as Meta's "AI info" label), advertisers may rely on platform-applied labels provided they meet IAB conspicuousness standards. The advertiser remains accountable for ensuring the platform label adequately fulfills disclosure requirements.

Implementation Requirements

Visual (Display, Video, Social):

For text labels:

- **Language:** Plain language using "AI-generated" terminology
- **Text size:** Sufficiently large to be clearly readable
- **Contrast:** Minimum Web Content Accessibility Guidelines (WCAG) AA contrast ratio (4.5:1 for standard text)

For visual indicators (watermarks, badges, icons):

- **Size:** Sufficiently large to be clearly readable
- **Contrast:** Sufficient contrast with background to ensure visibility across contexts

All visual disclosure methods must:

- **Persistence:** Remain visible throughout content exposure or appear on first frame/screen
- **Placement:** Positioned to avoid obstruction by platform UI elements

Audio (Radio, Podcast, Streaming):

- **Clarity:** Spoken at normal pace in clear, intelligible voice

Accessibility Requirements:

- All visual labels and indicators must include alt text for screen readers
- Audio disclosures should have visual equivalents when video is present

- Captions must include disclosure text

Provenance Metadata (C2PA)

What it is: Coalition for Content Provenance and Authenticity (C2PA) is an emerging industry-standard technical format that embeds invisible metadata into ad assets, like digital nutrition labels that travel with the content. This metadata creates a verifiable record of how the asset was created and whether AI was involved.

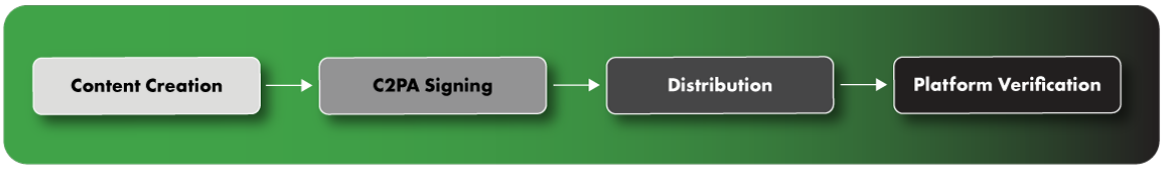

Why it matters: While consumer-facing labels tell audiences "this used AI," C2PA metadata enables platforms to verify those claims, supports audit trails for regulatory compliance, and can trigger automated disclosure systems at scale.

IAB Implementation:

Before distribution, the advertiser's disclosure determination must be encoded in C2PA metadata using two IAB custom assertions:

1. com.iab.threshold — whether AI involvement met the threshold to require disclosure (met/not met)
2. com.iab.disclosure — whether a consumer-facing disclosure was applied (yes/no)

Metadata should also capture AI involvement type (e.g., "synthetic voiceover," "background replaced") and tool identification. While certain fields, including timestamps, tool versions, signer entity, and some AI involvement types, may be captured automatically, the disclosure classification remains a human-determined judgment.

Important: This metadata serves verification and audit functions, not public consumer disclosure.

Platform-Specific Implementation

Advertisers should embed disclosure directly in creative assets to ensure labels travel across all distribution channels. Platform-applied labels may supplement but do not replace advertiser disclosure obligations. When platforms provide their own labeling mechanisms (based on C2PA metadata or advertiser declarations), these function as additional transparency layers rather than substitutes for IAB-required disclosure.

To avoid redundant labeling, advertisers may rely on platform-applied disclosure when the platform automatically generates labels from C2PA metadata that meet IAB standards for clarity and conspicuousness. In such cases, the advertiser remains accountable for ensuring the platform label adequately fulfills disclosure requirements.

2. To Whom to Disclose (Audience)

AI transparency in advertising serves multiple audiences: consumers, business partners, and regulators.

Consumer-facing disclosures are required only when AI materially shapes content in a way that could mislead a reasonable viewer about authenticity, identity, or representation. These cues appear directly within the ad experience and follow IAB format-specific disclosure guidance.

B2B disclosures are embedded in metadata and accompany every AI-involved asset, whether consumer AI labeling is required or not. This ensures agencies, publishers, platforms, and auditors can verify usage and alignment with policy without relying on visible labels.

Regulators and oversight bodies may request access to provenance data for compliance or investigation purposes. The C2PA framework provides an audit trail that remains tamper-evident and machine-readable.

3. When to Disclose (Timing/Lifecycle)

Pre-Campaign (Planning & Creation)

Before distribution, human reviewers evaluate AI involvement against IAB materiality thresholds. The decision must be encoded in C2PA metadata, creating an auditable record that travels with content throughout its lifecycle.

During Campaign (Distribution)

Consumer-facing disclosures must appear whenever and wherever content is displayed, traveling across platforms, formats, and distribution channels. A synthetic influencer post requires disclosure on Instagram, TikTok, YouTube, and any platform where it appears, not relegated to one-time notices. Disclosure must be present at the point of consumer exposure.

Ongoing Relationship Disclosure

Content representing ongoing relationships, such as AI avatars serving as brand spokespersons or synthetic influencers maintaining social presences, requires disclosure in every appearance. Each consumer encounter with the AI entity requires its own disclosure.

Post-Publication (Archiving & Audit)

Localization and Personalization: When AI generates variations of human-created content, materiality must be assessed for each campaign approach. AI translation or minor localization typically falls below disclosure thresholds.

Dynamic Content:

- **Dynamic Creative Optimization (DCO):** Dynamic content AI-assembled from pre-approved, human-reviewed components, does not require disclosure.
- **AI Personalization:** Real-time AI generation of personalized creative variants requires C2PA metadata documentation for all AI-involved content. Consumer-facing disclosure is required when generated content meets the disclosure thresholds established in this Framework. For high-volume personalization campaigns in which manual labeling is operationally impractical; C2PA metadata enables platform verification and compliance monitoring while automated labeling systems mature.

When to Disclose (Timing / Lifecycle)

Disclosure occurs throughout the ad lifecycle to maintain consistency from creation to distribution:

1. **Pre-Campaign (Planning & Creation):**
   - Human reviewers apply IAB's threshold determination to note if disclosure is required.
   - Agencies document AI tools used, purpose of use, and initial threshold trigger assessment.
   - C2PA metadata is appended automatically at export from creative tools, supplemented by the advertiser's IAB custom assertions through manual human addition.
2. **During Campaign (Distribution):**
   - If disclosure is required, the label or cue is displayed wherever the ad appears.
3. **Post-Publication (Archiving & Audit):**
   - Provenance and disclosure records persist with the asset to enable compliance review, reuse, or reporting.
   - Updates, personalized variants, or localized versions generated with AI require C2PA metadata documentation. Consumer-facing disclosure is required when versions meet disclosure thresholds in this framework unless operationally infeasible.

Automation supports this process as metadata capture (tool, signer, timestamp) can occur automatically, *but whether disclosure is required remains a matter of human judgment.*

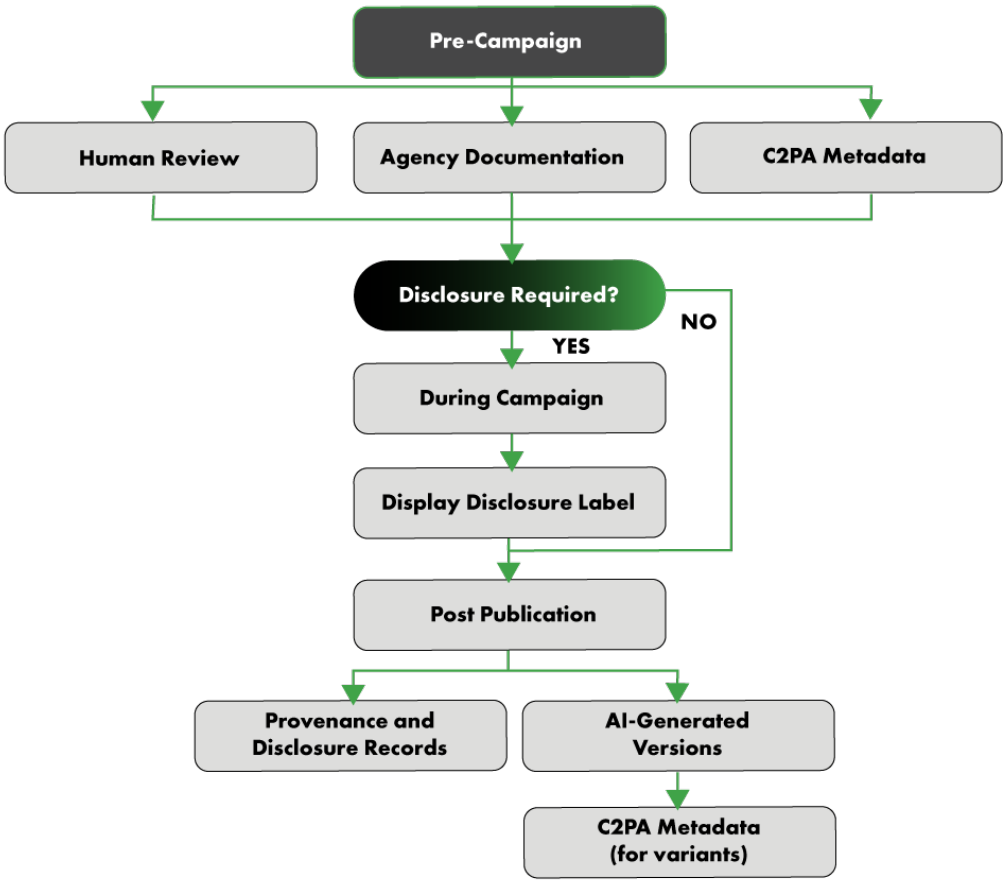

4. Who Is Responsible for Disclosure (Accountability)

Advertiser Accountability (Primary)

Ultimate responsibility rests with the advertiser, defined as the organization that creates the advertising content. This entity bears liability for determining disclosure thresholds following IAB guidance, implementing consumer-facing disclosures, and ensuring accurate provenance metadata. This accountability cannot be delegated.

In cases where agencies create and control advertising content on behalf of clients, both the agency and the client may share advertiser responsibility.

Platform Enforcement (Secondary)

Advertisers should embed IAB-compliant disclosure directly within their creative assets to ensure consistency across all distribution channels. Platform-applied labels based on their own policies may supplement, but should not replace advertiser disclosure obligations under this Framework.

As C2PA adoption matures, platforms can use custom IAB assertions (com.iab.threshold, com.iab.disclosure) to verify disclosure consistency and may flag or reject ads where metadata indicates disclosure is required but no consumer-facing label is present.

Platforms serve as secondary enforcers within their ecosystems, but do not assume advertiser liability for disclosure accuracy.

Regulatory & Standards Alignment

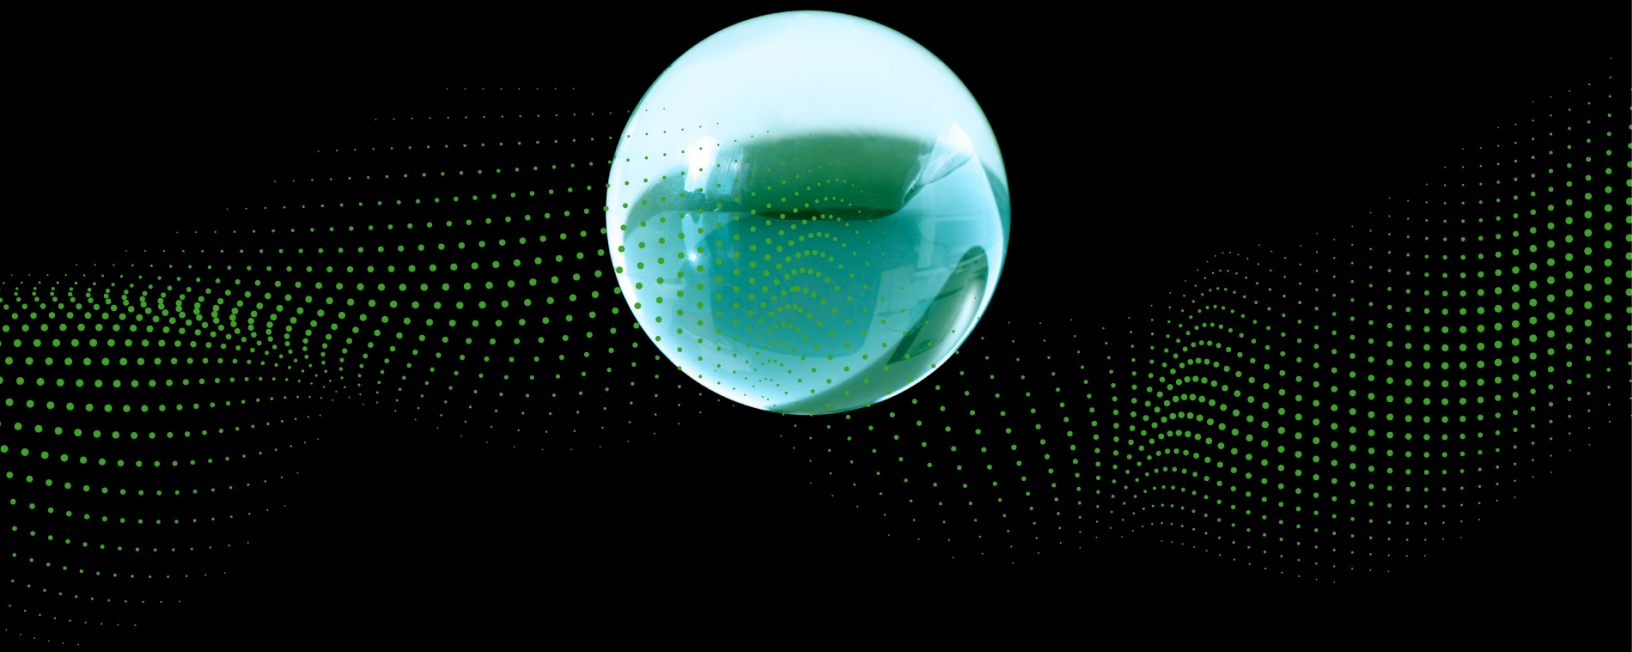

iab.

AI Transparency and Disclosure Framework

Regulatory & Standards Alignment

Global transparency rules for AI in advertising are converging on a common principle: truthful representation, but diverging in enforcement method. The regulatory landscape is evolving rapidly, and the specifics outlined here reflect the state of policy as of November 2025. This section summarizes regulatory trends and outlines how industry frameworks provide the practical bridge between law, technology, and market adoption.

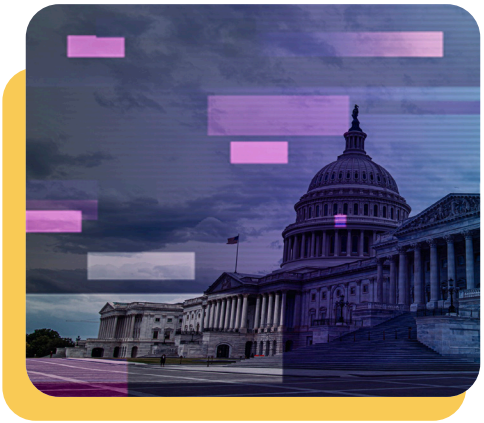

Regulatory Landscape

AI advertising disclosure exists in a fragmented global environment. The EU AI Act establishes mandatory baseline transparency requirements, requiring disclosure for AI-generated content with limited exceptions. In contrast, the U.S. maintains principle-based federal regulation focused on deception rather than mandatory labeling, though state activity is creating jurisdictional fragmentation. Regardless of regulatory approach, advertisers remain liable for misleading content. What is changing: explicit transparency obligations are intensifying as AI use in advertising scales.

United States: The FTC applies existing truth-in-advertising law to AI-generated content. Disclosure is required when a representation, omission, or practice would mislead a reasonable consumer and is likely to affect a decision or conduct, a "materiality" requirement for deception. Presenting synthetic personas or AI testimonials as authentic may trigger enforcement, while AI used as a production tool typically does not. The U.S. federal model remains principle-based, focused on deception rather than mandatory labeling.

State activity creates fragmentation. For example, California's SB 942 mandates generative AI disclosures (effective August 2026) whereas New York's Synthetic Performer Law (signed December 2025, effective June 2026) requires conspicuous disclosure of AI-generated performers in advertisements, with civil penalties of $1,000-$5,000 per violation. Brands face divergent rules across jurisdictions which present numerous challenges in adapting.

European Union: The EU AI Act, adopted in 2024 and entering into full force in 2027, with some provisions already active, mandates "AI-generated" labels as a consumer protection baseline. Content generated or modified by AI must be clearly disclosed, with limited exceptions for satire, art, and fiction. Unlike the U.S. state-led regulation on case-by-case instances, the EU establishes transparency as the default expectation. The Digital Services Act and Code of Practice on Disinformation extend this to platform accountability, requiring labeling of synthetic content.

APAC Region: China leads with assertive mandates requiring explicit labeling of all AI-generated synthetic content plus embedded metadata. South Korea enacted advertising-specific disclosure requirements through revisions to its Telecommunications Act (effective early 2026), mandating blanket labeling of all AI-generated photos and videos in advertisements regardless of materiality or consumer deception risk. The law emerged in response to widespread deceptive advertising using deepfake celebrities and fabricated experts, with platform operators held responsible for advertiser compliance and punitive damages up to 5x actual losses for malicious distribution of false content. This blanket approach contrasts with the materiality-based framework recommended here, which focuses disclosure on AI use that could mislead consumers rather than requiring labels on all AI-involved content. Japan, Singapore, India, and Australia favor principle-based, voluntary approaches, though evolution toward formal regulation is expected. Over 16 APAC jurisdictions have initiated AI governance.

Industry and Technical Alignment

C2PA and Content Credentials

The Coalition for Content Provenance and Authenticity (C2PA) standard provides the most mature technical basis for AI provenance. Its "Content Credentials" use cryptographically signed metadata to record how assets were created and edited. Major platforms, including Google, Meta, TikTok, LinkedIn, and OpenAI, are integrating it into content workflows. C2PA verifies origin and edit history, not truthfulness, but gives advertisers and regulators an auditable foundation for disclosure enforcement.

Current limitations include low adoption, documented bypass vulnerabilities, no accuracy verification, and watermark removal challenges. Provenance alone cannot tell consumers which specific elements were AI-generated. Technical standards provide the audit layer, but cannot substitute for human-determined disclosure.

Advertising Self-Regulation and IAB Initiatives

The ICC Marketing Code (2024) and ICAS' "Beyond Simple Labelling" extend long-standing advertising ethics to AI. Both reject blanket labeling in favor of case-by-case disclosure when AI materially alters perceived authenticity. This framework adopts the same materiality-driven approach, requiring disclosure when AI creates deception risk, rather than for every production tool.

Legal vs. Self-Regulation

Current law across the FTC and state systems prohibits deception and mandates ad identification, but it does not define detailed disclosure procedures. Self-regulation fills that gap by harmonizing interpretation across platforms and jurisdictions. Industry codes have evolved faster than legislation, balancing trust and feasibility while reducing "label fatigue" and the implied-truth effect, in which unlabeled content might seem more credible or truthful by comparison.

Platforms have established their own policies, subject to ongoing updates and refinements:

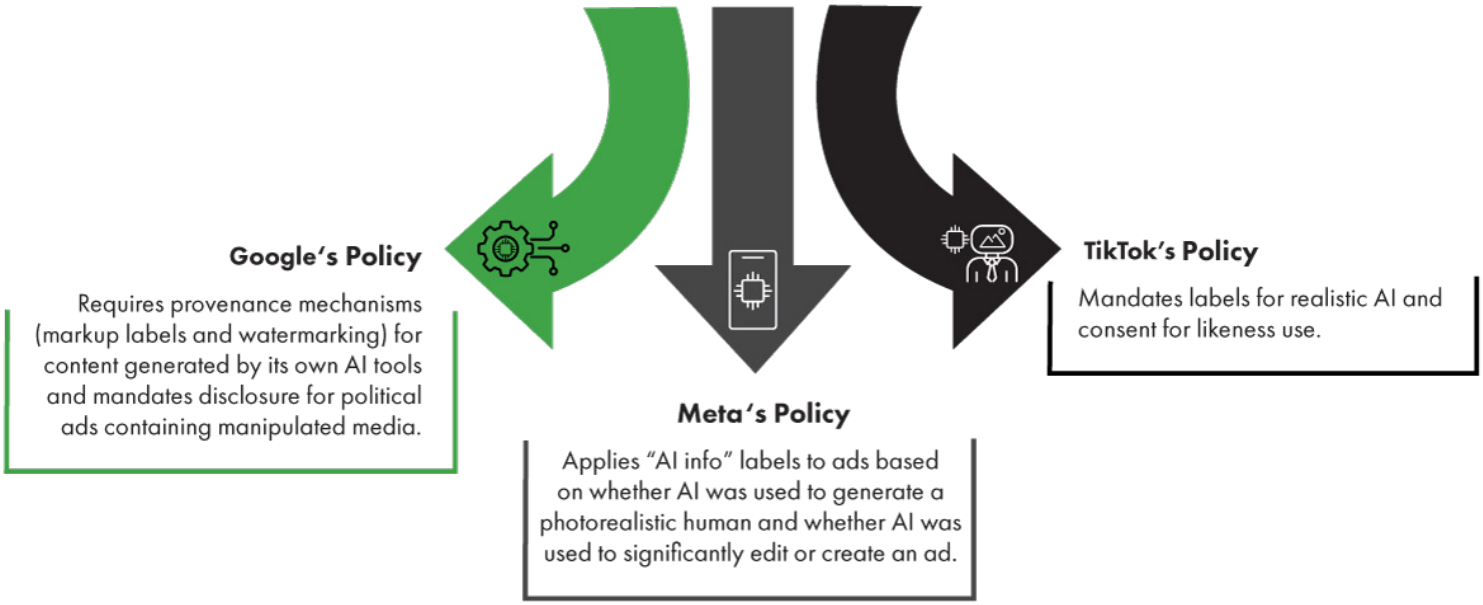

Self-regulation works best complementing rather than substituting legal compliance. Advertisers must meet mandatory requirements first and then apply industry frameworks to establish consistency and fill the grey areas. This positions the industry proactively, building trust through transparency while avoiding the operational burden of over-disclosure.

Implementation Roadmap

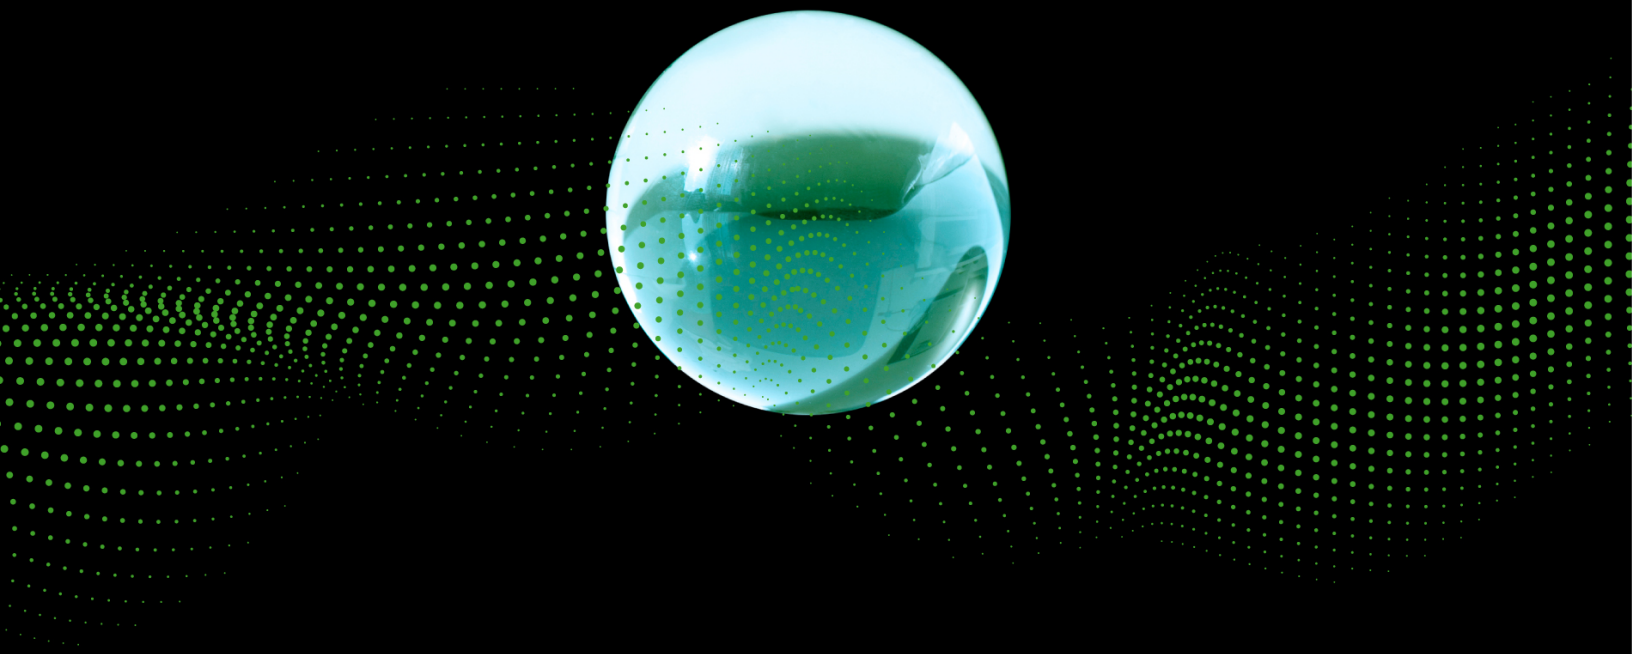

iab.

AI Transparency and Disclosure Framework

Implementation Roadmap

Building a comprehensive AI disclosure framework requires coordinated action across the advertising ecosystem. This roadmap translates the principles and requirements outlined in previous sections into a practical, phased implementation plan that addresses the core question: how do organizations operationalize the "Do we disclose? What? When? How?" framework at scale?

Phase 1 (0 – 6 Months): Foundation

Objective: Establish clear accountability, team readiness, and pre-publication review before large-scale implementation.

Key Actions

- **Assign a Lead:** Organizations should designate one person within marketing operations, brand safety, or compliance as the AI Disclosure Lead. This person serves as the go-to resource for disclosure questions and coordinates with legal on ambiguous cases. They are not meant to be an approval bottleneck. The goal is to make disclosure decisions routine across all teams rather than centralize every decision through one person.
- **Team Training:** The AI Disclosure Lead conducts concise sessions with creative, production, media, and comms teams so each function can apply disclosure thresholds directly in daily work.
- **Quality Assurance Processes**
  - Use a Pre-Launch Checklist Every campaign using AI should pass through a brief disclosure checklist before launch:

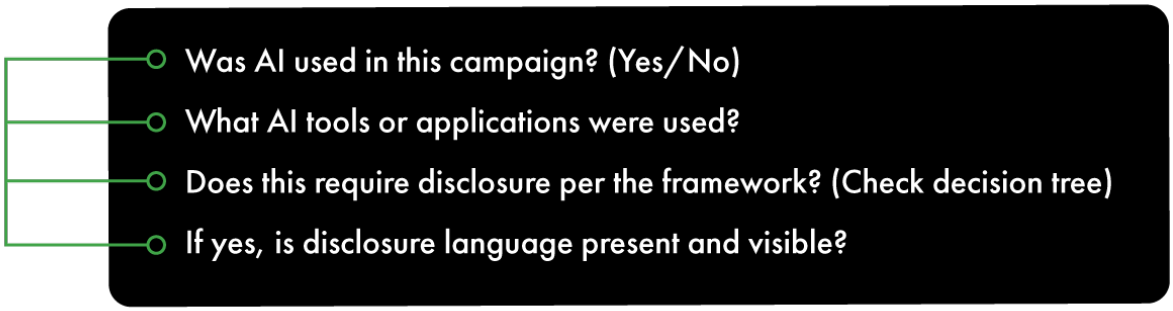

  - This checklist should live in existing workflow tools (project management software, trafficking systems, approval platforms) rather than requiring a separate process.
- Integrate Disclosure Review into Existing Legal Review Legal teams already review creative for claims substantiation, regulatory compliance, and brand risk. Add AI disclosure to this existing review process. Legal confirms whether disclosure is required per the framework and whether the disclosure language is appropriate. This ensures compliance without creating a separate approval layer.

Deliverables

- Training materials and recorded sessions

Phase 2: Pilot and Operationalization (Months 6 – 12)

Objective: Test disclosure mechanisms in live campaigns, evaluate consumer and operational impacts, and begin building technical infrastructure to support disclosure at scale.

Key Actions

- **Pilot Campaigns:** Launch controlled pilots across different AI use cases and disclosure methods. Start with one or two content types (e.g., image and video) to keep pilots manageable.
- **Measure What Matters:** Track pilot metrics across three dimensions:

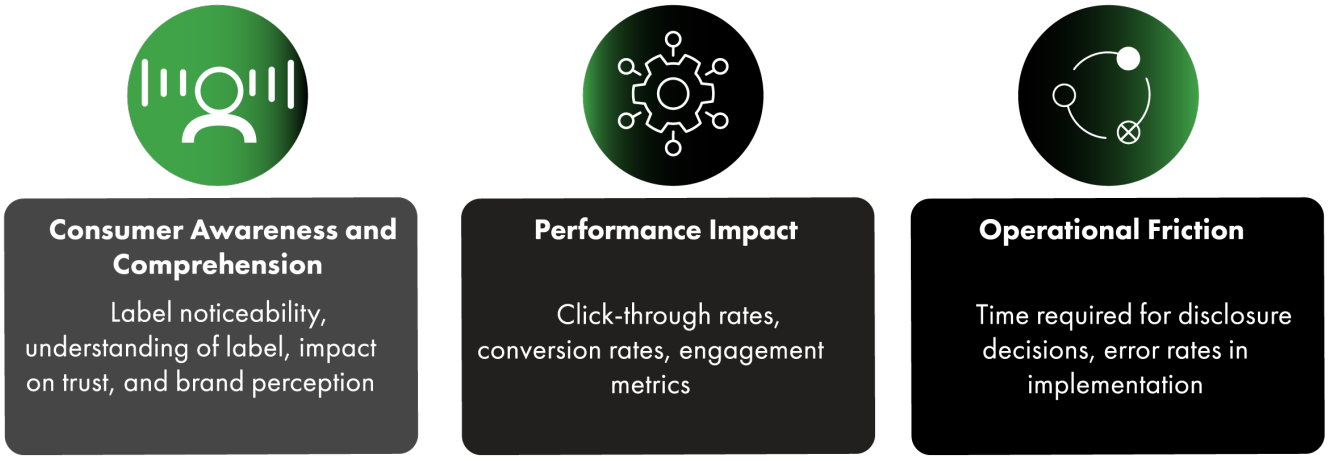

- **Cross-Partner Collaboration:** Coordinate with platforms (e.g., Meta, YouTube, TikTok, programmatic partners) to understand platform policy requirements and technical capabilities to ensure disclosures appear consistently across ad placements and formats.
- **Vendor Compliance:** Update partner contracts and creative briefs to include AI disclosure requirements. Organizations with technical maturity may begin requiring C2PA-aligned metadata from vendors.

Deliverables

- Pilot campaign performance report (including impact on performance and operational friction)
- Revised vendor and agency agreements

Phase 3: Scale and Refinement (Months 12 – 24)

Objective: Scale disclosure practices across all relevant campaigns and content while establishing ongoing measurement and governance to ensure sustained compliance.

Key Actions

- **Enterprise Rollout:** Scale disclosure practices across all relevant campaigns and content, operationalizing lessons learned from pilots. Streamline approval workflows and automate compliance checks wherever possible to reduce manual burden.
- **Measurement Framework:** Establish ongoing measurement assessing effectiveness (campaign performance and consumer sentiment). Recognize that consumer response to AI disclosure will evolve as familiarity increases.

Initial reactions may reflect novelty rather than long-term impact. Track trends over time rather than drawing conclusions from early data.

- **Policy Updates:** Monitor regulatory developments and adjust disclosure guidance as new requirements emerge.
- **Build Technical Infrastructure:** Begin evolving systems to support disclosure at scale. Creative management systems need fields to flag AI-generated assets with disclosure requirements automatically surfacing during trafficking. Ad servers require the capability to append disclosure language based on campaign tags. CMS platforms need automated labeling for AI-generated content.

Deliverables

- Performance and sentiment tracking report
- Automated compliance tooling integrated into existing systems

Success Metrics

Implementation effectiveness should be evaluated using measurable indicators:

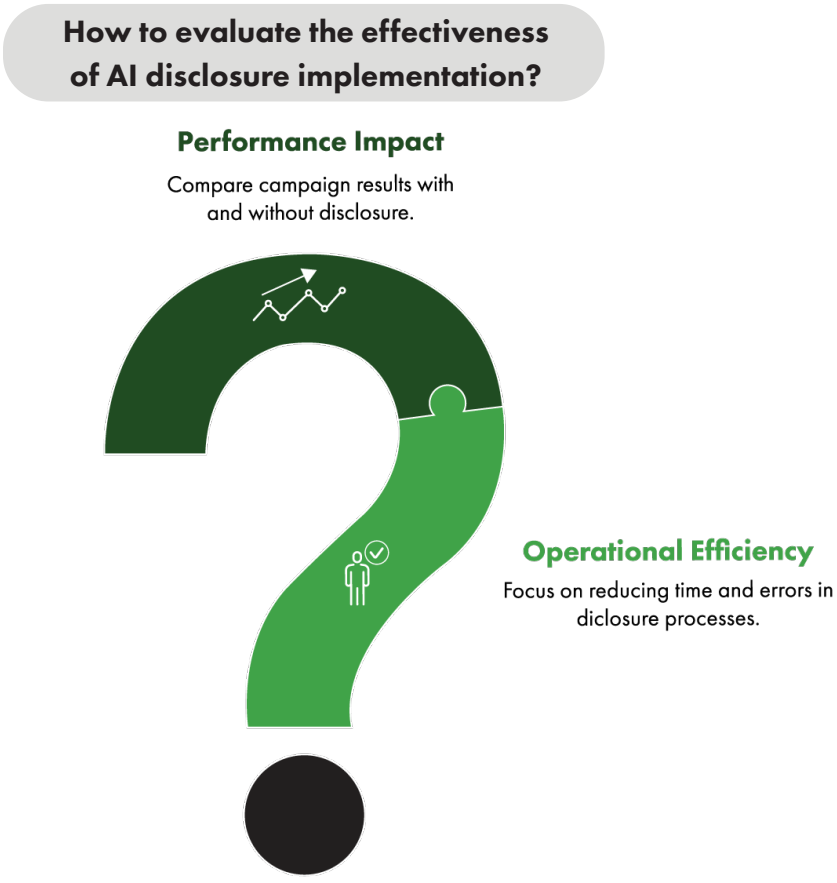

Key Takeaways

Successful implementation requires treating AI disclosure as an ongoing operational capability, not a one-time compliance project. The phased approach allows organizations to build competency progressively while adapting to evolving expectations. The goal is institutionalizing disclosure practices so thoroughly that they become automatic, positioning organizations to innovate with AI confidently while maintaining consumer trust and regulatory compliance.

Future Outlook

The next phase of AI transparency will be shaped by three converging forces: real-time content generation, advances in provenance technology, and evolving consumer and regulatory expectations. Each is already reshaping, and will continue to redefine, what constitutes sufficient disclosure and how it can be operationalized at scale.

Real-Time Personalized AI Advertising

AI is increasingly generating ads in real-time, customized for each user. Cadbury's campaign produced 130,000 unique AI videos for local Indian businesses, demonstrating both scalability and disclosure complexity. When creative is generated in the hundreds of thousands, pre-launch human review becomes impractical. Organizations will need automated compliance mechanisms which trigger disclosure based on pre-determined thresholds, embedded within AI systems. Personalization complicates this further as audiences expect tailored content, yet effectiveness depends on authenticity. The industry must establish conventions for signaling AI involvement without compromising the appeal of personalization.

Advances in Provenance, Watermarking, and Detection

Technical transparency is maturing. Platforms including Google, Adobe, and OpenAI are expanding the use of C2PA content credentials and watermarking systems such as SynthID to provide assertive provenance metadata at creation. These mechanisms are moving toward interoperability across media formats and integration into publishing workflows, enabling continuous verification of origin and modification history. However, provenance alone cannot guarantee audience understanding; metadata is invisible to consumers and susceptible to removal or manipulation. Thus, effective adoption pairs assertive provenance with visible human-facing labels and contextual information, as outlined in this Framework.

Regulatory Evolution and Enforcement

The regulatory landscape is fragmenting. The EU AI Act mandates "AI-generated" labeling as baseline, with limited exceptions. U.S. regulation remains principle-based federally, but state activity has been accelerating: California's SB 942 requires generative AI platforms with over 1M users, such as DALL-E, Sora, and Google's Nano-Banana to embed metadata disclosures and provide optional visible labeling for AI-generated image, video, and audio content beginning in 2026,while New York's Synthetic Performer Law (effective June 2026) requires disclosure of AI-generated talent in advertising. In Asia-Pacific, China leads with comprehensive requirements, while South Korea has enacted blanket labeling mandates for all AI-generated photos and videos in advertisements (effective early 2026), exemplifying the risk of over-regulation when industry fails to self-regulate effectively. Over 16+ APAC jurisdictions have initiated AI governance.

Per regulation updates in Dec.

This fragmentation creates dual challenges:

1. **Consumers face confusion** navigating inconsistent label types and meanings across platforms and jurisdictions.
2. **Advertisers bear significant operational burde**n complying with divergent requirements for cross-border campaigns.

Brands failing to disclose material AI involvement will increasingly risk enforcement, backlash, and reputational damage.

Evolving Consumer Expectations

As AI usage becomes ubiquitous, the meaning of disclosure will evolve from novelty to norm. Early friction, including confusion, skepticism, and/or "label fatigue," is expected to diminish once standardized labels are consistent and contextual. Research from Yahoo and Publicis shows that visible, credible disclosures increase trust by up to 73% while IAB research finds that consumer perceptions of advertisers using AI are mixed, as some view it as creative and others as inauthentic. In addition, over-labeling can erode perceived authenticity. Periodic industry calibration will be needed to reassess thresholds as familiarity rises and risk tolerance shifts.

Takeaway

Future transparency will rely on automation, interoperability, and adaptive governance. Disclosure will move from static labeling toward system-level assurance with machine-readable provenance verified at scale and human-facing clarity delivered consistently.

That said, this Framework reflects a moment in time within a rapidly evolving landscape. Consumer expectations, regulatory interpretations, and platform practices will continue to shift as AI becomes normalized. Maintaining trust will require ongoing recalibration of disclosure norms rather than adherence to a static rule set.

Recommendations & Next Steps

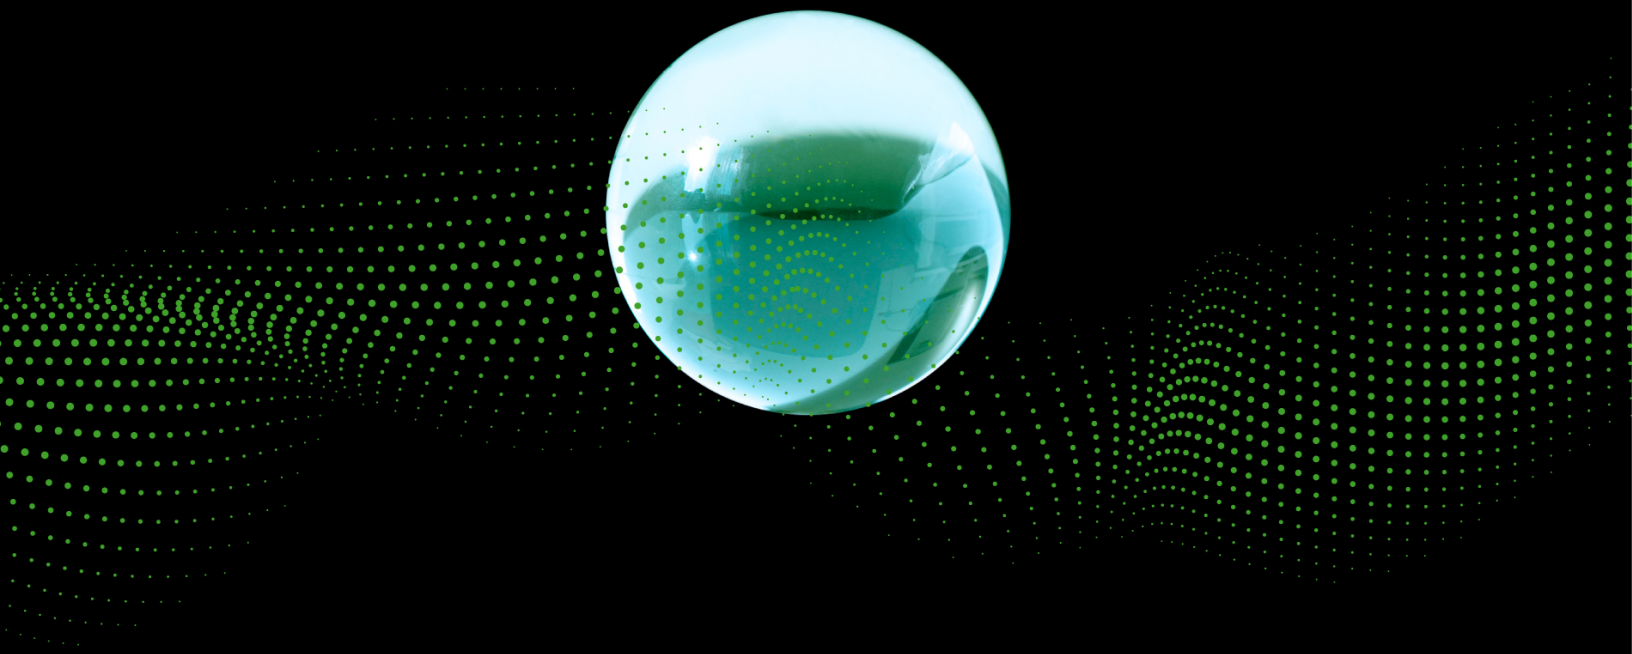

iab.

AI Transparency and Disclosure Framework

Recommendations & Next Steps

Minimum Viable Practices (MVP)

IAB recommends that advertisers, agencies, publishers, and platforms implement the following baseline actions:

- **Assign an AI Disclosure Lead** within marketing operations or compliance to oversee materiality assessments, documentation, and internal training.
- **Adopt the two-layer disclosure model outlined in this Framework:** (1) human-facing disclosures when AI involvement materially misleads a consumer, and (2) C2PA metadata to ensure machine-readable provenance.
- **Integrate IAB disclosure thresholds into creative and production workflows using a pre-launch checklist:** Was AI used? What tools? Does this trigger disclosure? Is the label present and visible? Disclosure decisions should occur pre-launch, not retroactively.
- **Document AI** use across campaigns to support auditability
- **Use standardized labels** ("AI-generated image," "AI-generated person") applied visibly and consistently, per IAB and accessibility guidelines.

These practices require minimal infrastructure and can be implemented within 60 days.

Industry Alignment Actions

Cross-sector coordination is essential to avoid fragmentation and establish trust:

- **Establish vendor alignment.** Update creative briefs and contracts to include AI disclosure. Agencies and vendors should flag AI involvement and apply disclosure when thresholds are met.
- **Pilot C2PA integration.** Test IAB custom assertions (com.iab.threshold, com.iab.disclosure) across platforms to verify metadata accuracy and enforcement consistency.
- **Platform enforcement.** Publishers and platforms are encouraged to flag ads lacking disclosures or accurate C2PA metadata for advertiser review and remediation.
- **Education and training.** Establish role-specific training for staff on consumer transparency and disclosure obligations.

Priority Areas for Future Standards Development

The following areas require ongoing standards work and cross-industry research:

- **IAB Tech Lab metadata specification.** Develop advertising-specific C2PA assertions (com.iab.threshold, com.iab.disclosure) integrated with OpenRTB, VAST/VPAID, and ad serving protocols to enable programmatic disclosure verification.
- **Verification protocols.** As third-party mechanisms become available to confirm disclosure claims, detect metadata tampering, and validate label-metadata alignment, and integrate them into best practices.
- **Cross-jurisdictional consistency.** Continue engagement with ICAS, EASA, and other trade bodies to harmonize AI disclosure standards across markets while maintaining a risk-based approach.

Call to Adoption

IAB encourages all members to:

- Publicly commit to the Framework principles of transparency, proportionality, consistency and clarity.
- Pilot implementation across at least one campaign within six months.
- Share outcomes and feedback through the appropriate IAB working group to inform 2026 standards development.

Early adoption positions the industry to maintain consumer trust, ensure responsible innovation, and avoid regulatory mandates. By embedding transparency into creative, operational, and technical layers, the advertising ecosystem can lead in defining credible, scalable AI disclosure practices that sustain both trust and growth.

Acknowledgements

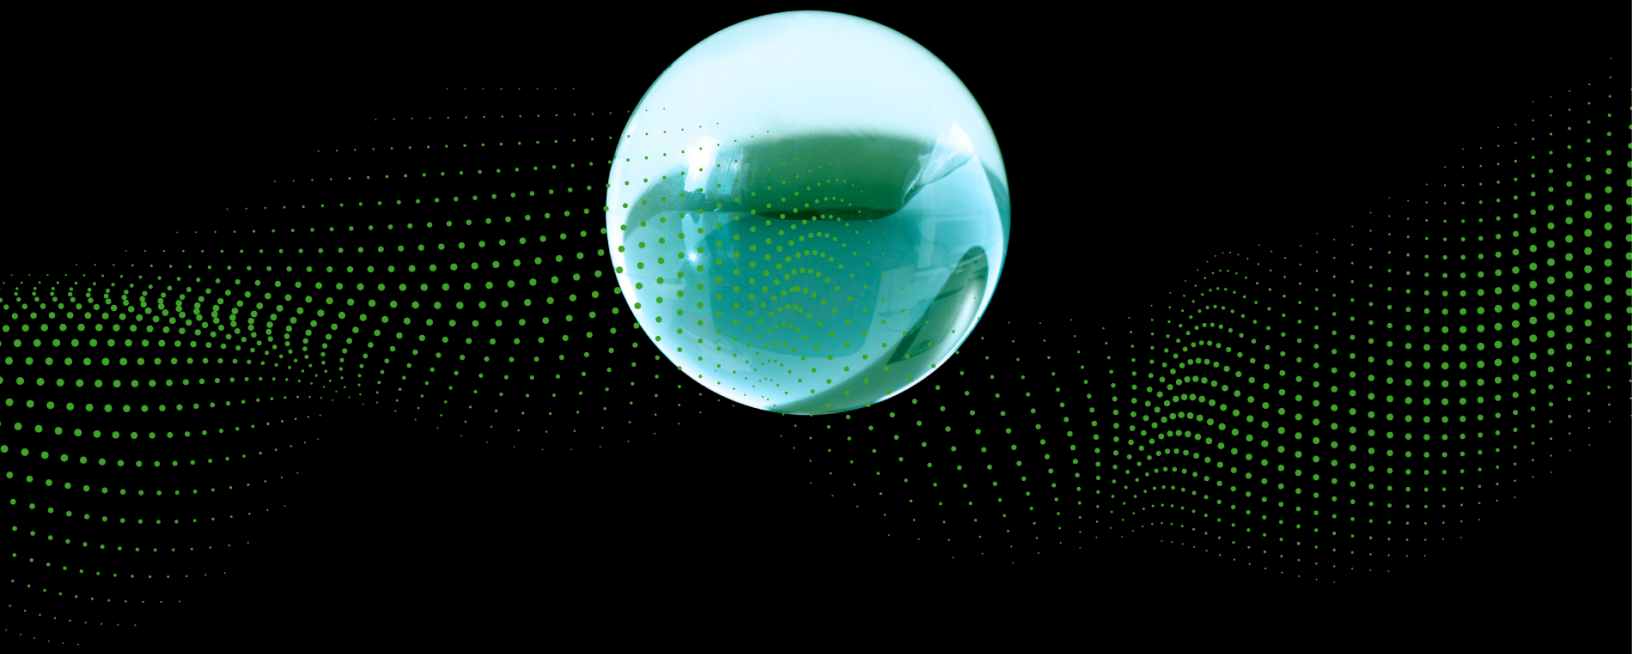

iab. AI Transparency and Disclosure Framework

Acknowledgements

The IAB AI Transparency & Disclosure Working Group is a cross-industry team of leaders from brands, agencies, publishers, and technology platforms who collaborated to create this framework. The group's goal is to provide the industry with practical, human-centered guidance on when, how, and to whom to disclose AI-generated content in advertising, establishing transparent frameworks that build consumer trust, ensure regulatory compliance, and balance innovation with accountability while avoiding undue operational burden on advertisers.

This working group is led by Caroline Giegerich, VP, AI at IAB, with special thanks to members:

Adam Buhler
Executive Vice President, Creative Technology
Digitas

Caroline Giegerich
VP, AI
IAB

Graham Wilkinson
EVP, Chief Innovation Officer, and Global Head of AI
Acxiom

Ken Fisher
Founder & Editor-in-Chief, Ars Technica
AI Council Editorial Lead
Condé Nast

Molly Crawford
Executive Vice President & Head of Public Policy
IAB

Steve Bagdasarian
Chief Commercial Officer
Comscore

Yury Adamov
Director – Global Corporate Strategy
Warner Brothers Discovery