LAYER 4 AI ADVERTISING TRUST STACK

Technical Conformance

From intention to evidence.

What companies must be able to prove as AI acts and adapts.

AUTHORIZED

What was it allowed to do?

DISCLOSED

What were people told?

ACTUAL

What did it really do?

Jay Mandel and Sean Meier

The Intention Advantage | Institute for Advertising Ethics

IEEE Standards Association, Knowledge Partner

12 October 2026 / Working proposal for discussion

THE TRUST STACK / LAYER 4

Trust depends on the integrity of the evidence.

BURGER

Ingredients and sourcing

DIAMOND

Origin and authenticity

DIGITAL TRADING CARD

Issuer and ownership

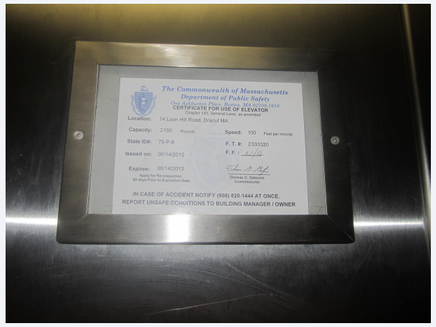

ELEVATOR

Inspection and safety

Provenance is the record of origin and what happened along the way.

Certification needs credible checks, clear scope and an accountable issuer.

Our mandate: preserve evidence to test actual behavior against authorized intent, permissions and disclosures.

Layer 1 defines intent.

Purpose and boundaries.

Layer 3 defines controls.

Approvals and limits.

Layer 4 supplies evidence.

What actually happened?

Layers 5 and 6 use the evidence for accountability and ongoing assurance.

Analogies illustrate different checks. A badge alone proves nothing.

Images: cyclonebill / Commons (CC BY-SA 2.0), Rawpixel (CC0), NBA Collect by Topps (card), NWRGeek / Commons (CC0).

CONTENT RECORD

AI ASSISTED

DISCLOSURE EXAMPLE

VIEW RECORD

SUMMIT I / THE PROBLEM LAYER 4 MUST ADDRESS

An ad agent can hit its target and exceed its authority.

An AI ad agent is software that creates ads, selects audiences or spends budget.

	CUSTOMER ASKS	ADVERTISER ASKS
PERSONAL DATA IN AD TARGETING	Why did you use my personal data to target this ad?	Which data and permissions supported this targeting?
CHANGING BEHAVIOR	Why did the offer or message change?	Did the ad agent stay within its approved audience and budget?
ACCOUNTABLE OWNER	Who can put it right when the ad agent gets it wrong?	Who approved the action, and could they stop it?

LAYER 4 CONNECTS FIVE EVIDENCE FIELDS

SOURCE | PERMITTED USE | RESPONSIBILITY | CHANGES | VERIFICATION

Layer 4 records permitted and actual use. Other layers define the rules and response.

32% mistook human-created content for AI-generated content in the study.

Scenarios: IAE Summit I report, 22 Apr 2026. Statistic: Zefr + OM Media Trials, US and Canada, 2026 (click to read).

TECHNICAL CONFORMANCE / SOURCE TO RECEIPT / PROPOSED FLOW

From stated intent to checkable evidence.

Capture what went in, what was permitted and what the ad agent actually did.

1. COLLECT

Approved task and limits.
Audience data: supplier and permitted use.
Images, voices, music: permission to use.
Creator disclosures and tools used.
Existing signed origin/edit records (C2PA).
Watermark and AI-detection results.

→

2. REVIEW

Compare proposed targeting, content and spending with purpose and permissions.

Record checks, exceptions and approval or rejection.

→

3. VERIFY

Human approval: identity, presence, authority and power to reject or stop.

Link actual actions and method changes to inputs, approvals, versions and time.

→

4. ISSUE

Record what was authorized, disclosed and done.

Include evidence links, scope, date and accountable owner.

Content Credentials are signed history records. Watermarks are signals embedded in media.

A supported watermark may help recover a record. AI-detection results are estimates, not proof.

CARRY + RECHECK

Record changes, human check-ins and interventions across handoffs.
Recheck after changes, withdrawal or expiry. Link credentials to permission and decision records.

Record automation explicitly. Protect private identity evidence. Click here for C2PA explanations of credentials and watermark recovery.

REGULATORY ALIGNMENT / AS OF 11 OCT 2026

Translate obligations into evidence.

RULE AND TIMING	WHAT IT REQUIRES	WHAT THE RECORD CAN SHOW
EU AI Act, Article 50 In effect since 2 Aug 2026*	Machine-readable marking for covered AI outputs; disclosure of deepfakes.	AI involvement, credentials, disclosure and the exact content version.
New York synthetic performers In effect since 9 Jun 2026	Disclosure in covered ads using AI-generated synthetic performers.	Performer type and the disclosure attached to the published ad.
California synthetic performers SB 1050; effective 1 Jan 2027	Clear disclosure in covered ads that prominently feature synthetic performers.	The disclosure, reviewer, release date and version that ran.

The receipt supports review. It does not establish legal compliance on its own.

Commercial value: answer buyer questions, investigate problems and protect rights.
Violations may bring fines or restrictions. A seal cannot replace the required conduct.

*Scope and exceptions apply. Older EU systems have a marking grace period until Dec 2026. Click each row for its source.

WORKSTREAM PROGRESS / DECISIONS FOR THE COMMITTEE

A proposed record. Clear decisions still to make.

WHAT WE HAVE PROPOSED

Five minimum evidence fields.

A flow that preserves proof across handoffs.

Public checks with protected identity evidence.

WHAT THE COMMITTEE MUST DECIDE

What proof is sufficient for each risk?

Who verifies, and how long is evidence kept?

Who answers for a missing handoff?

Possible application: an AI LUMAscape linked to evidence.

COMPANY
Brand, agency, AI or adtech

→

REQUIREMENT
Purpose, permission, controls

→

EVIDENCE + STATUS
Reviewer, scope, date, result

Follow the workstream: jaymandel.com/trust-stack

Receipt, seal and directory are proposals, not an adopted standard. Test the record in a real cross-company workflow.

Example site uses fictional companies. Any assessment must state its scope and does not establish legal compliance.