BRIEFING 17 · 2026

The AI Disclosure Gap

Why political campaigns need a control system for synthetic media, not a label at the end

ARTEFACT91

Campaign Intelligence + Strategic Communications

artefact91.com

ARTEFACT91 BRIEFING 17

In this briefing

Why political campaigns need a control system for synthetic media, not a label at the end

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 17

Executive Brief

Current signal

The Wesleyan Media Project has identified at least 164 U.S. political ads generated or enhanced with AI in the 2025–2026 cycle, with just under $80 million in known spending. Associated Press reporting on 25 September found that 69% of those identified ads did not disclose AI use. State regulation is expanding, but disclosure practice remains inconsistent and the legal scope varies by jurisdiction.

Political campaigns are normalizing AI-assisted creative production faster than transparency systems are becoming operational. The issue is not simply whether synthetic media exists. It is whether campaigns, consultants, platforms and regulators can reliably identify when AI materially shaped a communication, determine what rules apply, disclose it in a way audiences can actually notice, and preserve evidence of compliance.

As of 4 September, the Wesleyan Media Project had identified at least 164 political advertisements generated or enhanced with AI during the current U.S. election cycle, representing just under $80 million in spending. The project explicitly calls this an undercount because AI use is often difficult to detect and there is no universal federal disclosure rule.

Associated Press reporting published 25 September added a sharper operational signal: 69% of the identified ads did not disclose AI use. When the project cross-referenced the ads with state regulation, only 29% of ads in states with AI rules carried a disclosure, compared with 32% in states without such laws. Those figures must be interpreted cautiously: not every AI-assisted ad falls within every statute, and some laws took effect after particular ads aired. The finding is therefore evidence of an uneven disclosure environment, not proof that 71% of regulated ads were unlawful.

SIGNAL	FIGURE	READING
Identified AI ads	164+	Wesleyan describes this as an undercount
Known spending	≈ $80M	Tied to the identified ads only
Carried no AI disclosure	69%	Of the identified sample, not of all advertising
States regulating political deepfakes	31	Mostly disclosure requirements rather than bans

These figures use different units and should not be read as a single rate series.

The regulatory map is nevertheless expanding. The National Conference of State Legislatures reported in June that 31 states had enacted laws regulating deepfakes in political messaging. Most use disclosure requirements rather than outright bans. California added further election-related AI rules in September, extending and expanding disclosure and deceptive-media provisions.

The strategic problem for campaigns is that “put a label on it” is not a sufficient governance model. The content may travel across broadcast, connected television, social platforms, influencer posts, email, fundraising and earned media. The legal threshold may depend on whether the content depicts a real person, whether the depiction is materially deceptive, when it runs relative to an election, who paid for it, and the jurisdiction in which it appears.

Briefing 06 examined the institutional response gap once misinformation is already circulating. Briefing 12 examined the infrastructure of AI-enabled influence operations. Briefing 14 examined manufactured stakeholder voice. This briefing addresses a different operating question: how legitimate campaigns and advocacy organizations govern their own use of synthetic and AI-assisted content before publication — and how they prove what they did afterward.

SECTION 01

The Disclosure Problem Is Larger Than Deepfakes

Campaign teams increasingly use AI across a spectrum: background replacement, image extension, voice cleanup, synthetic performers, generated b-roll, cloned voices, fully generated scenes, translation, captions, script development and personalization. Laws often regulate only a subset of that spectrum. A campaign can therefore be ethically exposed even when it is technically outside a specific statutory definition.

The key distinction

Synthetic-media governance needs two tests: a legal test — “must we disclose?” — and a trust test — “would a reasonable viewer want to know this was materially generated or altered?” The second test can be stricter than the first.

THE OBSERVABILITY PROBLEM — The Wesleyan Media Project identifies AI ads through news reporting, explicit disclaimers, vendor data and visual review, and acknowledges that many uses will remain invisible. The market cannot rely on outside detection as a control. Campaigns need internal records of AI use at asset level.

THE JURISDICTION PROBLEM — The 31 state laws differ on timing, media type, intent, depiction, disclosure language, remedies and penalties. Some require a label; some prohibit defined deceptive uses during an election window; Maryland uses a broader prohibition model. A national campaign can face different rules for the same asset across placements.

THE NOTICE PROBLEM — Even a legally compliant disclaimer may not function as effective notice. Canadian experimental research from The Dais found that small AI-content labels had no meaningful effect on trust or sharing behaviour, while full-screen interstitial labels were more noticeable and effective. The lesson is not that every campaign should use a full-screen block, but that legibility, prominence and placement are part of disclosure quality.

SECTION 02

The Artefact91 Political AI Disclosure Control Chain

ARTEFACT91 POLITICAL AI DISCLOSURE CONTROL CHAIN

INVENTORY → CLASSIFY → MAP LAW → LABEL → PROVENANCE → APPROVE → VERIFY LIVE → ARCHIVE

1 — INVENTORY — Record every campaign asset in which generative AI materially contributed to a visual, voice, representation, translation, edit or scene. Do not depend on creators remembering at launch.

2 — CLASSIFY — Classify what AI changed: cosmetic edit, synthetic object, synthetic environment, synthetic human, cloned or altered voice, impersonation, depiction of conduct that did not occur, or fully generated content.

3 — MAP LAW — Map the jurisdictions, media and election windows in which the asset will run. Record the applicable disclosure language, placement rules, exceptions and enforcement route.

4 — LABEL — Apply the disclosure in the form required by law and test whether it is actually legible in the final medium. A disclosure that disappears in crop, autoplay, compression or reposting is not an operational disclosure.

5 — PROVENANCE — Preserve original inputs, model or tool used, prompts or production notes where appropriate, human edits, approvals and final export. Where supported, retain content credentials or other provenance metadata.

6 — APPROVE — Create a named human approval point for synthetic media. High-risk uses — especially depictions of real candidates, election officials or public figures — should require legal and senior communications review.

7 — VERIFY LIVE — After publication, verify the live asset. Check the label, sponsor disclaimer, crop, captions, thumbnail, repost behaviour and platform rendering. The final delivered file is not the same as the live communication.

8 — ARCHIVE — Keep an auditable record linking asset, approval, jurisdiction, disclosure, placement and dates. If challenged, the organization should be able to reconstruct what happened without relying on staff memory.

SECTION 03

The Artefact91 Synthetic-Media Risk Test

Before a political or advocacy asset is released, assess eight dimensions. The test is designed to support judgment and escalation; it does not replace jurisdiction-specific legal advice.

IDENTITY — Does the asset depict or imitate a real candidate, officeholder, election official, public figure or identifiable private person?

REALITY CLAIM — Could a reasonable viewer infer that depicted speech, conduct, events or circumstances actually occurred?

MATERIALITY — Is AI use central to the persuasive claim or merely incidental production assistance?

TIMING — Is the communication inside a statutory pre-election window or subject to electioneering rules?

PLACEMENT — Will the asset appear in broadcast, connected television, social, email, influencer, phone, fundraising or another regulated format?

DISCLOSURE QUALITY — Is the notice clear, conspicuous, durable through crops and reposts, and understandable to a normal viewer?

PROVENANCE — Can the campaign demonstrate how the asset was made, altered, approved and distributed?

CONSEQUENCE — If challenged, could the asset create legal, platform, press, donor, coalition or candidate-trust consequences?

SECTION 04

What Campaigns Should Not Do

- Treat “AI-assisted” as a single category. Generating a background and fabricating a candidate’s voice create radically different risks.
- Assume a platform’s built-in AI label satisfies campaign law. Platform metadata and statutory campaign disclaimers are different controls.
- Add the disclosure only at the master-file stage. Cropping, aspect-ratio changes, reposts and influencer edits can remove or obscure it.
- Use the presence of a label as a substitute for truthfulness. A disclosed synthetic claim can still be misleading, defamatory or otherwise unlawful.
- Assume parody or satire is automatically exempt everywhere. Exceptions vary by statute and context.
- Delete production records after launch. In a complaint, provenance may become part of the campaign’s ability to explain intent and compliance.

SECTION 05

Build a Two-Tier Approval Standard

TIER	WHAT IT COVERS	MINIMUM CONTROL
Tier 1 — Assistive	Resizing, cleanup, background extension, translation, captions, generic generated imagery	Asset inventory, disclosure check, normal creative approval
Tier 2 — Representational	Synthetic people, cloned voices, realistic scenes, depictions of real persons or events that did not occur	Legal review, senior communications approval, provenance record, live disclosure verification

SECTION 06

The Comms Team Needs a Live Disclosure Register

For every synthetic or materially AI-altered campaign asset, maintain a simple register.

ASSET ID — Master file and all variants.

AI USE — What was generated or altered.

PERSONS DEPICTED — Real, synthetic or composite.

JURISDICTIONS — Where the asset will be distributed.

LEGAL BASIS — Relevant rule, exception or counsel note.

DISCLOSURE — Exact language and placement.

PROVENANCE — Production record or content credentials where available.

APPROVER — Named human owner.

LIVE CHECK — Screenshot or recording of final placement.

RETENTION — Archive location and period.

Operational rule

If the team cannot answer "where did this asset run, what AI materially changed, what disclosure applied, who approved it and what did audiences actually see?", the campaign does not have a disclosure system — it has a labeling habit.

SECTION 07

What Leaders Should Measure

DISCLOSURE COVERAGE — Share of AI-material assets with required disclosure before launch.

LIVE DISCLOSURE INTEGRITY — Share of placements where required notice survives final rendering, crop and repost.

JURISDICTION MATCH — Share of assets mapped to the correct state or local rules before trafficking.

PROVENANCE COMPLETENESS — Share of high-risk assets with a complete production and approval record.

EXCEPTION RATE — How often campaigns rely on parody, satire, incidental-use or other exceptions.

CORRECTION LATENCY — Time from identifying a disclosure failure to correcting or pulling the asset.

TRUST SIGNAL — Whether audiences understand what the disclosure means, not merely whether text was technically present.

SECTION 08

A 30-Day Readiness Sprint

WEEK 1 — INVENTORY & POLICY — Inventory current AI uses and define which production activities must be recorded. Build a simple policy distinguishing assistive and representational uses.

WEEK 2 — JURISDICTION MAP — Create a campaign-specific map of disclosure and synthetic-media rules for every state or locality where paid or coordinated communications will run.

WEEK 3 — WORKFLOW & TEMPLATES — Embed AI-use fields in creative briefs, traffic sheets and approval forms. Create disclosure-safe templates for common aspect ratios and platforms.

WEEK 4 — SIMULATION & AUDIT — Run a simulated challenge: select one live synthetic asset and ask the team to reconstruct its AI use, legal analysis, approval and final placement within 30 minutes. Fix whatever cannot be produced.

ARTEFACT91 BRIEFING 17

Artefact91 Point of View

AI disclosure is becoming a campaign operating-system problem. A visible label is only the final output. The real control is the chain that identifies AI use, classifies the content, maps applicable law, preserves provenance, obtains approval, verifies the live placement and keeps an auditable record.

Political AI governance is moving from “Can we generate this?” to “Can we defend how this was made, disclosed and distributed?” The campaigns that treat synthetic media as a controlled production system — rather than a novelty tool — will be better positioned to move quickly without sacrificing credibility, compliance or evidentiary discipline.

ARTEFACT91 BRIEFING 17

Sources & Research Base

- Wesleyan Media Project — “At Least $80 Million Spent on AI Ads This Cycle,” 4 September 2026.
- Associated Press / Washington Post — “AI ads are popular in the 2026 midterms, but state laws haven’t stopped a lack of disclosure,” 25 September 2026.
- National Conference of State Legislatures — “Artificial Intelligence (AI) in Elections and Campaigns,” updated 23 June 2026.
- National Conference of State Legislatures — “2026 Election Enactments,” 2026.
- Office of Governor Gavin Newsom — new laws to protect California elections, 19 September 2026.
- The Dais, Toronto Metropolitan University — “Human or AI? Evaluating Labels on AI-Generated Social Media Content,” 2025.

Research Note

This briefing synthesizes current public research, enacted and reported legal developments, and Artefact91 analysis. It does not provide legal advice. Because political synthetic-media laws vary materially by jurisdiction and continue to change, campaigns should verify the law applicable to each placement before release. The 69% disclosure figure is an observed characteristic of the Wesleyan Media Project’s identified sample; it should not be treated as a universal estimate of all 2026 political advertising or as proof that every undisclosed ad violated a law.