


The Synthetic Spokesperson Governance Gap

When a brand can recreate a person, consent is only the first control



In this briefing

When a brand can recreate a person, consent is only the first control

·	Current Signal
·	Executive Brief
01	What Changed
02	Why This Is a Communications Governance Problem
03	The Artefact91 Synthetic Identity Control Stack
04	The Synthetic Identity Risk Matrix
05	Consent Is Not the Same as Audience Understanding
06	Disclosure Is Becoming a Governance Standard
07	The Rights Landscape Is Fragmented
08	Build an Identity Rights Ledger
09	Posthumous Use Needs an Extra Test
10	The Spokesperson Contract Has Changed
11	When Synthetic Identity Becomes a Crisis
12	A 30-Day Readiness Sprint
13	The Strategic Question
·	Artefact91 Point of View
·	Sources & Research Base
·	Research Note

HOW TO READ IT

A briefing, not a report: the executive brief carries the argument, the numbered sections carry the evidence and the method, and every external figure is attributed to its source.

ARTEFACT91 BRIEFING

Current Signal

On 2 October 2026, Reuters reported that Swedish manufacturer SKF had used AI to recreate Greta Garbo in a new advertisement, developed with the Garbo estate and family. Three days later, Reuters reported that Italian Prime Minister Giorgia Meloni had applied to register a short clip of her voice as an EU trademark as an additional response to AI-enabled misuse. The EU AI Act's Article 50 transparency duties have applied since 2 August 2026. [1][2][4]

ARTEFACT91 POSITION

SYNTHETIC IDENTITY IS BECOMING A COMMUNICATIONS ASSET THAT NEEDS ITS OWN GOVERNANCE ARCHITECTURE.

Permission matters, but so do scope, audience understanding, provenance, disclosure, approval and retirement.

ARTEFACT91 BRIEFING

Executive Brief

Generative media has lowered the production barrier between a real person and a commercially usable simulation of that person. A spokesperson, executive, performer, founder, subject-matter expert or other recognizable figure can now be recreated in new contexts without making a new physical appearance or recording every execution.

That creates obvious benefits: localisation, speed, continuity, creative flexibility and lower production friction. It also creates a new category of reputation risk. A synthetic use can be contractually authorised and still mislead audiences, exceed the approved scope, survive beyond the relationship that enabled it, or attach a person's identity to claims they never personally reviewed.

This briefing is distinct from Artefact91 Briefing 22, which addressed unauthorized executive impersonation and authentication. Briefing 25 examines the opposite side of the identity problem: what happens when an organization itself commissions, approves or distributes a synthetic human identity.

SECTION 01

What Changed

The strategic shift is not simply that AI can imitate people more convincingly. It is that identity can now be separated from a specific performance and reused as a production input.

First: Identity Can Be Reused Without a New Performance

Once enough reference material exists, a recognizable voice, face, cadence or persona can be generated into new executions. A person may appear to deliver words, gestures or messages they never physically performed.

Second: Permission Can Be Valid and Still Reputationally Fragile

The Garbo campaign illustrates the distinction. SKF worked with the estate and family, yet the campaign still prompted debate over whether a deceased person's identity should be reactivated for contemporary commercial use. Permission answered one governance question. It did not answer every audience question. [1]

Third: Public Figures Are Trying to Formally Control Identity as an Asset

Meloni's EUIPO application and Taylor Swift's 2026 U.S. trademark filings show public figures experimenting with additional legal tools to create clearer boundaries around voice and likeness. Reuters reports that these approaches are novel and should not be treated as a universal or settled shield against deepfakes. [2][3]

ARTEFACT91 FRAMEWORK

REAL PERSON → SYNTHETIC ASSET → NEW MESSAGE → AUDIENCE INFERENCE

SECTION 02

Why This Is a Communications Governance Problem

The conventional production model ties a spokesperson to a discrete act: a recorded interview, approved script, photo shoot, speech, endorsement or campaign appearance. Synthetic media weakens that connection. A voice can be translated, an image reanimated, a face inserted into a new setting, and a persona made to deliver new lines long after the original recording.

REPLICABLE — Once a usable likeness model or set of reference assets exists, new executions can be produced repeatedly and sometimes by parties far removed from the original relationship.

TRANSFERABLE — Files, models, prompts, source recordings and derived assets can move between agencies, vendors, platforms, production partners and internal teams unless access and onward-use rights are controlled.

DURABLE — Synthetic identity can outlive a campaign, contract, employment relationship, spokesperson role or even the individual.

The result is a new communications perimeter. Identity is no longer only something a person performs. It can become something an organization stores, licenses, generates and retires.

SECTION 03

The Artefact91 Synthetic Identity Control Stack

ARTEFACT91 FRAMEWORK

RIGHT → CONSENT → SCOPE → CREATE → DISCLOSE → PROVE → RETIRE

RIGHT — Who has the legal or contractual authority to permit the use? The answer can differ for a living person, employee, performer, estate, rights-management company or other representative.

CONSENT — What was actually authorised? Consent for a real recording should not automatically be treated as consent for a voice clone, avatar, synthetic translation, posthumous recreation or materially new message.

SCOPE — Where, for how long, in which territories, on which channels and for what purposes may the identity be used? Scope should also identify prohibited contexts.

CREATE — What system, reference material and human process produced the output? Who approved the script, voice, image, gestures and final context?

DISCLOSE — What does the audience need to know to avoid being materially misled about what they are seeing or hearing? Disclosure may also be legally required in some jurisdictions.

PROVE — Can the organization demonstrate provenance after publication: rights documentation, source assets, process records, approvals, disclosures and the final distributed file?

RETIRE — Can the organization reliably stop future use when the term ends, the relationship changes, the campaign closes or rights are withdrawn where the agreement allows it?

The stack turns a broad ethics debate into a communications operating system.

SECTION 04

The Synthetic Identity Risk Matrix

Use two dimensions before publication: certainty of authority and the risk that an audience will infer something untrue about the performance, endorsement or timing.

CLEAR AUTHORITY + LOWER DECEPTION RISK — Proceed with controls.

CLEAR AUTHORITY + HIGHER DECEPTION RISK — Disclose and limit.

AMBIGUOUS AUTHORITY + LOWER DECEPTION RISK — Pause and clarify.

AMBIGUOUS AUTHORITY + HIGHER DECEPTION RISK — Stop and escalate.

Legal permission and audience understanding are separate dimensions. A campaign can be authorised and still be misleading.

SECTION 05

Consent Is Not the Same as Audience Understanding

Audiences do not see the contract. They see the output. That creates four common inference failures:

PERFORMANCE INFERENCE — People may assume the person physically performed or recorded the material when they did not.

ENDORSEMENT INFERENCE — People may assume the person personally considered and endorsed a specific product, cause, policy or claim when the synthetic output was generated under a broader licence.

TIMING INFERENCE — People may assume the message reflects the person's current views, current health, current role or current availability.

AUTHENTICITY INFERENCE — People may assume that because the face and voice are recognisable, every word and gesture originated with the person.

The disclosure question should therefore be framed as: what would a reasonable audience infer if we said nothing?

SECTION 06

Disclosure Is Becoming a Governance Standard

The EU AI Act's Article 50 transparency obligations have applied since 2 August 2026. The European Commission says providers of certain generative systems must add machine-readable marks to enable detection of AI-generated or manipulated content, while deployers must disclose certain deepfakes and other covered AI outputs. Limited transitional provisions apply in defined circumstances. [4]

For communications teams, disclosure is therefore more than a creative preference. It can sit at the intersection of law, platform policy, audience trust and brand credibility.

IAB's 2026 research provides a useful industry signal. In a survey of 505 U.S. Gen Z and Millennial consumers and 104 advertising executives, 83% of executives said their companies had deployed AI in the creative process. Yet the research found a 37-point perception gap: 82% of ad executives believed younger consumers felt positive about AI-generated ads, compared with 45% of surveyed consumers who actually reported feeling positive. The study also found strong consumer interest in disclosure for fully AI-generated ads, AI video/images, and substantial interest in disclosure of AI voices and avatars. [6]

These findings should not be generalized to every audience or market. They are strong enough to support a practical conclusion: communications teams should not assume that internal enthusiasm for synthetic creative equals external acceptance.

SECTION 07

The Rights Landscape Is Fragmented

There is no single global rule for synthetic identity. Tennessee's ELVIS Act, effective since July 2024, expanded the state's protection of personal rights to include voice and expressly addresses AI-enabled simulation. Other jurisdictions rely on different combinations of publicity rights, privacy law, trademark, copyright, contract, consumer-protection law, platform rules and emerging AI-specific legislation. [5]

That fragmentation creates a communications-management problem even where lawyers are advising on the legal position. A campaign that is acceptable in one jurisdiction may require a different disclosure, rights structure or distribution decision elsewhere.

SECTION 08

Build an Identity Rights Ledger

Every planned synthetic identity use should have a standing record. The ledger should travel with the asset across internal teams, agencies, vendors and distribution partners.

RIGHTS HOLDER — Who can authorise the use and on what basis.

IDENTITY ELEMENT — Face, voice, image, gesture, signature phrase, avatar, synthetic translation or other distinctive characteristic.

SOURCE MATERIAL — Reference recordings, images or other assets used to produce the simulation and who controls them.

PERMITTED / RESTRICTED USE — Exact campaign, purpose, claims, products, causes, channels and prohibited contexts.

TERRITORY + TERM — Where use is authorised, start date, expiry date and renewal conditions.

DISCLOSURE STANDARD — Audience-facing label, contextual statement and machine-readable marking required by channel or jurisdiction.

APPROVAL OWNER — Who must sign off each materially new execution.

PROVENANCE RECORD — Where the final asset, script, approvals and production evidence are stored.

RETIREMENT — How generation and distribution stop when permission or relevance ends.

INCIDENT CONTACT — Who leads if the use is disputed, copied, altered or exceeds scope.

SECTION 09

Posthumous Use Needs an Extra Test

Deceased-person likenesses deserve an additional layer of review because consent is necessarily mediated through estates, heirs, contracts or other rights structures.

- **LEGITIMATE AUTHORITY** — Who is authorising the use, and is that authority documented for the relevant territory and medium?
- **KNOWN INTENT** — Is there credible evidence that the proposed use is consistent with the person's known preferences, past endorsements, professional identity or restrictions?
- **CONTEXT FIT** — Would the identity support a product, claim or cause the person never encountered and could not evaluate?
- **AUDIENCE DISCLOSURE** — Will the execution make it clear that the performance is synthetic or reconstructed where disclosure is appropriate or required?
- **LEGACY CONSEQUENCE** — Could the campaign materially reshape public memory of the person in a way the estate, family, collaborators or audience are likely to contest?

This is not an argument against posthumous creative use. It is an argument for treating legacy as a stakeholder, not merely an asset.

SECTION 10

The Spokesperson Contract Has Changed

Traditional spokesperson and ambassador agreements focus on appearance, recording, usage rights and term. Synthetic media creates additional governance questions that should be resolved before production begins.

Can the organization generate new speech the person never recorded?

Can the voice be translated into another language while preserving vocal identity?

Can facial movement, gesture or context be altered?

Can the likeness be paired with claims the person has not separately approved?

Can the synthetic version continue after the person leaves the organization?

Can the material be used for political, advocacy or public-interest communications?

Can the identity be licensed onward to agencies, platforms or technology providers?

Who controls reference files, model outputs and derivative assets after the term ends?

What approval is required for a materially new context?

SECTION 11

When Synthetic Identity Becomes a Crisis

An authorised synthetic campaign can still become an incident. Common triggers include a rights-holder dispute, an out-of-scope edit, audience confusion about whether the person actually spoke, a model output that changes approved meaning, reuse after expiry, misleading translation, or third parties copying the synthetic asset into unauthorised content.

ARTEFACT91 FRAMEWORK

PAUSE → VERIFY RIGHTS → FREEZE → CLARIFY → CORRECT / WITHDRAW → ARCHIVE → REVIEW

The first public objective is not aesthetic defence. It is to restore a verifiable account of what was authorised, what was synthetic, what circulated, what remains uncertain, and what the organization has stopped or corrected.

SECTION 12

A 30-Day Readiness Sprint

WEEK 1 — INVENTORY: Identify every current or planned use of synthetic human identity across advertising, social media, public affairs, internal communications, training, events, customer service and executive content.

WEEK 2 — GOVERN: Create the Identity Rights Ledger, define high-risk categories, establish mandatory legal/reputation review triggers and specify who can approve synthetic use.

WEEK 3 — DISCLOSE + PROVE: Adopt a disclosure standard by channel, capture provenance records, preserve final approvals, and address machine-readable marking obligations where applicable.

WEEK 4 — TEST + RETIRE: Run a tabletop scenario involving an out-of-scope synthetic spokesperson use. Test takedown, rights-holder contact, public clarification and asset retirement. Remove expired or unsupported synthetic assets from active libraries.

SECTION 13

The Strategic Question

Generative AI makes it technically possible to keep a recognisable identity speaking long after the original performance, campaign, employment relationship—or even lifetime—has ended. The communications advantage is obvious. The governance burden is equally real.

Who owns the right?

What exactly was consented to?

Where does the permission end?

What will the audience think happened?

What must be disclosed?

What evidence will we preserve?

How does the identity get retired?

If those answers are unclear, the synthetic execution is not ready.

ARTEFACT91 BRIEFING

Artefact91 Point of View

The next reputation risk is not simply that someone will fake your spokesperson. It is that your own organization will use a synthetic identity faster than its governance can explain, prove or stop.

ARTEFACT91 POSITION

TREAT IDENTITY AS INFRASTRUCTURE.

ARTEFACT91 BRIEFING

Sources & Research Base

- [1] Reuters, 2 October 2026, “Swedish ball-bearing ad revives Greta Garbo with US, Chinese AI.” <https://www.reuters.com/business/media-telecom/swedish-ball-bearing-ad-revives-greta-garbo-with-us-chinese-ai-2026-10-02/>
- [2] Reuters, 5 October 2026, “Italy PM Meloni seeks EU trademark for her voice amid AI deepfake risks.” <https://www.reuters.com/technology/italy-p-m-meloni-seeks-eu-trademark-her-voice-amid-ai-deepfake-risks-2026-10-05/>
- [3] Reuters, 27 April 2026, “Taylor Swift files to trademark her voice, likeness to ward off AI deepfakes.” <https://www.reuters.com/legal/litigation/taylor-swift-files-trademark-her-voice-likeness-ward-off-ai-deepfakes-2026-04-27/>
- [4] European Commission, Article 50 transparency guidance and FAQs; obligations apply from 2 August 2026. <https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-transparency-obligations>
- [5] State of Tennessee, ELVIS Act materials on voice and likeness protection. <https://www.tn.gov/governor/news/2024/1/10/tennessee-first-in-the-nation-to-address-ai-impact-on-music-industry.html>
- [6] Interactive Advertising Bureau, “The AI Ad Gap Widens,” 15 January 2026. <https://www.iab.com/insights/the-ai-gap-widens/>

ARTEFACT91 BRIEFING

Research Note

This briefing synthesises current reporting, primary regulatory material and industry research for strategic communications planning. Legal rights in voice, likeness, trademark, publicity, privacy and synthetic media vary materially by jurisdiction. The analysis does not treat a trademark filing as settled protection, estate consent as automatic reputational legitimacy, or survey findings as universal evidence of audience behaviour. This briefing is not legal advice.