Record sufficiency explained: the ten questions a decision record must answer to be explainable and challengeable

24 November 2026 · 8 min read · blogs.iictp.org/record-sufficiency-decision-records-ten-questions-explained

Short answer: “record sufficiency” is a test proposed in an arXiv preprint submitted on 17 November 2026: a decision record is sufficient if a reader who has only the record and the items it cites can explain the status, bound it, re-derive it and challenge it. The author proposes ten questions a record must answer and four mechanical checks. It matters for legal and compliance tooling because a log of approvals is not the same as a record of why a claim became an operative status. This is an arXiv preprint. It has not been peer reviewed, and its numbers are the authors’ own claims.

TL;DR The short version

- 'Record sufficiency' says a decision record is sufficient if a reader with only the record and its citations can explain, bound, re-derive and challenge the status it created.
- The author proposes ten questions a record must answer and finds that schema validation alone misses most defects.
- It is a single-author preprint on synthetic records, but you can test one real decision record against the ten questions today.

Read online, listen to it, or change the text size: https://blogs.iictp.org/record-sufficiency-decision-records-ten-questions-explained

What is record sufficiency?

The paper, “Record Sufficiency: What a Decision Record Must Preserve When an Unsettled Claim Becomes an Operative Status”, starts from the point that organisations keep evidence, logs and approvals but often not the step by which an unsettled claim became an operative status, such as an accepted label or a released image. Later readers then cannot tell what the status was allowed to mean.

How it works

- **Four uses of a record:** explain the status, bound it, re-derive it and challenge it.
- **Ten questions** follow from those uses, and the author maps them onto an earlier JSON Schema for decision records (DRK v0.1).
- **Four mechanical checks** are defined that a program can run on a record.

- **Four levels kept apart:** conformance (decided by the checks), adequacy (whether answers are specific enough to use, which needs a reader), fidelity (which no check on the record decides) and merits (which sufficiency does not require).

Results reported by the source

- On 164 single-defect variants of three synthetic records, schema validation, the only check DRK v0.1 implements, rejected every variant with a deleted required property and none of the others. It also accepted a record whose claim, rule, scope, limits and challenge route were empty. The author says these are properties of the schema on seeded defects, not estimates for real records, and that a false record passes every check.
- The statement-of-reasons schema of the EU Digital Services Act Transparency Database answers 4 of the ten questions in full, 2 in part and 4 not at all. The unanswered ones concern why the matter was unsettled, what the status depends on, whether it can be re-derived and how to challenge it.
- In 25 real statements from 25 platforms, drawn by a seeded rule and coded by a language model and checked by the author, free text addressed any of those four questions in only 2, and in 18 of them (24 under a stricter reading adopted during coding) no free text said why the item fell under the rule.

Limitations and open questions

- Single author, preprint, synthetic records for the schema test, and a small real sample of 25 statements. The paper says no deployment or legal claim is made.
- The coding of the real statements used a language model checked by the author, so the figures depend on that coding.
- Sufficiency is not correctness. A well-formed record can still be wrong; the paper keeps fidelity and merits outside what its checks decide.

How to apply it

1. Pick one decision your organisation makes that turns a claim into a status, such as approving a model release or a content-moderation action.
2. Ask the ten questions of the existing record, and note which are answered, partly answered or missing.
3. Add the missing fields, such as why the matter was unsettled, what the status depends on and how to challenge it.
4. Run mechanical checks for completeness, and have a person review adequacy.

What it means for legal and compliance teams

Good records support appeals and audits. A record that cannot show why a rule applied cannot be challenged properly. This connects to architecture decision records in engineering, and to evidence for agent delegation. This is general information, not legal advice.

Frequently asked questions

- **What is a sufficient decision record?** One from which a reader with only the record and its cited items can explain, bound, re-derive and challenge a status.

- **How many questions and checks are proposed?** Ten questions and four mechanical checks.
- **Did the DSA statement-of-reasons schema pass?** It answered 4 of the 10 questions in full, 2 in part and 4 not at all, according to the paper.
- **Is it a legal standard?** No. The paper makes no deployment or legal claim.

Related reading

- Architecture decision records explained
- The imputation rule and the delegation record
- What is a blameless postmortem?

Sources and currency

Sources were read on 24 November 2026. Products, rules and standards change, so check the linked originals before you rely on any detail.

Related credential: Certified Technical Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. Record Sufficiency: What a Decision Record Must Preserve When an Unsettled Claim Becomes an Operative Status. arXiv 2611.23422, 17 November 2026 (preprint). arxiv.org/abs/2611.23422.

Published by IICTP, an independent technical credential body, on 24 November 2026. The current version of this post is at https://blogs.iictp.org/record-sufficiency-decision-records-ten-questions-explained. More guides: blogs.iictp.org.