

Resemble AI Expands Platform with AI Watermarking for Audio, Video, Image, and Text

Resemble Watermarker now embeds imperceptible, tamper-resistant marks across every content modality - giving AI output a verifiable chain of custody

MOUNTAIN VIEW, CA, UNITED STATES, June 24, 2026 /EINPresswire.com/ -- Resemble AI, the generative AI security platform, today expanded [Resemble Watermarker](#) from audio-only watermarking to full multimodal coverage, adding image, video, and text. The expansion is powered by [PerTh Multimodal](#), a ground-up rebuild of the model behind the company's existing audio watermarker, which is also [available as open source](#). Now with PerTh Multimodal, enterprises and creators can include imperceptible, tamper-resistant marks across every

content type they generate. The update arrives as synthetic media threats reach an inflection point and organizations face mounting pressure to prove the origin and authenticity of the content they create and distribute. In a June 2026 report, Gartner named deepfakes one of four critical threats where enterprise defenses are most overmatched, the cases where the attacker now holds the advantage.

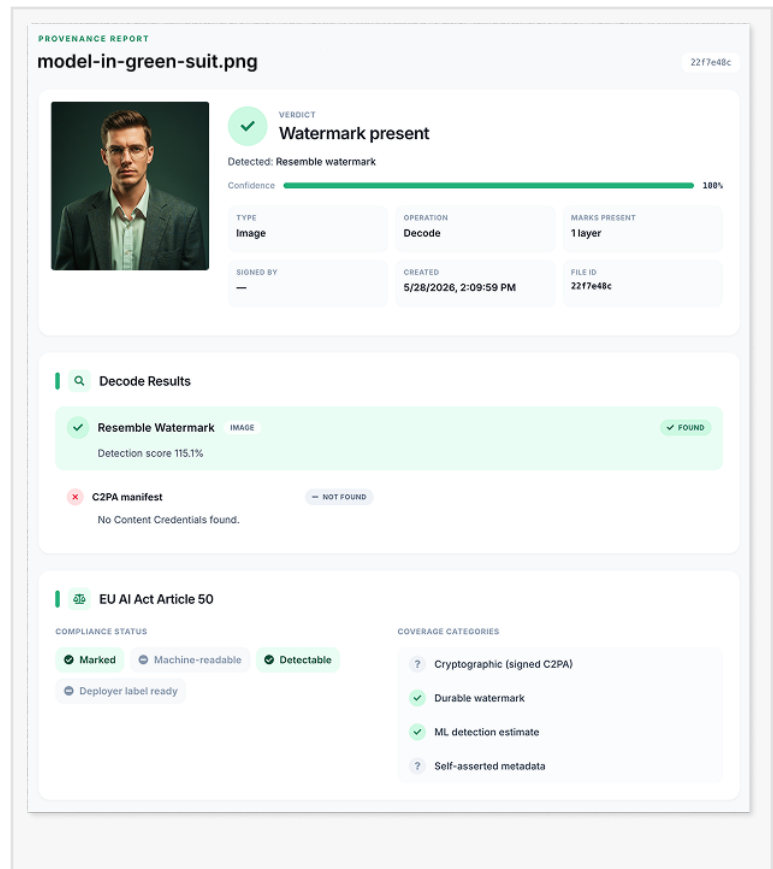
Impersonation and fraud are now a core business risk

The risk is what bad actors do with synthetic media: deepfaking your assets and impersonating your brand, your executives, and your talent to commit fraud and damage reputation. As found in The 2025 Resemble AI Deepfake Threat Report, brand and reputation attacks generated 156.3 billion impressions last year, more than all other attack categories combined. Across every category, synthetic media drove 296.4 billion impressions, the equivalent of roughly 53 for every internet user on earth, which means the damage a fraudster can do is far-reaching and as fast as



their access to social media.

What makes it hard to fight is a combination of proof of ownership and the ability to remove the unauthorized content: when a fake of your brand or an executive spreads, most organizations have no way to show which content is authentically theirs and which is not. Research presented at ACM in 2025 found that only about 38% of AI content generators watermark their outputs adequately, so most AI-generated assets carry no verifiable chain of custody. While deepfake detection catches the fakes other people make, watermarking proves the authenticity of the content you make, so a genuine asset can be told apart from an impersonation, and the question of who made something, and whether it has changed, is answerable by the content itself.



Resemble Watermarker: the authentication layer for AI-generated content

Resemble Watermarker is a declaration of ownership placed inside the content itself. It proved the approach in audio first, and this update extends the same imperceptible, in-signal mark to image, video, and text, so one product now covers every content type an enterprise generates. Where metadata-based approaches can be stripped the moment a file is re-uploaded, compressed, or re-encoded, Resemble Watermarker embeds the mark directly into the content.

A properly watermarked asset can answer four questions at any point in its lifecycle:

- Who issued this content?
- Was it AI-generated, AI-altered, or authentic?
- Has it been tampered with since creation?
- Are other provenance marks present, including SynthID or C2PA credentials?

Resemble Watermarker reads SynthID and C2PA markers, and it writes and signs C2PA credentials, so it delivers both layers regulators are converging on: a provenance manifest any compliant tool can read, and an in-signal watermark that survives when that metadata is stripped. Assets can also carry an explicit attribution mark that encodes issuer identity, content date, and authenticity status. Support for on-device, VPC, and air-gapped deployment lets enterprises run it inside their own environments.

“A watermark is a declaration of ownership placed inside the content, ideally at the moment of

creation, so the question of who made something and whether it changed can be answered by the file itself, not the channel or handle it came from," said Zohaib Ahmed, CEO and Co-Founder of Resemble AI. "Deepfake detection answers whether content is AI-generated or altered, a watermark proves who it belongs to, and that is a foundational authentication layer every organization needs."

Built for the regulation that is arriving, in the EU and beyond Resemble Watermarker is designed to map directly to the obligations enterprises face as AI content rules move from rule-making to enforcement. The EU AI Act's Article 50 transparency obligations take effect August 2, 2026, requiring AI-generated and manipulated content to carry machine-readable marking, with penalties up to 15 million euros or 3% of worldwide annual turnover.

The EU is not moving alone. China's mandatory labeling rules took effect September 1, 2025, requiring all AI-generated content to carry both a visible label and an embedded watermark or metadata identifier, the same two-layer pattern the EU is now adopting. More jurisdictions are expected to follow as the approach becomes the global default. Resemble Watermarker produces the machine-readable marking these rules call for across all four modalities through a single API, giving compliance, legal, and trust and safety teams an authentication layer that travels with the content itself.

Availability

Resemble Watermarker powered by PerTh Multimodal is available today through a self-serve API at resemble.ai, and audio watermarking remains available as open source. Enterprise licensing, including VPC and on-premises deployment, is available through Resemble AI's enterprise sales team. Resemble Watermarker integrates with Resemble AI's broader generative AI security platform, including Resemble Detect for deepfake detection.

About Resemble AI

Resemble AI is generative AI security that gives organizations the tools to verify, detect, and act on synthetic media attacks and fraud attempts — anywhere they appear. Built on our foundational models, Resemble provides a layered trust stack that turns model output into decisions companies can understand, audit, and defend. Founded in 2019, Resemble AI is trusted by global enterprises and government agencies to secure the full lifecycle of synthetic media across audio, video, image, and text. Learn more at resemble.ai.

Clint Bagley

Resemble AI

clint.bagley@resemble.ai

This press release can be viewed online at: <https://www.einpresswire.com/article/921719977>

EIN Presswire's priority is source transparency. We do not allow opaque clients, and our editors try to be careful about weeding out false and misleading content. As a user, if you see something

we have missed, please do bring it to our attention. Your help is welcome. EIN Presswire, Everyone's Internet News Presswire™, tries to define some of the boundaries that are reasonable in today's world. Please see our Editorial Guidelines for more information.

© 1995-2026 Newsmatics Inc. All Right Reserved.