AI Vendor Claims & Evidence Audit

ILLUSTRATIVE PRODUCT DECISION GUIDE | PUBLIC-SOURCE AND SERVICE METHODOLOGY

No completed client matter is represented. This document demonstrates questions, evidence requirements, limitations and possible deliverables. It contains no verified findings about a named buyer or company. A matter, source access, fixed fee and timing require separate written acceptance.

---

1. Cover copy: question and promise

AI Vendor Claims & Evidence Audit

The first decision: Can vendor AI assertions be traced to evidence that supports the buyer's decision?

Use this review when: AI procurement, public performance claim, model release or new AI policy.

Target reader: AI buyers, vendors, risk committees and disputes teams.

Illustrative subtitle: A clearly bounded external-evidence review identifying what original records substantiate, which claims are contested and what the client must request before relying on an assertion.

2. The buyer decision and research boundary

The decision owner must resolve: **Can vendor AI assertions be traced to evidence that supports the buyer's decision?** The matter is activated by a documented event such as **ai procurement, public performance claim, model release or new ai policy**. This specimen is intentionally subject-neutral. For the real commission, the research team must insert the actual named entity, decision date, jurisdiction, exact wording of the challenged claim and period before claiming anything factual.

Agreed scope model: One named subject or event; up to three material claims; agreed source jurisdictions and date window.

What is not automatically included: unbounded jurisdictions, confidential database access, professional regulated opinions, private surveillance, hardware or onsite testing, urgent escalation and any undisclosed third-party subscription costs.

3. What the buyer provides before instruction

- Exact name and identifiers of the relevant company, asset, product, person or project and its jurisdiction.
- The decision they must make, their deadline and the particular assertion or document that prompted it.
- Lawfully sharable materials, with any personal and confidential information minimised.
- Details of known adverse records, disputes and contradictory accounts, rather than a one-sided brief.
- Approval of the written scope, access conditions, final fee, permitted reliance and independent reviewer.

4. Investigative questions prepared for the real commission

- Can vendor AI assertions be traced to evidence that supports the buyer's decision?

- Which model card establishes the relevant subject, version and reporting period?
- Could unreplicable accuracy claim overturn the decision-relevant statement?
- What original vendor technical document supports an independent cross-check rather than repeating the same source?
- Which dated records would distinguish missing evaluation population from a genuine discrepancy?
- What source result would justify a specialist referral rather than a research conclusion?

5. Evidence acquisition order and verification method

- **Resolve the subject.** Start with `model card` and capture issuer, entity match, record date, retrieval timestamp and original identifier.
- **Triangulate.** Compare `vendor technical document` against the first source, testing whether both genuinely derive from independent originals.
- **Establish sequence.** Examine `NIST AI RMF` and order assertion, contemporaneous record, amendment and current status.
- **Seek negative evidence.** Test `unreplicable accuracy claim`, `missing evaluation population` and `unsupported security promise` with documented competing hypotheses.
- **Document the cutoff.** Name what is inaccessible, subscription-controlled, outdated or outside jurisdiction, and stop when the evidence cannot support more.
- **Human review.** Route legal, audit, security, engineering, medical, valuation or licensing questions to an authorised professional if applicable.

Why an evidence analyst is involved: Audit claim-to-source traceability, identify counterexamples and document information gaps. The value exists only when the recorded work improves the buyer's decision beyond readily accessible search results, an internal analyst or a self-service AI summary.

6. Ready-to-typeset example evidence matrix

Question/claim; First applicable original record; Best competing evidence; Illustrative state before casework; Next evidence action

Identity, legal capacity and exact subject; `model card`; Alias/period mismatch; Not yet examined; Obtain verified legal or asset identifier

Main decision assertion; `vendor technical document`; `unreplicable accuracy claim`; Not yet examined; Request original dated statement and counter-record

Material event timing; `NIST AI RMF`; `missing evaluation population`; Not yet examined; Reconcile dated event and publication chronology

Adverse exception; Independent contrary source; `unsupported security promise`; Not yet examined; Ask precise targeted contradictory question

Meaning of statuses (for the eventual real report):

Evidence state; Meaning for the decision owner

Supported; Applicable original records support the narrowly worded claim for the stated time and entity.

Partially supported; Some elements are established, while material qualifiers, scope or periods remain open.

Contradicted; A more reliable applicable record conflicts with the proposed assertion.

Unresolved; Accessible evidence is absent, ambiguous, outdated or conflicting; further documents are required.

Outside mandate; The question requires a reserved professional function or additional agreed research.

7. Concrete specimen outputs (without fictional findings)

Output committed for this product: Source coverage sheet, first three contradictions or gaps, clear next-record requests.

Suggested report artifacts:

- **Claim-To-Test Map**: rows for evidence question, entity match, issuer, record identifier, effective date, retrieval date, contrary document, confidence explanation and reviewer disposition.
- **Documentary Gap Register**: a table distinguishing event dates, record publication dates, claim dates, subsequent amendment and unresolved chronological gaps.
- **Vendor Clarification Pack**: a controlled clarification list addressed to the relevant source owner or instructing professional.
- A concise executive decision note stating what may be relied upon for the specified purpose and which matters remain unresolved.

Editorial visual specification: Include one actual evidence matrix or original-record timeline only when underlying evidence exists. Never insert decorative charts, fabricated relationships, arbitrary ratings or imaginary completed cases.

8. Source acquisition catalogue for this offer

These are real portals and competitor-method references to evaluate, not proof of any specific subject. Access, freshness, jurisdiction and licensing must be rechecked during production.

Ref; Primary source / comparator; Link; Evidence limitation and status

EU-AI-2026; European Commission AI Act Service Desk: AI Act staged enforcement and application; https://ai-act-service-desk.ec.europa.eu/en/ai-act/faq/when-does-enforcement-start; WEB CHECKED 2026-10-12. The dates differ by obligation; do not treat all AI Act requirements as simultaneously applicable

NIST-AI; U.S. NIST: AI risk management framework; https://www.nist.gov/itl/ai-risk-management-framework; WEB CHECKED 2026-10-12. Voluntary risk management framework; not an AI certification

9. Deliberately adverse questions and professional boundary

- If `unreplicable accuracy claim`, does the strongest original record still justify the same assertion?
- Could `missing evaluation population` result from a date mismatch, name collision, incomplete publication or source-update lag?
- What would demonstrate `unsupported security promise` and require a 'cannot determine' outcome?
- **Professional boundary:** No AI certification, conformance assessment, penetration testing or blanket accuracy guarantee.

10. How a reader commissions the real product

The reader identifies the subject, decision date and three most material questions and requests a **written scope and fixed fee** through the existing website pathway: `/specialist-divisions/ai-governance-model-risk-evidence-assurance/offers/evidence-first-look/`. An initial enquiry must not contain confidential attachments. No work begins, payment occurs or information is treated as professionally verified until human acceptance, conflicts, lawful sourcing, pricing and the secure handoff process are complete. This illustrated manuscript does not constitute an offer to undertake regulated activity.

11. Production / reliance quality checklist

- Confirm the appointing client, exact research question, relevant location, decision date, conflicts, privacy, sanctions and licensing.
- Check original source identity, authority, issue date, effective date, retrieval timestamp, version and lawful access.
- Reproduce important calculations and compare at least one plausible contrary explanation.
- Keep allegations, regulatory findings, historical claims and unresolved matches distinct.
- Mark each proposed finding supported, partially supported, contradicted, unresolved or outside mandate; include missing-source requests.
- Confirm that illustrations do not imply a real engagement, completed investigation, certified outcome or affiliation with named comparators.
- Obtain final human sign-off on accuracy, legal boundaries, confidentiality, claims and publication version before release.

Visual publication policy: When instructed to generate a PDF later, follow the approved Aperture PDF visual master, original approved brand art, document control page and separate final human sign-off. There is no PDF yet.

Professional boundaries: Independent research support only. No legal, tax, investment, audit, valuation, brokerage, investigation requiring a specific professional licence or regulated opinion is implied. No confidential material should be sent before instructions.