

The TAG logo consists of the letters "TAG" in a bold, white, sans-serif font, centered within a solid dark blue rectangular box.

TAG

ANALYST REPORT

COMPETITIVE ASSESSMENT OF COMMERCIAL VENDORS ADDRESSING DEEPFAKES

THE TAG ANALYSTS

LEAD ANALYST: DR. EDWARD AMOROSO
CHIEF EXECUTIVE OFFICER, TAG,
RESEARCH PROFESSOR, NYU

GetReal.



COMPETITIVE ASSESSMENT OF COMMERCIAL VENDORS ADDRESSING DEEPFAKES

THE TAG ANALYSTS

LEAD ANALYST: DR. EDWARD AMOROSO,
CEO, TAG, RESEARCH PROFESSOR, NYU

This TAG report begins with an up to date summary of the status of the commercial market for security solutions targeting the risk of synthetic media fraud, including the malicious use of deepfakes. The report then focuses on a competitive assessment of leading vendors based on a set of comparison criteria. As will be shown several excellent commercial options exist which fare well against our set of criteria.

INTRODUCTION

As readers of this report already understand, synthetic media fraud has moved quickly from novelty to operational risk.¹ The use of deepfake video, cloned voice, AI-generated identity documents, manipulated images, and real-time impersonation has created a new class of digital trust problem for enterprise security, fraud, compliance, legal, human resources, and executive protection teams.

Security teams now recognize that the core challenge is no longer whether an attacker can create convincing synthetic media. That question has been answered. Instead, the practical issue is whether an enterprise can detect, authenticate, verify, and respond to media-based deception before it enables a fraudulent transaction, unauthorized access decision, executive impersonation, hiring fraud event, customer scam, or other reputational crisis.

¹ Note our use of "synthetic media fraud" in the title of this report and in the discussion. While this reference is not perfect, and is certainly less familiar than the more common use of "deepfake," we believe it captures the essence of the actual risk. As will be shown, this manifests in fraud scenarios such as identity impersonation during live video, which is an increasingly nagging issue.

At TAG, we have come to view synthetic media fraud as more of an identity assurance problem than a media (or political) problem. Traditional controls such as passwords, MFA, device fingerprinting, liveness checks, and call-backs remain useful, but they were never designed for a world in which the person speaking on a video call, leaving a voice message, submitting a selfie, or participating in a remote interview might be digitally generated.



Figure 1. Liveness Evidence of a Deepfake Video Image²

In addition, our TAG customers report that this risk is especially concerning in financial services, business process outsourcing, remote hiring, customer support, call centers, government services, law enforcement, public communications, and executive decision workflows. The enterprise challenge is thus to determine whether a digital person, digital voice, or digital artifact should be trusted at the point of interaction.

This report is intended to help enterprise security teams, as well as adjacent executives and teams in legal departments, public relations, corporate communications, and information technology (IT), better understand both the problem as well as the specific commercial vendor options that are available. We also introduce and use a set of comparison criteria to call out the commercial vendors which we found to be especially strong.

MOST COMMON METHODS FOR ADDRESSING THE RISK

We have spent considerable time at TAG reviewing how commercial vendors approach the challenge of deepfakes, and while the technology continues to advance, there are some common themes that have emerged. In this section, we provide a glimpse into the major methods being deployed and used by the best available vendors in this important category of risk management.

1. Direct Inspection – The first method involves the inspection of images, audio, video, and documents for artifacts that suggest manipulation. In video, this can include facial movement inconsistencies, lip-sync mismatches, lighting anomalies, frame discontinuities, eye reflection irregularities, and biological inconsistencies. In audio, this can include unnatural cadence, spectral artifacts, breathing irregularities, and signs of voice cloning.

² Example taken from <https://fortune.com/2022/09/03/live-deepfakes-detect-methods-zoom-fraud/>.




Figure 2. Spotting Irregularities in Eye Images in Deepfakes³

- 2. Real-Time Detection** – The second method involves evaluation of synthetic indicators while a video call, interview, meeting, or customer interaction is occurring. Real-time detection is important because many deepfake-enabled fraud events occur interactively. A fake candidate in a remote interview, a synthetic executive requesting a payment, or a cloned customer voice authorizing account changes must be detected while the event is still actionable.
- 3. Identity Verification** – The third method involves continuously determining whether the person in a workflow is who they claim to be. The best implementations combine face, voice, device, behavioral, document, and contextual signals. This is an important distinction because media detection alone might say that a video looks suspicious, but identity verification can help determine whether the individual is trustworthy in the context of a transaction or session.
- 4. Provenance Capture** – The fourth method focuses on where media came from, how it was captured, whether it has been edited, and whether cryptographic credentials can be used to validate authenticity. This is promising for journalism, public sector and media evidence, insurance claims, and field verification. Its success relies on the capture device, editing tools, platforms, and relying parties supporting the standard.⁴
- 5. Workflow Integration** – The fifth method is driven by the need for synthetic media protection to become embedded into communications platforms, fraud systems, HR systems, investigation tools, identity verification workflows, and SOC or case management processes. A detector that works only as a standalone upload portal is useful for analysts, but it will not be sufficient for high-volume enterprise protection.

VENDOR LANDSCAPE

Our main focus at TAG since 2016 has been to review, assess, rate, and compare commercial security vendors to assist enterprise teams with source selection and security architecture implementation. To that end, we have expended considerable effort in the area of deepfake detection and synthetic media fraud mitigation. The list below provides an overview of the key vendors we've review – ordered alphabetically.

Clarity

Founded 2022 | HQ: Tel Aviv

This seed funded startup is focused on protecting hiring and interview workflows from AI-enabled deception, including deepfake candidates, AI cheating, and identity fakes. Clarity focuses on so-called vertical specialization, which can be powerful when the use case is well-defined. The hiring process is a major area of risk because attackers can use synthetic media, proxy interviewers, AI-generated answers, forged credentials, and identity manipulation to gain employment, contractor access, or insider placement.

³ See <https://physicsworld.com/a/the-eyes-have-it-how-to-spot-the-difference-between-a-deepfake-portrait-and-a-real-picture/> for a description of how this is done.

⁴ The standard supporting this description of media authenticity and tracking is C2PA (Coalition for Content Provenance and Authenticity). C2PA is an open technical standard that creates a tamper-evident chain of custody from content creation through distribution. See <https://c2pa.org/>.

Clarity's strength is therefore its high level of focus on specific HR and recruiting workflows rather than more generic media detection. Its limitation is that this approach to specialization may make it less suitable as an enterprise-wide synthetic media fraud platform across executive communications, customer fraud, public communications, and investigative use cases. Time will tell whether the company ever decides to expand its focus.

GetReal Security

Founded 2022 | HQ: Austin, Texas

GetReal is a trust and authenticity platform focused on deepfake detection, impersonation detection, and continuous identity verification across voice and video interactions. The company aligns with the enterprise shift from simple detection to identity-centered trust. GetReal is not merely asking whether a clip or image is synthetic. It allows organizations to continuously assess whether the person or media in a sensitive interaction should be trusted. This approach is important for executive impersonation, hiring fraud, help-desk account-recovery attacks, financial authorization, customer interaction, and high-risk communications. GetReal's strongest attributes include multimodal coverage across face, voice, and session; real-time orientation, continuous rather than point-in-time identity-centric framing, and executive-level relevance. All of these advanced media forensics capabilities stem from groundbreaking academic research from its iconic founder, Dr. Hany Farid.⁵

Reality Defender

Founded in 2021 | HQ: New York City

This startup places emphasis on detecting AI-generated and manipulated media across audio, image, and video through API, SDK, and enterprise workflow integration. TAG sees Reality Defender as one of the more credible detection-first providers because it is developer-friendly, enterprise-oriented, and built for multiple media types. Its strength is broad technical detection and ease of embedding into applications.

It is well-suited for organizations that need to add synthetic media detection into fraud, platform, or content workflows. The main limitation is that a detection engine, even a strong one, must still be connected to identity assurance, case disposition, and business decision logic. Reality Defender should therefore be viewed as a powerful detection layer that may require additional enterprise controls around identity, workflow, and response.

Resemble AI

Founded in 2019 | HQ: Santa Clara

This startup is notable because it combines generative voice capabilities with detection and security controls. This dual role gives the company useful insight into the mechanics of synthetic voice and media generation. TAG views Resemble AI as particularly relevant for audio and voice-oriented deepfake risk, including cloned voice fraud, synthetic customer calls, fake interviews, and manipulated spoken content.

Its strengths include voice expertise, multimodal detection claims, and an understanding of responsible generation and verification. Its limitation is that enterprise anti-fraud buyers will need to distinguish between creative synthetic voice governance and full-spectrum identity spoofing prevention. Resemble AI is strong where voice and generated media security are central, but it may need to be paired with identity, transaction, or meeting security controls.

⁵ Our team was familiar with the work of Dr. Hany Farid from UC-Berkeley (see <https://www.getrealsecurity.com/teams/execs/dr-hany-farid>), where he had established a reputation in the computer science community as a leader in dealing with media fraud. Since then, TAG has worked with both Ballistic Ventures, which funds GetReal, as well as GetReal Security itself on a variety of different projects including presentations at events. That said, we do our best to remain objective, and we present (we hope) all vendors in this report in a fair and unbiased manner.

Sensity AI

Founded in 2018 | HQ: Amsterdam

This company has long been associated with deepfake detection and media intelligence, and it is relevant for organizations that require forensic review of manipulated media, threat intelligence around deepfake campaigns, and analysis across image, video, audio, and identity-related content. TAG views Sensity as technically credible and useful for teams that need investigative depth and broad coverage of synthetic media risk.

Its strengths include media analysis orientation, research depth, and applicability to fraud, law enforcement, intelligence, and trust and safety contexts. Its relative weakness is that the enterprise market is moving toward integrated real-time identity trust, and buyers will need to evaluate how well Sensity fits into live collaboration, call center, HR, and transactional workflows.

Sentinel AI

Founded in 2019 | HQ: Estonia

This company represents the class of vendors focused on detecting manipulated media and supporting democratic, media, public sector, and trust-oriented use cases. TAG views Sentinel as historically relevant to the deepfake detection market, particularly in the context of AI-based detection of forged video and synthetic media.

Its strengths include early market focus on manipulated media, public-sector relevance, and conceptual fit with misinformation and influence protection. Its weakness is that current enterprise buyers will need to validate product maturity, commercial momentum, integration support, and whether the offering is sufficiently current for real-time fraud, identity, and collaboration use cases.

Truepic

Founded in 2019 | HQ: San Diego

This team occupies a different but important position in the market. Rather than focusing primarily on detecting deepfakes after the fact, Truepic emphasizes content authenticity, capture integrity, and provenance. This makes Truepic especially relevant where the goal is to prove that a photo or video was captured authentically at a known time, on a known device, and without unauthorized manipulation.

TAG views this as highly valuable for insurance, field inspection, journalism, public sector reporting, investigations, compliance documentation, and other evidence-oriented workflows. Truepic's strength is trust at capture and provenance-based assurance. Its weakness, relative to deepfake fraud prevention, is that provenance does not solve all adversarial scenarios. If media enters from uncontrolled sources, then provenance alone may not be enough.

PROPOSED COMPARISON CRITERIA

TAG recommends that enterprise teams compare synthetic media fraud prevention vendors using a set of criteria that reflect business risk rather than laboratory accuracy alone. The criteria requirements listed below were created by TAG based on our review and assessment of what is available commercially. We try to do our best to point out the vendors from our list above who are particularly well-suited to that requirement.

1. Performance (Top Performer: GetReal) – The first criterion is real-time performance, including latency, scalability, and effectiveness during live meetings, calls, and customer sessions. For this requirement, we especially like GetReal because it was built for live interaction protection, executive communications, and continuous trust during sessions.

- 2. Assurance (Top Performer: GetReal)** – The second criterion is identity assurance depth, meaning whether the product can connect media authenticity to whether the individual is who they claim to be. In our view, GetReal best addresses this requirements given their focus on determining whether a participant is genuinely who they claim to be.
- 3. Breadth (Top Performer: Reality Defender)** – The third criterion is modality breadth, meaning whether the platform can analyze video, audio, image, document, and live interaction signals. Deepfake fraud is increasingly multimodal, and Reality Defender analyzes multiple modalities and detecting deepfakes across enterprise workflows.
- 4. Governance (Top Performer: GetReal)** – The fourth criterion is governance readiness, including auditability, privacy controls, deployment options, false positive management, human review, and regulatory defensibility. GetReal has the strongest alignment with enterprise trust, risk management, review processes, and operational governance.
- 5. Explainability (Top Performer: Sensity AI)** – The fifth criterion is forensic explainability, including whether the product can show why something was flagged and whether analysts can defend the result. Sensity AI is strong in its emphasis on investigation, forensic review, and analyst-driven examination of manipulated content.
- 6. Integration (Top Performer: Reality Defender)** – The sixth criterion is enterprise integration. Buyers should examine APIs, SDKs, browser extensions, video conferencing support, call center integration, identity platform integration, and policy automation. Reality Defender focuses on embedding detection into apps, workflows, and platforms.
- 7. Provenance (Top Performer: Truepic)** – The seventh criterion is provenance and authenticity support, including use of cryptographic content credentials, trusted capture, chain-of-custody, and standards-based verification. Truepic is the clear commercial leader in trusted capture, content credentials, and chain-of-custody verification.
- 8. Robustness (Top Performer: GetReal)** – The eighth criterion is adversarial robustness, meaning the vendor's ability to keep pace with new generative models, real-time face swaps, voice cloning, and model evasion. GetReal is best positioned around evolving enterprise fraud scenarios, including impersonation, deepfakes, and identity attacks.
- 9. Specialization (Top Performer: Clarity)** – The ninth criterion is workflow specialization, because a product optimized for hiring fraud might outperform a generic detector, while a provenance vendor might outperform a detector in evidence capture. Clarity has the most specialized offering, focused on hiring fraud and synthetic candidate detection.
- 10. Maturity (Top Performer: GetReal)** – The tenth criterion is operational maturity, including customer support, deployment flexibility, accuracy, and ability to support enterprise-scale programs. GetReal appears best aligned with large-scale enterprise deployment, executive protection use cases, and operational support requirements.

TAG ASSESSMENT

Our initial assessment is that no vendor in this market should be viewed as a complete solution to synthetic media fraud in every context. The category is dynamic, the attack techniques are evolving quickly, and enterprise workflows differ across financial services, government, hiring, media, customer support, and executive protection. This implies a moving target when trying to establish clear comparisons, so readers should keep that point in mind.

That said, we do believe that vendors can be differentiated by how well they support the most important dimensions included in the set of criteria introduced above. When we make such criteria-based assessments (see Figure 3), we are obliged to make somewhat subjective determinations, but we do based on available evidence, and discussions with vendor leadership teams and their customers.

Performance	GetReal	Integration	 Reality Defender
Assurance	GetReal	Provenance	 truepic
Breadth	 Reality Defender	Robustness	GetReal
Governance	GetReal	Specialization	 Clarity
Explainability	 sensity	Maturity	GetReal

Figure 3. Summary of Criteria Assessment of Deepfake Security Vendors

Against these criteria, it turns out that GetReal Security appears to score highest in our assessment, presumably because its approach seems most aligned with how we view this market as evolving. Our view is that the core enterprise problem is not simply fake media detection, but instead the establishment of continuous trust in the various type of human digital interactions that occur on a day-to-day basis.

FINAL NOTE

TAG Research as a Service (RaaS) customers are always encouraged to reach out to TAG analysts via their portal accounts for more information or insights on any aspect of cybersecurity and artificial intelligence. We look forward to hearing from you.

ABOUT TAG

Recognized by Fast Company, TAG is a trusted next generation research and advisory company that utilizes an AI-powered SaaS platform to deliver on-demand insights, guidance, and recommendations to enterprise teams, government agencies, and commercial vendors in cybersecurity and artificial intelligence.

Copyright © 2026 TAG Infosphere, Inc. This report may not be reproduced, distributed, or shared without TAG Infosphere’s written permission. The material in this report is comprised of the opinions of the TAG Infosphere analysts and is not to be interpreted as consisting of factual assertions. All warranties regarding the correctness, usefulness, accuracy, or completeness of this report are disclaimed herein.