

VIKRAM JHA - DEFENSIBLE AI WORKSHEET

APPROVAL-RECORD SCORING WORKSHEET

Version 1.0 - 23 August 2026

<https://vikramjha.work/artifacts/approval-record-scoring-worksheet>

Score five real approval records out of four. The interesting number is not the total - it is which question fails across the set.

WHY IT EXISTS

This does not measure whether you have approvals. You do; everybody does. It measures whether an approval you already collected can answer 'who authorized this?' about a specific action, months later, to someone not inclined to be generous. A human genuinely approved something and an action genuinely happened - the question is whether the record binds the second to the first. Take the five most recent records, or five at random from the last quarter. Not your best five; that defeats the exercise.

WHAT IT CONTAINS

1. Four questions, one point each, applied to five real approval records
2. The distinction between naming an account and naming a principal
3. Why a stated boundary is the load-bearing question, and the one that fails most often
4. The difference between an expiry field and an annual review meeting
5. How to read the result by column rather than by row

THE FOUR QUESTIONS

Each is worth one point. Score five records. A single record scoring 2/4 tells you about one record; a column that fails five times out of five tells you about your system.

1. DOES THE RECORD NAME THE PRINCIPAL?

Passes when: It identifies who or what acted in a way that distinguishes it from anything else that could have acted - a specific service identity, a specific human, a specific delegation between them.

Fails when: It names a shared service account, a role, a team or an integration. If four things in your estate could have produced this record, it does not name a principal.

The common failure: An agent runs under a service account that predates it, shared with two batch jobs and a monitoring integration. The record is honest and useless.

2. DOES THE RECORD NAME THE SPECIFIC ACTION?

Passes when: It describes the action at the granularity a person would need to decide whether it was acceptable - which record changed, which payment released, which customer contacted.

Fails when: It names a capability, a category or a workflow. 'Agent executed customer-update workflow' tells you a class of thing happened.

The common failure: The test: could someone who was not there tell you whether the action was the right one? If they would have to go and find out what actually happened, it is a category.

3. DOES THE RECORD STATE THE BOUNDARY OF WHAT WAS APPROVED?

Passes when: It states what the approval does not cover - an amount ceiling, a set of accounts, a time window, a class of counterparty, a required condition.

Fails when: It records assent without extent. Someone approved; what they approved is inferred from context that is no longer available.

The common failure: An approval without a stated boundary expands to fit whatever the system later does under it. Nobody decides to widen it - there was never an edge to widen.

4. DOES THE RECORD STATE WHEN IT EXPIRES?

Passes when: The approval has an end: a timestamp, a run count, a terminating condition, or a review date with teeth.

Fails when: The approval is standing, or its expiry is a policy statement elsewhere rather than a field on the record.

The common failure: If the system stopped honoring this approval tomorrow, would anything break that should not? If yes, 'we review those annually' describes a meeting, not a control.

READING THE COLUMNS

Read the columns, not the rows. The column that fails across the set is the finding worth having.

Q1 - PRINCIPAL

What it means: You have an identity problem, not an approval problem.

What to do about it: Fixing the approval format will not help until agents have principals. Run the reachability audit first.

Q2 - ACTION

What it means: Your logging granularity sits above your decision granularity.

What to do about it: The approvals may be fine and you cannot demonstrate it.

Raise the record, not the policy.

Q3 - BOUNDARY

What it means: The common case: approvals record assent without extent.

What to do about it: Add a stated boundary to the record format before adding anything else.

Q4 - EXPIRY

What it means: You have standing authority nobody re-decides.

What to do about it: Pair with the time-to-safe-state protocol - standing authority and slow revocation compound.

HOW TO USE IT

1. Pull five approval records. The five most recent, or five at random from the last quarter - not your best five.
2. Score each against the four questions. About fifteen minutes per record once you have found them, and finding them is often the first finding.
3. Read the result by column. One record scoring badly is a record; a column failing five times is an architecture.
4. Do not expect 20/20. Very few production estates score above 2/4 on a random sample, and a team scoring 1/4 that knows which column failed is in better shape than a team that has never counted.
5. Remember what it does not measure: it scores the record, not the decision. A record can pass all four and describe an approval that was substantively wrong.

Built from public standards and general practice. No client information appears anywhere in this document. Regulatory instruments move - check anything cited here against its primary source on the day you rely on it.

Free to use and adapt; attribution welcome, not required.