TECHNOLOGY

AI Transparency and Provenance Principles for Marketers

JUNE 2026

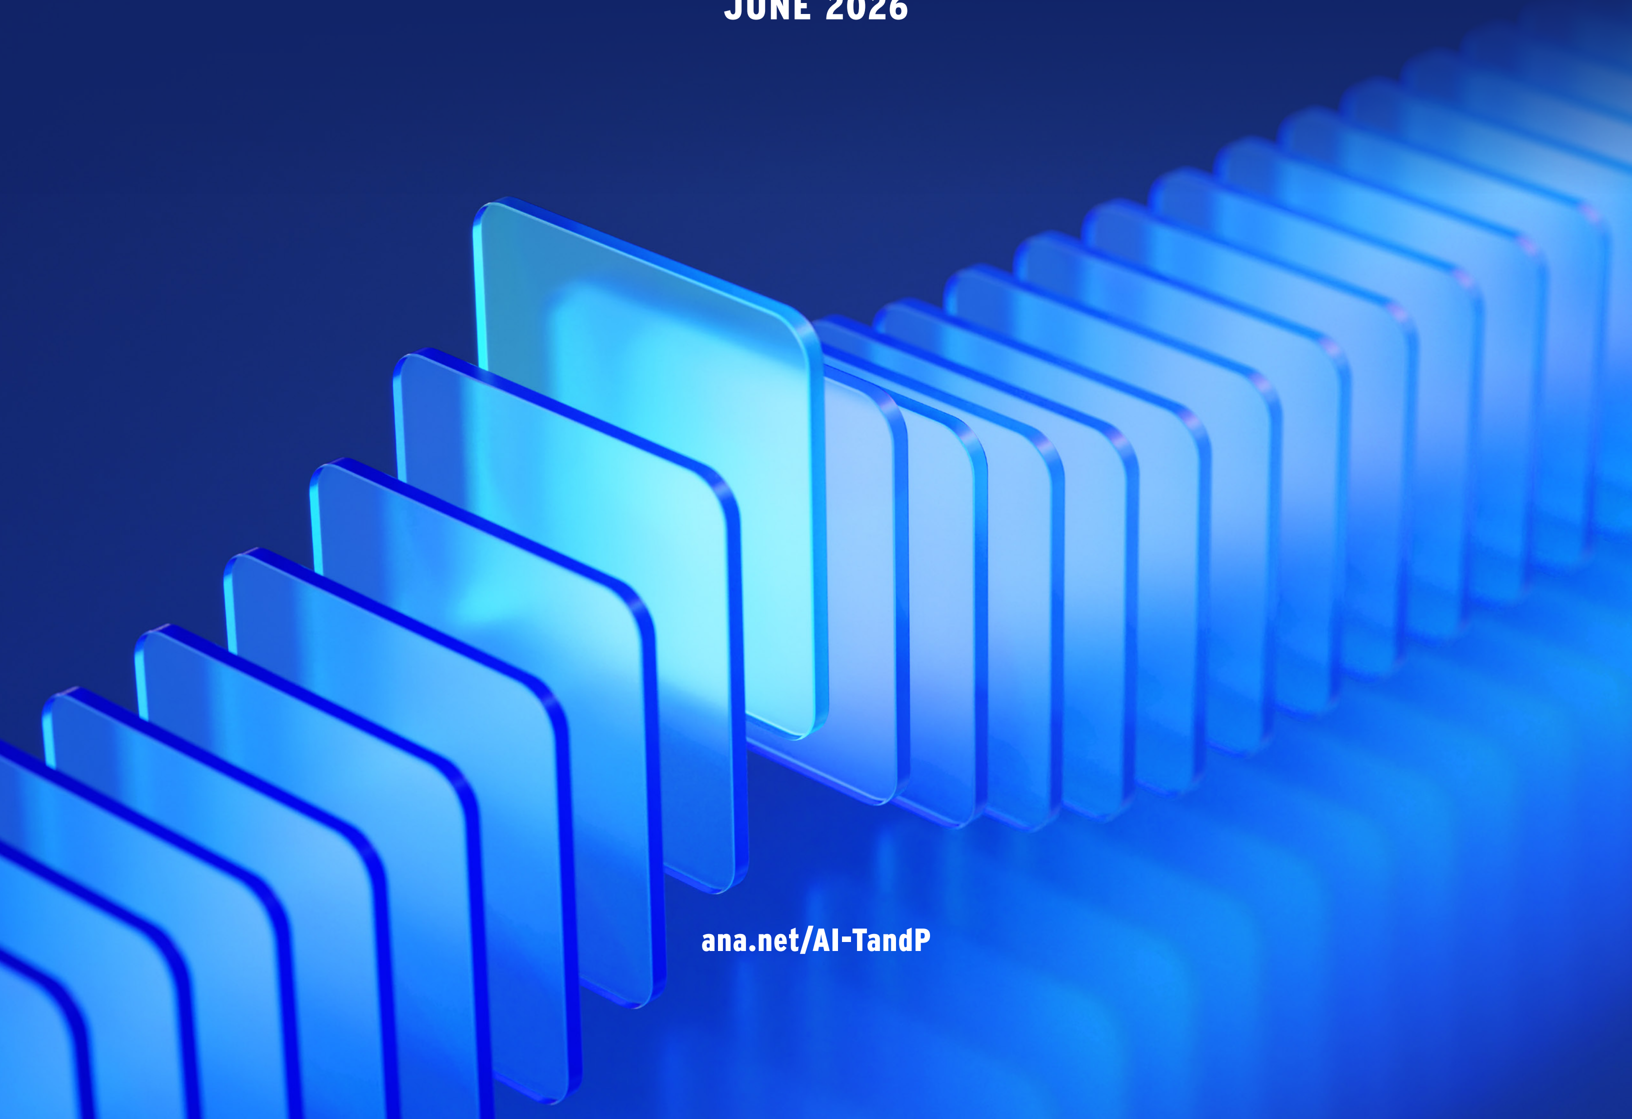

ana.net/AI-TandP

TABLE OF CONTENTS

EXECUTIVE SUMMARY

The irresponsible use of generative artificial intelligence by marketers has significant potential to damage the relationship between consumers and brands. As evidenced by social media sentiment, consumers can sense when a brand uses artificial intelligence (AI) technology, especially when it's used poorly. The responsible use of AI, where brands are transparent about their usage, respect creators, and put the same level of thought and creativity that they do with traditional advertising and marketing, has tremendous potential to strengthen their relationship with consumers.

Regulation of AI in advertising is beginning to take shape across multiple jurisdictions. While comprehensive federal oversight in the U.S. has not yet emerged, state governments are already introducing transparency requirements related to AI-generated content. At the same time, international regulatory frameworks, including the EU's AI Act and emerging rules governing synthetic media in several global markets, signal a clear movement toward disclosure, accountability, and responsible AI deployment.

These developments indicate that the regulatory environment for AI in marketing will continue to evolve rapidly, and likely unevenly across regions. For marketers, this creates both operational complexity and reputational risk. Waiting for fragmented mandates to dictate compliance may result in inconsistent practices, increased costs, and diminished consumer trust.

A proactive, industry-led approach offers a more effective path forward. By establishing voluntary principles for transparency, disclosure, and provenance, the marketing industry can create consistency across jurisdictions while reinforcing trust with consumers, regulators, and partners.

In response, the ANA, through a structured process with its members, has developed a set of consensus-driven guidelines designed to support responsible use of generative AI in advertising. These recommendations align existing legal principles with emerging standards, including efforts such as the Content Authenticity Initiative and the Coalition for Content Provenance and Authenticity (C2PA), to provide practical guardrails for disclosure, provenance, and ethical AI deployment.

Our goal is to equip ANA members with a clear, adaptable, voluntary framework that promotes transparency, protects brand integrity, and ensures that consumers remain informed when interacting with AI-enhanced marketing content. Marketers who adopt common principles now can lead responsibly as the regulatory landscape continues to evolve.

EXECUTIVE SUMMARY

The ANA established a working group within its AI Committee to address the question of if, when, and how to disclose the use of generative AI in marketing and advertising. The working group addressed transparency, which is the "above-the-glass" user experience (i.e., what the user can see), and provenance, which is the "below-the-glass" technology designed to authenticate creative work, distinguish original content from knockoffs, and reduce legal and reputational risk.

The working group shared its findings with the AI Executive Leadership Council, a collection of industry leaders who own the AI for marketing remit within their organizations, and they agreed that it is appropriate, respectful of consumers, and good for the brand to be transparent around the use of AI when it may be viewed as misleading. There is consensus among this collective group of marketers that the definition of misleading is relative to the AI's impact on the message and the context of the advertised product or service. Additionally, this group proposes alignment (with expansion of the currently tracked provenance elements) with the Content Credentials digital chain-of-custody (provenance) container created by the C2PA.

The transparency principles developed in this paper are goals to which members can aspire to put into practice. Many of these include broader considerations that may require separate policies or legal review. The ANA recommends that marketers consult their legal teams on these topics, as the points outlined in this framework may be affected at the marketer level by several factors, including the extent of that marketer's regulatory obligations.

Ongoing advances in AI are providing seasoned creative professionals with tools, personalization options, and the ability to significantly expand creative ideation in areas once constrained by time and budget. On the other hand, the affordability, simplicity, and quality of these tools are democratizing the ability for anyone to create polished content, including short- and long-form video around brands, enhancing the risk that unsanctioned messages about a brand could proliferate, create consumer confusion as to what is real, and, in turn, erode the hard-earned trust that marketers have built with their consumers. Brands need to protect themselves, both for authenticity and for legal reasons, through robust provenance.

The working group and AI Leadership Council have come together on the following provenance pillars:

1. The industry seems to be coalescing on the Coalition for Content Provenance and Authenticity and its technology, Content Credentials, which provides a foundation of verifiable provenance that may be helpful for members (although it is currently limited and has its flaws). Members should explore this option.
2. Marketers are stakeholders — not system architects — and should collaborate across the ecosystem.
3. Understanding that some teams operate under key bandwidth constraints, the elements of provenance should include core and optional tags. The more robust the provenance, the more protected the marketer can be in the future.

EXECUTIVE SUMMARY

> **Generative AI has advanced to the point that YouTube CEO Neal Mohan wrote in his 2026 annual letter, "It's becoming harder to detect what's real and what's AI-generated. This is particularly critical when it comes to deepfakes. We clearly label content created by YouTube's AI products, and creators must disclose when they've created realistic altered or synthetic content. Because labels aren't always enough, we remove any harmful synthetic media that violates our Community Guidelines. We're also building on the foundation of Content ID [YouTube's automated content identification system] — a system our partners have trusted for well over a decade — to equip creators with new tools to manage the use of their likeness in AI-generated content. Finally, we remain committed to protecting creative integrity by supporting critical legislation like the NO FAKES Act."**

In an era when generative AI is reshaping the business landscape, the marketing industry has a unique opportunity and responsibility to lead with integrity. By establishing voluntary principles and fostering cross-sector collaboration, the industry can build a foundation of trust that benefits consumers, brands, and the broader digital ecosystem.

Transparency and provenance should not be seen only as a defensive safeguard against risk, but also as an opportunity. Over time, repeated responsible disclosure creates a new expectation that brands are not just innovative in their use of technology, but trustworthy in how they apply it. In this sense, transparency becomes more than compliance; it becomes a brand asset, deepening loyalty and building long-term equity in an environment where consumer trust is increasingly fragile.

This guidance will be updated periodically to reflect changes in law, platform capabilities, and consumer expectations. As AI technology advances, so will risks, regulations, and best practices. The ANA will continue to convene members, monitor global developments, and refine this guidance.

Michael Donnelly
EVP,
Technology Practice
ANA

Jeff Minsky
VP, AI and Marketing Futures,
Technology Practice
ANA

TRANSPARENCY

Transparency helps ensure comprehension and contributes to authenticity between a marketer and a consumer.

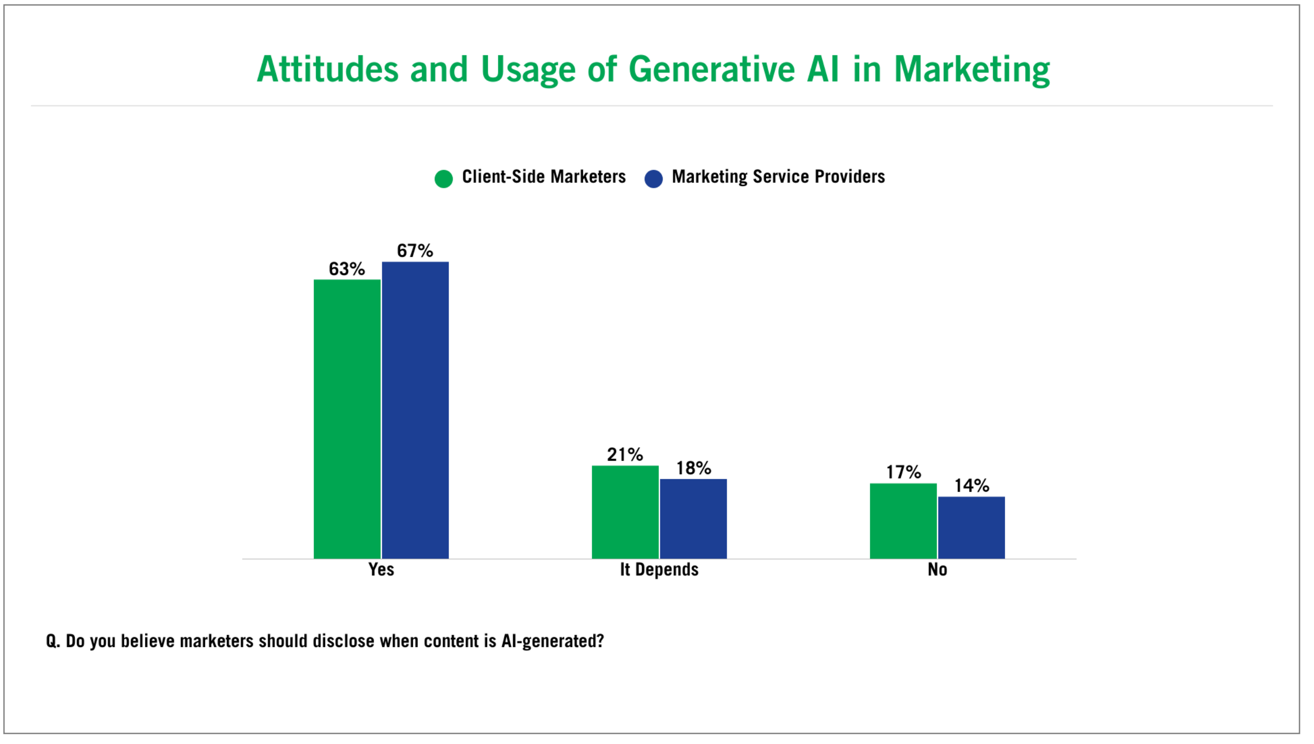

In 2025, a survey was fielded to the full membership of the ANA on "Attitudes and Usage of Generative AI in Marketing." One of the survey questions asked, "Do you believe marketers should disclose when content is AI-generated?"

Sixty-three percent of client-side marketers (CSMs) said "Yes," 21 percent said "It depends," and 17 percent said "No." Sixty-seven percent of marketing service providers (MSPs) said "Yes," 18 percent said "It depends," and 14 percent said "No."

Over time, transparency around the use of generative AI, implemented through repeated responsible disclosure, creates a new expectation that the brand is not just innovative in its use of technology but also trustworthy in how it applies it. In this sense, transparency evolves into a brand asset, one that deepens loyalty and builds long-term equity in an environment where consumer trust is increasingly fragile.

The ANA has identified several key tensions that should be considered in establishing effective AI disclosure practices (please note that this is a snapshot in time from Q4 2025 and may evolve in the coming months and years):

TRANSPARENCY

- **Technical feasibility challenges**. While metadata tagging and watermarking are built into generative AI creative applications from Meta, Google, and Adobe (Firefly), adoption across platforms and partners remains inconsistent.
- **Legal uncertainty**. Copyright, likeness, and intellectual property issues remain unsettled and subject to evolving interpretation, complicating decision-making.
- **Risk-based approach needed**. Not all AI uses carry the same level of risk, requiring differentiation between high-risk applications (where robust disclosure is essential) and lower-risk uses (where minimal notice may suffice).
- **Multi-jurisdictional complexity**. Global aspirations for consistency often conflict with varied local regulations and cultural expectations across different jurisdictions.

Transparency Principles

The ANA AI Marketer Leadership Council and AI Working Group believes that clearly defined principles are necessary and has drafted the list below:

1. A human should retain ultimate editorial, legal, and ethical accountability when generative AI is utilized for any element in the creation and/or production of an advertisement or commercial marketing asset.
2. Unless mandated by a law that applies (local, state, federal), disclosure of the use of generative AI is a matter of degree of use, intent, and the relative level of general regulation of industries. Context is essential.
3. The benchmark around disclosure of AI use in advertising creative should be centered around how prominent, in the eyes of a reasonable consumer, the use of generative AI is to the creation of the end product.
4. Marketers must ensure AI-generated depictions of people, products, or performance claims are not deceptive and remain consistent with existing advertising laws, especially when generative AI enhances realism.
5. Marketers should abide by jurisdiction-specific disclosures requirements, such as New York's Synthetic Performer Act requiring disclosure when AI is used to depict a person in an ad.
6. Marketers and their agencies should embrace commercially reasonable best efforts to help ensure the generative AI tools that they choose to use have been trained responsibly.
7. Marketers should default to a "be more, not less, transparent" mentality, which the working group believes will drive a greater sense of trust with consumers in the longer term.
8. If an AI-generated ad includes a disclosure on one platform (e.g., TV), the same disclosure should appear on derivative or repurposed versions (e.g., social, print, OOH).
9. Marketers should periodically survey consumers about their understanding and comfort levels with AI-generated content.
10. Where law or regulation does not already require disclosure, marketers should consider documenting their decision-making rationale on whether to disclose the use of AI to protect against future regulatory scrutiny.

TRANSPARENCY

11. Marketers should recognize that case law is not settled on the rights of tech companies to train LLMs, although progress has been made. Marketers should always take reasonable steps to ensure that proper indemnifications are in place, as reviewed by their own legal counsel.
12. Prior to using others' intellectual property, marketers should ensure that they obtain copyright and likeness permissions and license grants, especially with synthetic performers or voices.

For AI-generated audio/music, disclosure is not required when it is solely for background or functional purposes, and no representation is made regarding human composition or performance. Where the musical work is presented as an artistic expression or implies human authorship, disclosure, such as "music generated with AI technology," is recommended for transparency. Some states have enacted laws specifically designed to protect artists from AI-generated cloning of their voices and likenesses without their consent. For instance, Tennessee's Ensuring Likeness Voice and Image Security (ELVIS) Act includes key provisions that protect individual's voices, targeting platforms and services that produce AI-generated vocal or visual replicas and impose liability on distributors and platforms for unauthorized AI content. Notably, these protections apply to any person, not just celebrities or public figures.

In essence, explicit consent and licensing for AI-generated voices or likenesses of real people are now legally required in Tennessee, with other states likely to follow.

TRANSPARENCY

Additional Consideration: Risk-Based Disclosure Differentiation

A critical principle in shaping disclosure practices is potential for risk — the idea that not all applications of AI carry the same potential for consumer harm or misunderstanding. Uses, such as deepfakes, warrant the most rigorous standards, including prominent disclosures and robust provenance, to ensure consumers are not potentially misled. Other high-risk uses, such as synthetic endorsements, should be avoided entirely to help avoid risk under FTC rules prohibiting businesses from creating or selling fake AI-generated reviews or testimonials.

Medium-risk applications, such as artistic or fantastical content that substitutes for traditional creative methods, require a more balanced approach, with moderate disclosure sufficient to signal AI involvement without disrupting the experience. At the other end of the spectrum, low-risk uses that primarily support efficiency — such as background editing or grammar corrections — potentially pose little risk of consumer deception and therefore do not require disclosure. This tiered approach helps ensure that the rigor and visibility of disclosure scale is proportional to the potential for consumer impact, aligning protective measures with actual consumer risk.

Industry	Baseline Risk	Why It Matters	Minimum Disclosure / Provenance
Healthcare, Finance, Legal	High	High risk of consumer harm, heavy regulatory scrutiny, and significant ethical implications.	Prominent on-asset label; audible/visual notices for AI segments; C2PA credentials; audit trail; legal/compliance review; consent for likeness/endorsement.
Government & Public Sector	High	Public trust, democratic integrity, and statutory obligations demand maximum transparency.	Persistent on-screen or adjacent labeling; accessible public notice; C2PA with chain-of-custody; archival log and provenance verification.
Retail, CPG, Travel	Medium	Disclosure is needed when AI materially alters product or experience representations or steers purchase decisions.	Contextual labels near images/offers (PDP, "About this item"); flag AI-altered visuals; embedded metadata; elevate to High if claims/endorsements or realistic simulations are used.
Entertainment, Media & Gaming	Low	Typically lower stakes; disclosure becomes necessary when realism or persuasive/monetized intent is high.	Credits or description-level notice; embedded metadata; *escalate* to Medium/High for photorealistic likeness, synthetic endorsements, political content, or child-facing experiences.

Escalation Triggers

- Photorealism or realistic voice likeness
- Endorsements, advice, or claims that may be relied upon
- Vulnerable audiences (e.g., minors, medically sensitive contexts)
- Material financial or legal consequences

Good Practice

- Pair **direct** labels with **C2PA** credentials
- Keep disclosures proximate to the AI-influenced element
- Maintain an auditable log of AI tools, prompts, and reviews
- Pre-approve disclosure language with legal/compliance

Source: ANA AI Transparency and Provenance Working Group (AI-generated by Gemini)

TRANSPARENCY

Additional Consideration: High Importance

We consider AI usage to be highly important when AI-generated or AI-altered content substantially affects how an audience perceives the authenticity, credibility, or persuasive impact of advertising. It goes beyond routine creative enhancements (e.g., resizing, sharpening, color correction, background cleanup) to include AI involvement that could potentially mislead or confuse consumer decision-making. In practice, a change is present when AI alters the realism of what's shown, serves as the core persuasive element of the message, or creates ambiguity by blending synthetic and human-created media without providing context or necessary disclosures.

Three simple tests can be applied to determine if AI usage warrants disclosure.

1. **Realism**: AI-generated content resembles a real person, product, place, or event in a photo-realistic way and will most likely be mistaken for human-generated content. Some states, like New York, may require disclosure in this instance.

   Example: An auto brand uses AI to generate a photorealistic image of its new SUV driving through rugged terrain in Yosemite National Park, even though the car has never been to Yosemite and the photo shoot has never taken place. Because the AI rendering looks indistinguishable from a real photograph, consumers could believe it is an authentic event.

2. **Persuasive intent**: The AI-generated content is central to the message's persuasive intent and should be disclosed. Resizing, sharpening, color correcting, and background extensions are not prominent and do not need to be disclosed.

   Example: A skincare company deploys an AI-generated influencer to give a heartfelt testimonial about how the product cleared her skin. Unlike routine editing (e.g., sharpening or color correction), the AI avatar itself is the central persuasive device driving the ad's credibility. Some states, like New York, may require disclosure in this instance.

3. **Mixed media**: Combining AI-generated content with realistic human-created content can potentially confuse consumers. The combining of AI and human content could influence the audience's perception or decision-making regarding the product or service.

   Example: A travel brand produces a commercial combining real drone footage of a beach resort with AI-generated images of celebrities enjoying the property. The blend of authentic video with AI-fabricated experiences could mislead viewers about who stayed at the resort. Some states, like New York, may require disclosure in this instance.

Suppose the content is used in high-trust/highly-regulated contexts (e.g., health, finance, political messaging). In that case, any use of generative AI may require disclosure or may be prohibited under applicable law, regardless of the outcomes of the change test above.

TRANSPARENCY

Significant AI and Use Cases

Examples of significant use:

- Changing the appearance or language of an actor (outside minor corrections or touch-ups)
- Enhancing the look of a product being sold when such enhancements misrepresent or meaningfully alter how a reasonable consumer would perceive the product's qualities or performance
- Using synthetic actors or avatars, even in secondary or background roles where facial features are clearly recognizable (some states, like New York, may require disclosure in this instance)
- Scripts or advertising/marketing copy, AI-generated imagery, audio, or video that AI has predominantly created with minimal human oversight
- AI-generated transcripts or captions for accessibility. While this would seemingly be considered "non-significant," it is in the brand's best interest, given AI's penchant for hallucination and mistakes, to label this category as "AI-generated" if used. If transcripts and captions are reviewed and confirmed by qualified humans, then this category may be considered as non-significant.

Examples of non-significant use:

- Early-stage or background creative work
  - AI used for ideation, brainstorming, or inspiration/mood boards
  - Storyboarding assistance
  - Using AI for draft copywriting, such as first-draft headlines, alternative phrasing, or tone options, with final human editing
- Enhancements to a product's appearance are not significant when they constitute ordinary aesthetic refinements that do not misrepresent or meaningfully alter how a reasonable consumer would perceive the product's qualities or performance.
- Visual adjustments or post-production polishing
  - Lighting, color correction, and exposure balancing to create consistent tones across shots
  - Replacing or cleaning up backgrounds, such as removing trash, shading the sky a deeper (or lighter) shade of blue, or adjusting reflections
  - Eliminating camera artifacts or continuity errors (lens flares, dust spots, unintended shadows)
  - Minor aesthetic improvements to people or objects, such as removing wrinkles or blemishes on actors (except where it would be significant, like a dermatology ad), smoothing hair flyaways, removing lint, or brightening eyes; if these improvements don't misrepresent identity or product effectiveness for multi-platform formats (social, display, print)

TRANSPARENCY

- Technical or production efficiency
  - Generating temporary stand-in voices for timing or animatics to record actual human talent
  - AI-assisted subtitling or dubbing when translations are checked and approved by humans
  - Automating cropping, resizing, or layout optimization
- Operational and accessibility enhancements
  - Image alt-text generation for compliance or SEO, edited or verified by humans
  - Voice-over adjustment for clarity or pacing without altering the substance of the message
- Data-driven or analytical applications
  - Audience sentiment or trend analysis using AI to inform future creative direction
  - A/B testing automation to optimize messaging or visuals
  - Performance analytics predicting engagement or recall likelihood

HOW TO DISCLOSE: RIGHT TIME, RIGHT REASON, AND INCREASED TRUST

This framework is designed to cut through the noise by providing a practical, principle-based approach: disclosures should be used meaningfully, proportionally, and consistently so they help protect consumers, reinforce trust, and remain feasible to implement at scale. Effective disclosure requires a dual approach that combines direct, human-readable signals with indirect, machine-readable provenance.

There are two main methods to disclose the use of AI. “Above the glass,” which is visible to consumers, is through direct disclosure. “Below the glass,” which is visible to machines and data experts, is made through technical disclosure.

Direct disclosure ensures that consumers can immediately understand when AI has played a role in creating content. This may take the form of labels such as “co-created with generative AI,” captions or audio statements, platform-native disclosure tools (e.g., TikTok, Meta, YouTube toggles), or even on-pack and offline content labeling.

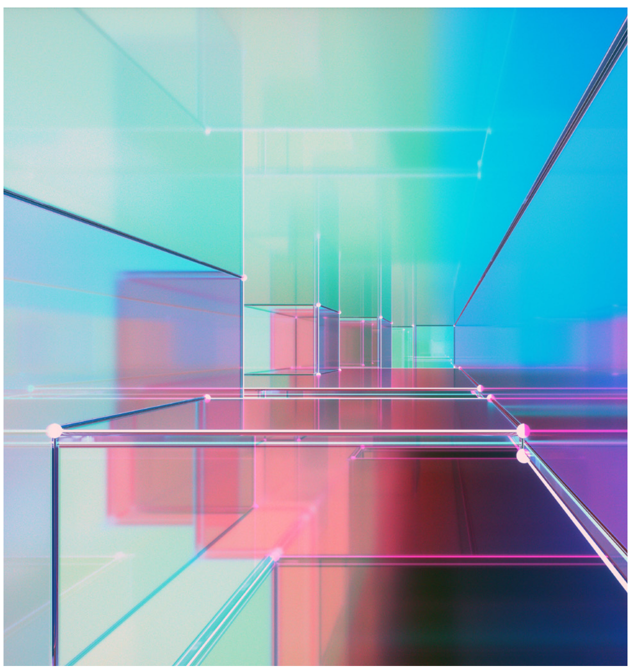

In parallel, technical disclosure provides the technical backbone by embedding C2PA credentials, metadata, cryptographic provenance, or watermarks into digital assets.

These machine-readable signals establish a verifiable chain of custody for content, enabling platforms, regulators, and auditors to validate authenticity.

Importantly, technical disclosure should be viewed as a complement, not a substitute, for direct disclosure when consumers are the intended audience.

Taken together, this layered approach balances transparency with credibility, building both immediate comprehension and long-term trust.

HOW TO DISCLOSE: RIGHT TIME, RIGHT REASON, AND INCREASED TRUST

Channel Specific Protocols

Channel	Disclosure Method
Social media	Platform toggle and post label (e.g., "This ad was co-created with generative AI")
Websites	Footer or FAQ link and embedded metadata
Email	Footer disclosure
E-commerce	"About this item" section and image metadata
Packaging	Label on the back of pack or promotional leaflet
Audio/Video	Audible or visual label during AI-generated segments
Accessibility Content	Alt text or captions with disclosure and link to more info

While disclosure is conceptually familiar — marketers already operate within norms of transparency through disclaimers, sponsorship tags, and consumer labeling — this form of provenance introduces a new, AI-specific safeguard.

Provenance helps ensure credibility and provides valuable critical and historical data that is useful for both internal marketing organizations and future potential compliance requirements.

PROVENANCE

Since the early days of user-generated content (UGC), brands have balanced risk and opportunity. Concerns around misalignment with brand values, inconsistent creative quality, and potential reputational harm were real and persistent. Over time, however, many brands learned to harness UGC, amplifying aligned content, drawing inspiration from consumers, and partnering with influencers to guide creation within defined guardrails.

Generative AI fundamentally accelerates this dynamic, especially as text-to-video tools get closer to cinematic-quality production. The democratization of the video-based creative process has arrived. Anyone can create high-quality ad units — from six-second clips to full 60-second spots and quickly accelerating to longer form content. While brands gain access to scalable, efficient content creation, they also face a surge in unsanctioned, AI-generated brand content that can appear authentic to consumers. When misleading or unauthorized ads circulate, distinguishing between official and unofficial creative becomes increasingly difficult.

As a result, provenance and transparency are now critical. Brands must establish mechanisms that allow consumers, platforms, and regulators to clearly verify the origin, change-history, and authenticity of marketing content. Standardized provenance frameworks can help protect brand integrity, reduce risk, and reinforce trust in an environment where the line between authentic and synthetic content is rapidly blurring.

Technical Standards and C2PA

The **C2PA** is a marketing organization for **Content Credentials**, which represents an essential step toward enabling verifiable provenance in marketing content. The C2PA Steering Committee members include (at the time of publication) Adobe, Amazon, Microsoft, Google, Intel, Meta, OpenAI, Publicis Groupe, Sony, BBC, ARM, and TruePic. Content Credentials, which is the technology, is a human-readable “receipt” or overlay showing the provenance of a particular creative unit that is cryptographically embedded in a creative unit.

The C2PA’s tasks include:

- Identifying content provenance requirements
- Applying those requirements in the development of content provenance specifications
- Developing best practices and reference designs for applying these standards
- Promoting selected specifications to become global standards
- Promoting global adoption of digital provenance techniques
- Promoting the adoption of the digital Coalition’s specifications and standards by social and media platforms
- Ensuring content remains accessible even with digital provenance techniques applied

PROVENANCE

However, as some of these technological aspects remain machine-readable and less transparent to the typical viewer, significant challenges remain in translating these technical implementations into consistent, real-world practice. Chief among them is the “last mile” problem; even when metadata is embedded, its effectiveness is limited if distribution platforms strip it upon upload or fail to make it visible to end-users, or if humans, whether intentionally or unintentionally, delete it. While this paper does not prescribe a single technical path forward, the C2PA standards offer one possibility for simplifying complexity by enabling provenance through cryptography-secured tagging carried across creative platforms, such as Adobe Firefly, and integrated into platforms like Google’s advertising network.

Content Credentials and C2PA are still works in progress. They do not currently solve all ANA members' desired provenance goals. It is not a data rights management (DRM) system. It does not prevent copying. It does not stop disinformation, and it does not currently validate copyright or ownership.

However, given Content Credentials’ and C2PA’s stated intent and their broad acceptance across both technology and content platforms, we recommend that, for the immediate future, we support C2PA's efforts while continuing to help it evolve to incorporate additional provenance elements into the Content Credentials framework.

Global Interoperability

While C2PA’s platform can be applied globally, marketers should always check local market guidelines to ensure compliance with evolving legal frameworks such as the EU AI Act, which places specific obligations on transparency and the disclosure of AI-generated content. At the same time, the industry should remain attentive to emerging standards worldwide, both regulatory and technical, and commit to adopting practices as these evolve. By reviewing international frameworks, ANA members can strive to simplify compliance, safeguard brand reputation, and promote as much consistency as possible in consumer transparency across markets.

Operational Realities

Implementing provenance and disclosure practices is not solely a matter of technical standards; it is equally an operational challenge. To progress as an industry, those developing new strategies should consider both large global organizations and smaller teams with limited resources. Success depends on embedding these practices into everyday ways of working, integrating Content Credentials as developed by C2PA into workflows and digital asset management (DAM) systems, defining clear lines of responsibility across agencies, production partners, and platforms, and updating contracts to reflect shared accountability. In this sense, disclosure is not only about guidelines on paper, but about practical shifts in how marketing ecosystems operate. It is recognized that there is no perfect system, but the ways forward outlined in this document are currently accepted and members should aspire to implement these or similar methods that align with the principles.

ANA Advertising Provenance Table

While all fields are important, they are tiered in terms of priority to include if constrained.

C2PA Layer:

Hard Binding = cryptographic

Assertion = human-authored claim

Soft Metadata = referenced, not bound

Tier	#	Data Type/ Field Name	Description	Why Include (Rationale)	C2PA Layer
Tier 1	1	**Technical Metadata**	**File hash, watermark, C2PA tag, hash algorithm used, and timestamp of cryptographic signature**	Promotes authenticity and integrity by enabling verification that the asset has not been altered post-signing. The cryptographic foundation of the C2PA manifest.	**Hard Binding**
	2	**Marketer Name (Brand and Parent Company)**	**Advertiser identity**	Promotes transparency and proper attribution, enabling clear accountability in the provenance record.	**Assertion**
	3	**Campaign Name/ID**	**Campaign identifier**	Supports reporting, analytics, and audits by linking the asset to a specific marketing initiative.	**Soft Metadata**
	4	**Ad Format/ Media Type**	**Video, static image, audio, DOOH, interactive, etc.**	C2PA manifest structure and consumer disclosure format vary by media type. Capturing format enables format-appropriate implementation guidance and compliance mapping.	**Soft Metadata**
	5	**Ad Unit Length**	**For video or audio**	Used for tracking and is an operational spec.	**Soft Metadata**
	6	**Creative Ad Unit Name**	**Unique name or code for this creative asset**	Links this creative to its campaign, trafficking system, and asset library, ensuring consistent tracking across platforms.	**Soft Metadata**
	7	**Date Created**	**Production date**	Establishes the compliance and rights timeline, anchoring when the asset and its AI components were originated.	**Assertion**
	8	**Was AI Used? (Y/N)**	**Top-level AI disclosure flag**	Establishes the disclosure trigger that activates downstream AI-specific fields and satisfies platform, regulatory, and consumer transparency requirements.	**Assertion**
	9	**AI Tools Used (Provider, Model Version)**	**Specifies which AI software and model version was used**	Promotes transparency and reproducibility. Enables accurate attribution if the asset is later audited, disputed, or requires rights verification.	**Assertion**

PROVENANCE

Tier	#	Data Type/ Field Name	Description	Why Include (Rationale)	C2PA Layer
Tier 1	10	**Purpose of AI**	**Editing, synthetic content, digital twins, voiceover, etc.**	Categorizes the nature and scope of AI involvement – from minor editing to fully synthetic content – enabling proportionate disclosure and risk assessment.	**Assertion**
	11	**Date of Last Modification**	**Date of the most recent creative change**	Maintains provenance tracking and digital chain of custody, particularly critical for compliance in highly regulated industries.	**Assertion**
	12	**Revision History (High-Level)**	**Overview of major AI-related edits and creative changes**	Supports chain-of-custody by documenting the sequence of material changes, enabling reconstruction of how the final asset evolved.	**Assertion**
	13	**Ingredient Asset ID**	**The provenance of individual assets such as stock photos, music, logos, and AI-generated backgrounds**	Provides component-level compliance across every element of an ad. Enables efficient identification and replacement of low-resolution, low-quality, or rights-expired individual assets.	**Hard Binding**
	14	**Copyright Attribution**	**Final creative ownership**	Avoids rights disputes by documenting who owns the finished asset, articularly when AI generation, human creative work, and third-party assets are combined.	**Assertion**
	15	**Rights and Consent for Synthetic Actors/ Digital Twins**	**Likeness rights and contractual terms**	Required for legal protection where AI generates or replicates human likenesses. Ensures compliance with performer agreements, state likeness laws, and emerging federal standards.	**Assertion**
	16	**Intended Expiry Date**	**Stop date**	Triggers review or deactivation to ensure AI-generated content does not remain in distribution beyond its rights, compliance, or contextual validity window.	**Soft Metadata**
	17	**Compliance and Review**	**Legal reviewer name and date of legal approval**	Creates an auditable record of legal and compliance sign-off, establishing a defensible chain of accountability for regulated or high-risk advertising categories.	**Assertion**
Tier 2	18	**Platform Disclosure Require-ments**	**TikTok, Meta, YouTube, and other platform-specific AI labeling flags**	Ensures platform-specific AI labeling requirements are met prior to trafficking, preventing takedowns, delivery failures, and potential policy violations.	**Soft Metadata**

PROVENANCE

Tier	#	Data Type/ Field Name	Description	Why Include (Rationale)	C2PA Layer
Tier 2	19	**Human Oversight Confirmation**	**Y/N and reviewer name and role**	Documents that a qualified human reviewed and approved AI-assisted content prior to distribution, satisfying emerging regulatory requirements for human-in-the-loop accountability.	**Assertion**
	20	**Geofencing/ Distribution Region**	**Intended market(s)**	Enables jurisdiction-specific compliance mapping, as AI disclosure and synthetic content requirements vary significantly by region (EU AI Act, FTC, CPRA, LGPD) and platform market.	**Soft Metadata**
	21	**Synthetic Versus Captured Indicator**	**Fully synthetic, hybrid, or captured**	Aligns with emerging regulatory norms and platform labeling standards. Differentiates disclosure obligations based on the degree of AI generation versus human-captured content.	**Assertion**
	22	**Training Data Visibility (High-Level)**	**Discloses whether the AI tool's training data is known/auditable and whether proprietary brand assets were used in fine-tuning.**	Documents responsible data sourcing and surfaces potential copyright or data rights exposure tied to the model's training corpus.	**Assertion**
	23	**Visual Asset Rights/ Generative Origin**	**Stock, photographer, or generative source**	Supports proper licensing by identifying the origin of every visual element, distinguishing between licensed stock, commissioned photography, and AI-generated imagery.	**Assertion**
Tier 3	24	**Optional: Agency/ Production Vendor**	**Name of the agency or production company that created or trafficked the asset**	Closes the chain-of-custody gap between the marketer and the C2PA credential issuer. In most workflows the agency applies the C2PA manifest and signs the assertion — this must be documented.	**Assertion**
	25	**Optional: Consumer-Facing Disclosure Format**	**The label, badge, or notice displayed to consumers (e.g., "Created with AI," platform AI badge, watermark)**	Documents how AI use was disclosed at the consumer level, not just internally. As FTC guidance and state laws evolve, the format and placement of the disclosure itself may become a distinct compliance requirement.	**Soft Metadata**
	26	**Optional: NAICS Code**	**Industry classification**	Used for analytical purposes and not provenance specific.	**Soft Metadata**
	27	**Optional: Creative Personnel**	**CD, copywriter, AD, and producer**	Provides attribution and allows for future data sorting.	**Soft Metadata**

Revised April 2026

PROVENANCE

Provenance provides useful safeguards — ensuring that when content is questioned, brands can authenticate their creative assets and defend their integrity with confidence.

Recent moves by major players illustrate why provenance is becoming increasingly important. Canon has piloted cameras that cryptographically sign images at the moment of capture, preserving a verifiable record throughout the editorial process. In contrast, Google's latest Pixel devices now embed C2PA Content Credentials directly into photos and edits, creating tamper-evident metadata that travels with the file. These efforts demonstrate how provenance can evolve beyond one-time disclosure to become a standard layer of accountability, safeguarding brands from deepfakes, unauthorized manipulation, and reputational harm by helping ensure content integrity throughout the content's entire lifecycle.

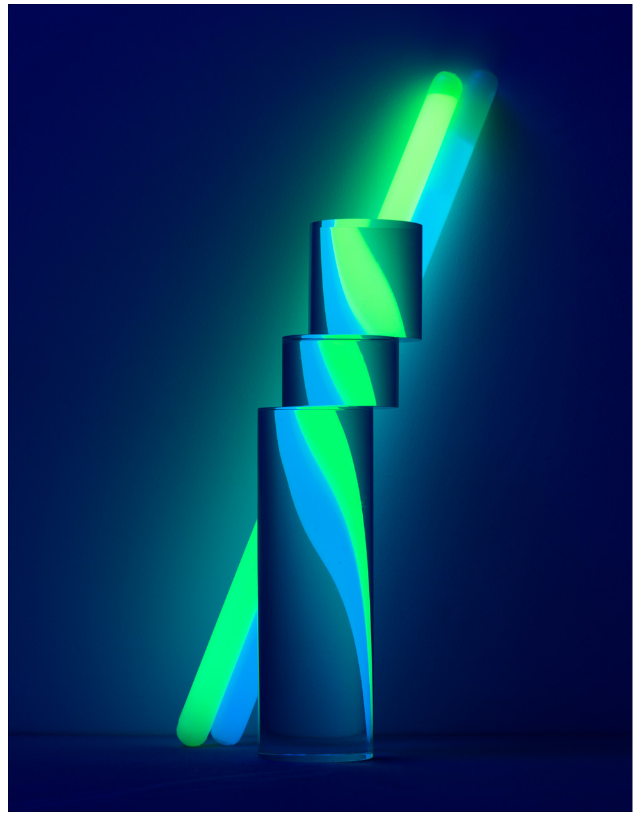

THE CASE FOR SELF-REGULATION

While comprehensive federal regulation of AI in advertising has yet to take shape, state legislatures are already moving to address transparency. For example, in New York, **Senate Bill S8420A**, known as the NY Synthetic Performers Act, requires advertisers to disclose the use of synthetic performers (image or voice), with escalating penalties for repeat violations. In California, a 2024 law (Assembly Bill 2355) requires political advertising to disclose when content has been generated or altered using AI, with financial penalties of up to $5,000 per violation for non-compliance. The disclosure must be clear and conspicuous, stating, “This ad was generated or substantially altered using artificial intelligence.” The state privacy landscape is likely to continue to rapidly evolve in this area, making it important for members to consider ways to stay on top of a shifting AI legal and regulatory landscape.

These state-level actions reflect a broader global trend. The **EU’s AI Act** introduces transparency requirements across high-risk AI applications, while **China’s Deep Synthesis Provisions** impose strict rules on synthetic media labeling and disclosure. In Australia, industry guidelines emphasize the importance of disclosure where omission may mislead, the need for human oversight, safeguards against bias, and respect for privacy and intellectual property rights. These guidelines illustrate how markets outside the EU and China are also formalizing expectations around responsible AI use in marketing. Together, these efforts underscore the inevitability of regulatory intervention and the risks of waiting for fragmented mandates to dictate practice.

For marketers, the implication is clear: the regulatory landscape is evolving rapidly, but unevenly. A voluntary, industry-led framework offers a path to consistency across jurisdictions while reinforcing trust and brand integrity. By adopting common principles now, ANA members can avoid the reputational and operational risks of piecemeal compliance and lead with a proactive stance that balances innovation with accountability.

These guidelines provide aspirational principles and practical guardrails that support industry self-regulation. By aligning well-established laws with evolving industry requirements and tenets, such as those from the Content Authenticity Initiative and C2PA, this initiative aims to foster trust, protect brand integrity, and help ensure that consumers are informed and empowered in their interactions with AI-enhanced content. The intended outcome is a practical, adaptable framework that the working group recommends ANA’s member organizations adopt to lead responsibly as the industry transforms and society adapts to AI.

Against this backdrop, the working group has developed consensus-driven recommendations that establish clear, consistent, and ethical principles for the disclosure and provenance of generative AI, providing actionable, industry-aligned guidance to ensure consumer trust today and protect brands for the future.

SUMMARY

The pace of generative AI's evolution is remarkable. Only three years ago, our industry was wowed by the idea of having a human-like conversation with a computer, and now we are talking about viable marketing creative content that comes close and soon will likely reach near-parity with some of the greatest marketing campaigns/creatives the industry has produced. Managing marketers' processes and consumer experience through this transition is a key challenge, and marketers' traditional pace of policy, workflow, and organizational structure changes must adapt faster than ever. Yet ensuring strong communication and trust between the marketer and the consumer must be at the center of this transition.

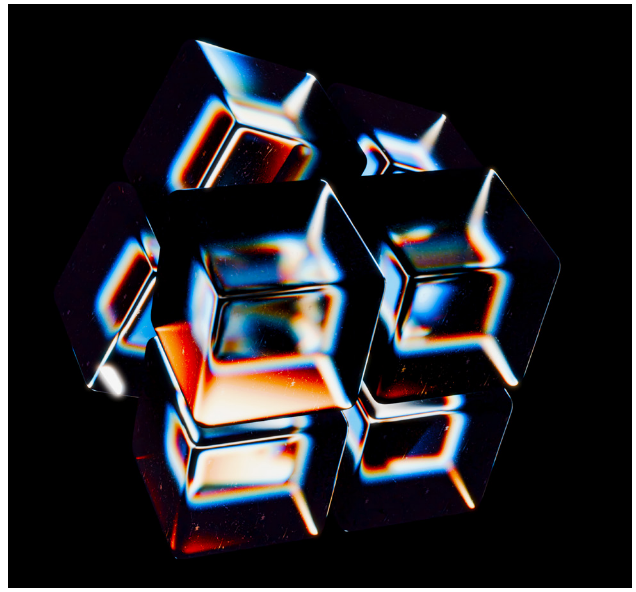

The AI Working Group believes that embracing the C2PA and Content Credentials products as one such approach is a logical next step. C2PA has already laid a strong foundation for transparency and provenance and has gained acceptance from key market players. With the ANA and its members, we believe we can work with C2PA and other industry organizations to achieve the level of self-regulation and provenance standards outlined in this paper.

ACKNOWLEDGEMENTS

The ANA appreciates the time, careful thought, debate, and foresight of the working group leaders, who spent countless hours above and beyond their already full plates to bring this paper to fruition. Their diverse perspectives and collaborative spirit were instrumental in shaping the recommendations and insights presented herein.

- Michael Therrien: Director, AI Strategy, Global Marketing, Kyndryl
- Jacqui Stephenson: Global Responsible & Purpose Marketing Officer, Mars
- Michael Ditter: Director, Artificial Intelligence Strategy and Emerging Technology, Diageo
- Rinita Datta: Director Product Marketing, Cisco
- Simon Duncan-Watt: Senior Director, Content and Creative, CARE

And thanks to members of our ANA AI Leadership Council from the following companies:

- **Adobe**
- **Anheuser-Busch**
- **Barclays**
- **Best Buy**
- **CBRE**
- **Cigna**
- **Constellation Brands**
- **Diageo**
- **Disney Experiences**
- **Esquire Bank**
- **General Motors**
- **Heineken**
- **HP**
- **Humana**
- **Huntington Bank**
- **IBM**
- **Inspire Brands**
- **Kellanova**
- **Kyndryl**
- **L'Oréal**
- **Lowe's**
- **Mars**
- **Mastercard**
- **McDonald's**
- **Medtronic**
- **Molson Coors**
- **Mondelēz International**
- **Nestlé**
- **Novartis**
- **P&G**
- **Paychex**
- **PepsiCo**
- **Qualcomm**
- **Sanofi**
- **The Clorox Company**
- **The Hershey Company**
- **Verizon**
- **Visit Orlando**
- **Wells Fargo**

ANA Contributors:

- **Michael Donnelly, EVP, Technology Practice**
- **Jeff Minsky, VP, AI and Marketing Futures, Technology Practice**
- **Melanie Kania, Senior Manager, Marketing and Communications**
- **Gabrielle White, Manager, Content**

APPENDIX

Links to AI Disclosure Policies of Platforms

- YouTube
- TikTok
- Meta
- Pinterest

Links to Relevant Government or Trade Industry Body Documents

- IAB: AI Transparency and Disclosure Framework
- 4As: The 4As Guide to Content Provenance
- WFA Responsible AI Principles
- European Union's AI Act

Other Transparency Efforts in the Marketplace

- **Google SynthID**: While not specific to advertising or marketing, SynthID is a new watermarking tool explicitly designed for AI-generated content. It empowers users to identify AI-generated (or altered) content. It embeds digital watermarks directly into AI-generated images, audio, text, or video.
- **Media Ratings Council AI Initiative**: Slated for the release of comprehensive AI-focused standards in 2026, this initiative focuses on six primary uses of AI (generative AI, agentic AI, ad verification and brand safety, AI-powered search engines, invalid traffic, and large datasets, and identity resolution in TV measurement). As it relates to our current effort, the generative AI topics that they will address include: principles around measurement of content/data ownership, permissibility and attribution, reliance on inaccurate or incomplete data and data quality issues.

ABOUT THE ANA

The Association of National Advertisers (ANA) is the definitive voice of the marketing industry. Since 1910, we have set and advanced the agenda for marketing transformation, connecting over 1,600 member companies to an influential global network, insights, and resources that drive growth. Our members represent 20,000 brands and $400 billion in annual marketing investment. Through industry-leading research, the CMO Growth Council, and our proprietary Growth Agenda and Practices, the ANA empowers marketers to shape the future of marketing and create lasting impact for their organizations and the industry.

TECHNOLOGY

AI Transparency and Provenance Principles for Marketers

JUNE 2026

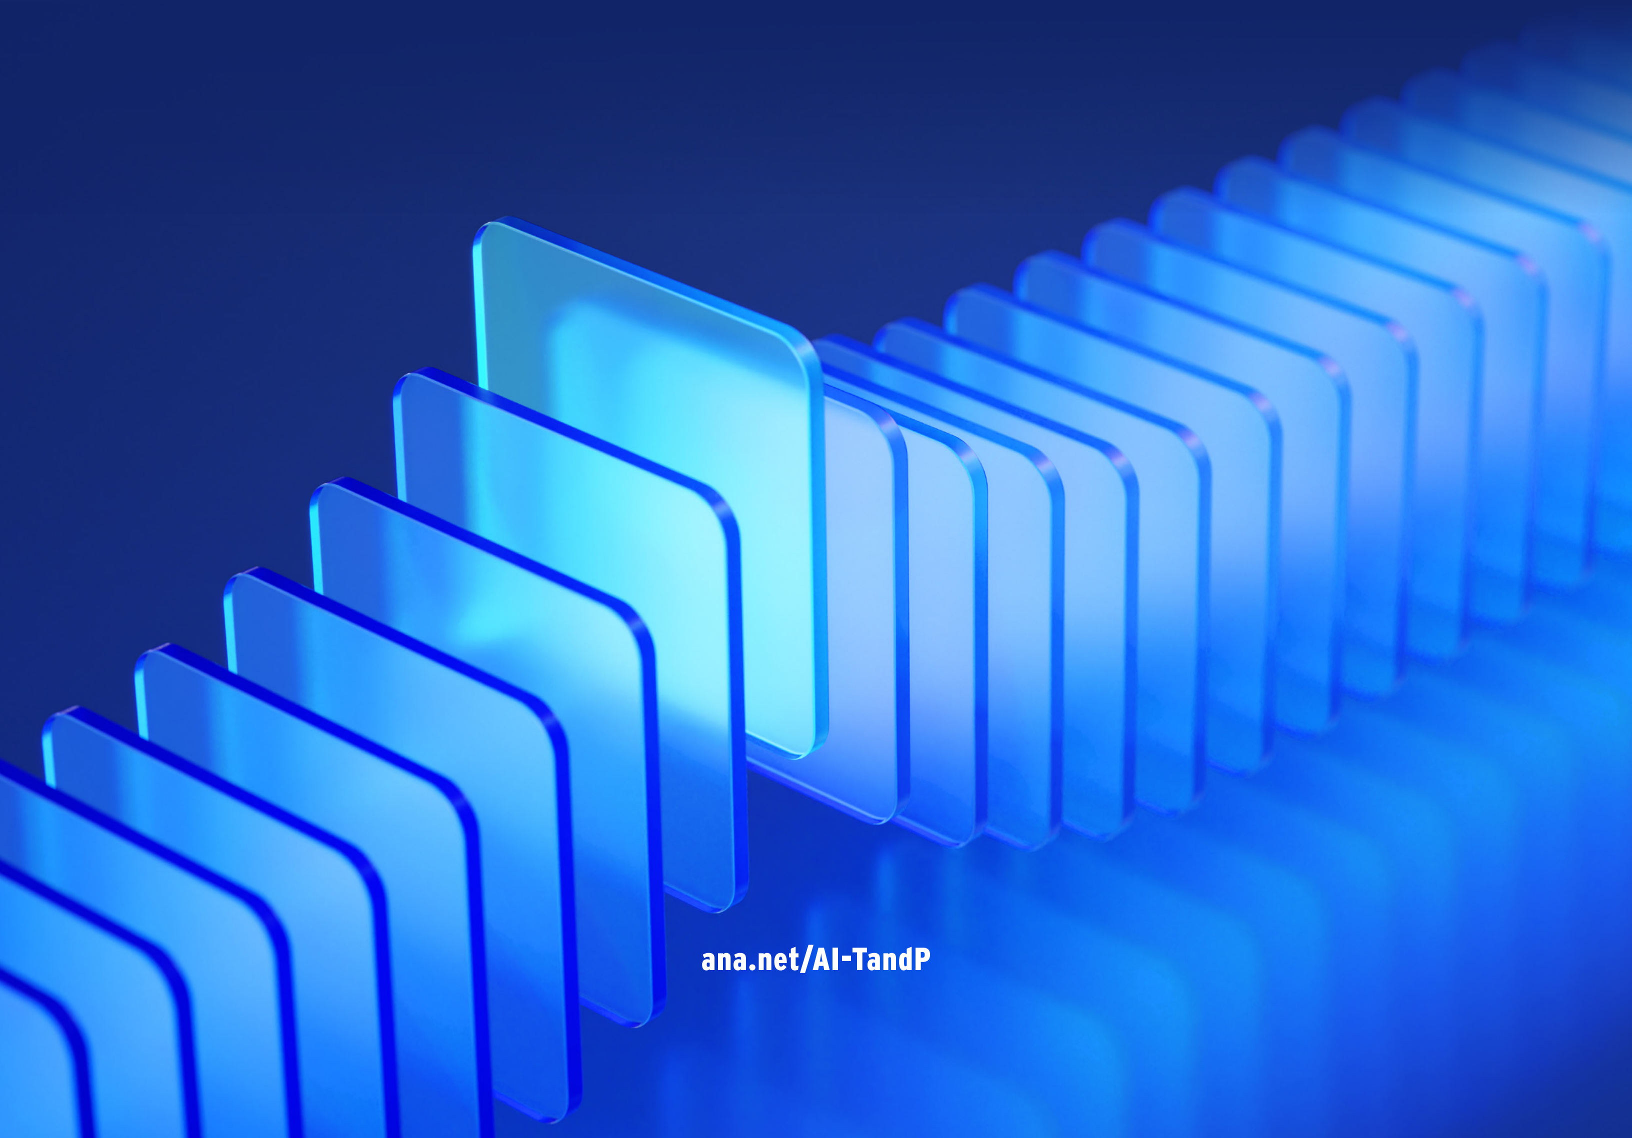

ana.net/AI-TandP