

Decision Custody

The missing clock in AI accountability — a one-page brief

THE PROBLEM

AI systems now make and shape decisions that outlive them. Models are retired within a year or two; suppliers change hands; reference data is overwritten in place; the people who ran the system move on. The decision itself lives on — a refused licence follows a person for a career; a land title outlives everyone who touched it. In Europe, the logs of a high-risk AI system may legally be kept for as little as six months.

THE CONCEPT

Decision custody is an institution's demonstrated ability to reproduce, explain, defend, and answer for every consequential decision made in its name — human or machine — for as long as the consequences last.

THE TWO CLOCKS

Consequence clock	how long the decision matters — can be relied on, challenged, or cause harm.
Evidence clock	how long the institution can still show how the decision was made.
Custody gap	consequence horizon minus evidence horizon — an unfunded liability with no line on the balance sheet.

CONSEQUENCES — DECADES

EVIDENCE — MONTHS

the custody gap

THE FOUR TESTS OF CUSTODY

- 1) **Reproduce** — can the decision be re-derived from the record alone, without the system that made it?
- 2) **Explain** — can it be explained to the person it affected, in their language, without reference to the model?
- 3) **Defend** — would the record stand before a body with power over the institution, with the supplier absent from the room?
- last) **Answer** — is there a named office — an office, not a person — that will still answer at the end of the consequence horizon?

THE BOARD QUESTION

“If this decision is challenged ten years from now, who will answer for it — and what evidence will they have?”