



IDΔAC™ NEWSLETTER

The Approval Record Problem

Why a complete AI decision file
can still be weak evidence

17 July 2026 | Public thought brief

The record says the decision was reviewed by a human.

An AI-assisted screening system recommends moving a candidate out of the active recruitment queue. A recruiter reviews the case. The organisation records the approval.

On paper, human oversight occurred.

Later, the decision is challenged.

The organisation can show:

- the system recommendation;
- the reviewer's name;
- the final status;
- an approval timestamp.

But it cannot clearly establish:

- which version of the recommendation the reviewer saw;
- what case-level evidence was visible at the time;
- whether contradictory information was available elsewhere;
- whether the recruiter could pause, modify or refuse the route;
- whether escalation was practically usable;
- why the recommendation was accepted rather than challenged;
- whether the recorded rationale existed before the outcome became operational.

The file is complete enough to describe what happened.

It may still be too weak to show that the decision was defensible when it was made.

This is the approval record problem.

Approval is an event, not a conclusion

An approval record establishes that an interaction occurred. Someone clicked, signed, confirmed or closed a task at a particular time.

That is useful evidence. But it is evidence of an event, not proof of the quality or legitimacy of the judgement behind it.

The presence of an approver does not automatically establish that the person:

- understood how the AI-enabled recommendation had been formed;
- saw enough evidence to challenge it;
- retained realistic alternatives;
- had authority proportionate to the consequence;
- could escalate without making the route operationally unworkable;
- recorded their reasoning before hindsight changed the decision frame.

This is why the language of "human approval" can create false confidence. It compresses several distinct questions into a single reassuring label.

A person may have approved the action. That does not yet tell us whether meaningful human judgement remained live.

As discussed in The Human-in-the-Loop Illusion, a human can be present in a workflow after the decision has already become difficult to reopen.

Record completeness is not evidentiary sufficiency

Organisations often assess their decision records by asking whether the expected fields are populated.

Is there a recommendation? Yes.

Is there a reviewer? Yes.

Is there a timestamp? Yes.

Is there a final outcome? Yes.

From a process-completion perspective, the file may look satisfactory. From an evidentiary perspective, the relevant question is different:

Does the record preserve enough of the decision state to show what could be known, challenged and changed before commitment?

That requires more than administrative completeness.

Three distinctions matter.

Evidence existed versus evidence was visible

Information may exist inside the organisation without being available to the reviewer.

A candidate record may contain contradictory employment information, an accommodation request, a manager note or a corrected document. If those elements sit in another system, behind a different permission layer or outside the interface used for review, their existence does not prove that they informed the decision.

Under scrutiny, “the organisation had the information” and “the reviewer could use the information” are not equivalent claims.

Evidence was visible versus evidence was decision-relevant

Even where information was technically displayed, the route may not show whether it was understood or treated as relevant.

Interfaces organise attention. A prominent risk label, ranking or recommended action can dominate the frame, while qualifying evidence remains visually secondary. The reviewer may see several data points but still operate inside a route whose structure makes one option appear normal, expected or already justified.

The question is therefore not only what appeared on the screen. It is what the route made actionable.

A rationale exists versus the rationale was contemporaneous

A later explanation may be honest and accurate. It may reflect what the reviewer believes they considered.

But once the outcome is known, the evidentiary conditions have changed. The chosen option often appears cleaner. Doubts become easier to organise. Alternatives look weaker than they felt before commitment.

A rationale added later remains relevant, but it should not silently acquire the same status as reasoning recorded while the decision was still open.

Good decision traceability preserves that distinction.

The missing decision state

Most weak records do not fail because they contain nothing. They fail because they preserve the endpoints while losing the state between them.

They show:

Recommendation → Approval → Outcome

But the defensibility of the route may depend on what happened inside the missing middle:

Recommendation → Evidence access → Human interpretation → Challenge → Escalation → Authority transition → Commitment → Outcome

Not every routine decision requires an elaborate dossier. Proportionality matters. But where an AI-assisted route affects employment, access, pricing, prioritisation, restriction, eligibility, safety or another material consequence, the organisation may need to show more than the final administrative sequence.

It may need to reconstruct the decision state.

That means identifying what the reviewer knew, what remained uncertain, what alternatives were live, what intervention rights existed and when the organisation became committed to the consequence.

Five questions to test an approval record

Before treating recorded approval as evidence of effective oversight, an organisation can ask five questions.

1. Can we reproduce what the reviewer actually saw?

Not what the current system displays. Not what the organisation can assemble later. What was available to that reviewer, in that interface, at that time?

Versioning matters. Evidence access matters. Missing or unretained information matters.

2. Can we show that alternatives remained operationally available?

A theoretical ability to reject is weak if rejection required disproportionate effort, delayed the case beyond a deadline or triggered a process the reviewer could not realistically use.

3. Can we distinguish review from confirmation?

Did the reviewer test the recommendation against case-level evidence, or simply verify that the workflow had produced an expected result?

The record should not assume that human contact equals human challenge.

4. Can we identify where final authority sat?

The person completing the task may not be the person capable of legitimately assuming its consequence. If exposure exceeded the reviewer's mandate, the route should show whether authority moved and who accepted final commitment.

5. Can we explain why the route continued without escalation?

Non-escalation is not automatically a failure. But where warning signals, uncertainty or authority gaps were present, the absence of escalation can become important diagnostic evidence.

The organisation should be able to distinguish a reasoned decision not to escalate from a route in which escalation was invisible, unusable or culturally discouraged.

What a stronger record looks like

A stronger approval record does not need to reproduce every internal action or retain unlimited data.

It should preserve enough attributable evidence to connect:

- the AI-enabled signal or recommendation;
- the version and timing of what was reviewed;

- the evidence available to the reviewer;
- material uncertainty or missing information;
- the reviewer's practical intervention rights;
- any challenge, override or escalation;
- the authority responsible for final commitment;
- the rationale available before operational effect.

The objective is not bureaucratic accumulation.

It is reconstructibility.

A large document archive can still fail that test. A smaller, well-bounded record may pass it if it preserves the elements that made the decision attributable and challengeable.

For a broader public examination of evidence, authority and escalation, see the IDΔAC™ Resource Library, including the public brief on operational human oversight and decision-route traceability.

Before the file is challenged

The approval record problem usually becomes visible after the organisation has already lost control of the timing.

A candidate complains. A customer disputes an outcome. Internal Audit asks for evidence. Legal needs to understand who authorised the route. A regulator or executive sponsor wants to know where human judgement entered and whether it could still change the decision.

At that point, the organisation begins reconstructing.

Screenshots are collected. Emails are searched. Policies are cited. People explain what they normally do. A rationale is written in clearer language than the decision probably contained at the time.

The resulting account may be plausible. It may even be substantially correct.

But the evidentiary question remains:

Is the organisation demonstrating the decision route, or reconstructing a defensible story after the route has already been challenged?

That difference is precisely why decision exposure should be examined before an incident, dispute or formal review forces a late reconstruction.

What IDΔAC™ examines

IDΔAC™ Exposure Diagnostic is a bounded, 14-day, evidence-first diagnostic for selected AI-assisted decision routes.

It examines whether a route can be attributed, challenged, reconstructed and defended under scrutiny, with particular attention to evidence access, human judgement, escalation readiness, authority and final organisational commitment.

It is not legal advice, certification, conformity assessment or technical model validation.

Explore the Exposure Diagnostic Use the Resource Library www.iddac.eu · contact@iddac.eu