

Decision Custody

Who Answers for a Decision After the
Machine That Made It Is Gone

*The institutional counterpart to decision sovereignty — and the
one question the accountability literature does not ask.*

The covenant is the promise; custody is the proof.

Sanjeet Kumar

Enterprise Architecture & Technology Governance Leader
Founder & Owner, Robotiq Technologies Ltd.

Sovereign Digital Resilience · Decision Custody · AI Governance

sanjeetkumar.com/papers · ORCID 0009-0003-1677-684X · 2026 · Views are the author's own.
DOI: 10.5281/zenodo.21959272

IN ONE PAGE

Ask a simple question of any organisation using AI to make decisions about people: if someone challenges one of those decisions three years from now, could you still show how it was made?

For most organisations the honest answer is no. Not because anyone did anything wrong — but because the technology that made the decision will be long gone. AI models are replaced within a year or two. Suppliers change hands. Reference data gets overwritten. The people who ran the system move on. Yet the decision itself lives on: a refused licence follows a person for their whole career, and a land title outlives everyone who touched it. In Europe, the law requires the logs of a high-risk AI system to be kept for as little as six months.

This paper gives that problem a name - **Decision custody**.

‘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.

In plain terms: making a good decision today is not enough. Someone must still be able to stand behind it in ten or twenty years — show what was decided, why, on what information, and under whose authority — long after the software, the supplier and the staff have all changed. Courts, archives and land registries have done exactly this with paper records for over a century. Nobody has yet asked it of the systems now making the decisions.

For a busy reader, the paper comes down to three things —

1. the custody gap — a simple measure of how far your evidence falls short of your obligations (Figure 2);
2. four plain tests that any decision record either passes or fails (§7); and
3. five questions a board can ask at its next meeting (§10). If the answers come back vague, that is the finding.

1. Two clocks that do not agree

When an institution makes a decision that matters — grants a licence, registers a title, approves a loan — two clocks start ticking at the same moment.

The first clock measures **how long the decision matters**. Call it the consequence clock. A licence granted today will still be relied on in twenty years. A refusal will still be shaping someone’s career in ten. A land title will still be the proof of ownership long after everyone who handled it has died. This clock runs for as long as anyone can still rely on the decision, challenge it, or be harmed by it.

The second clock measures **how long the institution can still show how the decision was made**. Call it the evidence clock. It runs only as long as the records survive and still make sense: the logs, the rules that were applied, the data that was used, and the memory of the people who were there.

For most of history, the two clocks kept the same time without anyone trying. The same paper file both recorded the decision and explained it. It sat in a registry or an archive, and nobody threw it away. If a question came fifty years later, someone fetched the file.

That is no longer true, and the gap is not small.

The consequence clock has not changed — decisions matter for exactly as long as they ever did. Licences, registrations, entitlements and credit decisions are the ordinary output of registries, regulators, licensing bodies, insurers, lenders and public authorities: the institutions that exist precisely because somebody needed decisions to outlast the people who made them.

The evidence clock, however, has collapsed. The AI systems now making or shaping these decisions are replaced within a year or two of being deployed. Suppliers merge or move on. Reference tables are overwritten in place. Configuration changes with every release. And a retired model rarely announces itself as a problem — the system keeps producing answers while the machinery underneath quietly becomes something else. Meanwhile, the legal minimum is being treated as a target: e.g., under the European AI Act, the logs of a high-risk system must be kept for as little as six months. Other jurisdictions do not even set a number: Canada’s proposed Artificial Intelligence and Data Act died on the order paper, and the United States regulates retention sector by sector — which leaves the six-month European floor as the strictest general rule anywhere.

Six months of evidence for a decision whose consequences run for sixty years is not a control. It is an accounting convention.

The regulator’s six months is written in **system time** — the lifetime of the software. The citizen’s question arrives in **consequence time** — the lifetime of the decision. An institution that clears the legal floor believes it has done something> which means, it has answered the wrong clock.

2. Custody is not a metaphor

Custody sounds like a borrowed metaphor. This section shows it is not: it is a working discipline, a century old, that has already answered the exact question, but now AI has reopened — **how do you stay responsible for something you no longer hold?**

The word being reached for right now, across security vendors and governance blogs, is custody. That reaching is not wrong, but it is shallow. It borrows the forensic image — the evidence signed from collection to courtroom — and stops there, producing another logging feature with a better name.

Custody has a deeper and far more useful history, and it belongs to archival science rather than forensics. Hilary Jenkinson formalised it in 1922; Theodore Schellenberg revised and redefined it at mid-century. In 1981 Gerald Ham challenged the settled conception directly, arguing in *The American Archivist* that the custodial posture — archivists as caretakers of what they physically hold — had become untenable as electronic records multiplied beyond any institution’s capacity to take them in, and naming what followed the post-custodial era. Terry Cook developed the argument through the 1990s. Frank Upward gave it structure with the records continuum model, which charts the many ways an entity can claim custody of a record it does not possess.

The important part, for our purposes, is what post-custodial theory did not say> It did not say custody ends. The professional literature is explicit that post-custodial is not non-custodial: the obligation persists but detaches from possession.

The institution remains answerable for records held elsewhere and discharges that obligation through oversight, specification, and involvement in how the records are made, rather than through physical control of them.

Nor did the idea stay theoretical. Digital preservation codified it: ISO 14721, the Open Archival Information System reference model, makes it a mandatory responsibility of any archive to obtain sufficient control of information to ensure its long-term preservation, and to keep that information independently understandable to a designated community — without the assistance of the people who produced it. The standard exists, is decades mature, and has simply never been read as applying to decisions.

That is an exact description of where every organisation now stands in relation to its decision machinery. It does not hold the weights. It does not hold the training data. It frequently does not hold the inference logs, the model version, the reference tables, or the vendor. It rents all of it, on terms and

timelines it does not set. Enterprise architecture has been treating this as a novel problem of AI governance. It is a well-worked problem in a neighbouring profession, and that profession reached a conclusion worth importing wholesale: losing possession does not discharge the obligation, it changes its form>.

This is also why ‘custody’ is the right word and traceability, auditability and observability are not. Those are properties a system either has or lacks. **Custody is a relationship between an institution and a record** — *one that has to survive changes of possession, of technology, of supplier and of staff, because that is the only condition under which it means anything.*

3. Defining decision custody

Here is the definition this paper stands on. The rest of the section takes it apart, one word at a time.

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.

Four verbs and one clause do the work, and each is chosen against a weaker alternative.

Reproduce

Not re-run. Re-running requires the model, and the model is the part guaranteed to disappear. Reproduction means that the recorded inputs, the applied rule, the operative context and the reasoning captured at the time are together sufficient to re-derive the decision without the system that originally produced it. A decision that can only be reproduced by the machine that made it is not in the institution’s custody; it is on loan from a supplier.

Explain

In the language of the person the decision was made about, not in the language of the system that made it. An explanation requiring the reader to understand model architecture is not an explanation, it is a deferral.

Defend

Before a body with power over the institution: a court, a regulator, an ombudsman, an auditor general, a legislative committee, a journalist with a document. Defence is adversarial by construction. A record that has only ever been read by its own authors has not been tested.

Answer for

Answerability is not explainability, and the two are routinely confused. A system can be perfectly explainable while no one is answerable for it — every step visible, every party pointing sideways. Answering for a decision means a named office accepts its consequences, including the consequence of having been wrong.

For as long as the consequences last

This is the binding clause and it is where the definition does its real work. Custody is indexed to consequence, not to retention policy, not to contract term, and not to the regulatory minimum. It is the clause that makes custody measurable, because a consequence horizon can be estimated in years and then compared against something.

One boundary should be drawn clearly. Decision sovereignty is a different and already-defined concept, developed principally in the context of state and military authority, which asks where the decision

boundary sits and whether a supplier has come to control it. That question is prior and it is not this one. Decision custody is the institutional counterpart: it asks what happens afterwards, and who is still standing to answer. Sovereignty concerns the moment of decision; custody concerns its whole life.

4. What sits beneath custody

Custody does not replace the existing accountability tools — logging, provenance, review. It sits above them and depends on them. Five layers are worth separating, because most institutional failures happen at the boundary between two layers rather than inside any one. Figure 1 shows the whole stack.

- **Layer 0 :: the record.** That a decision was made: when, by what, affecting whom. Almost every organisation has this, and almost every organisation mistakes it for the whole problem.
- **Layer 1 :: provenance.** How the decision was made: the chain of inputs, the nature of the processing, and the flow-on effects through the pipeline. This is well-founded ground. Singh, Cobbe and Norval set it out as decision provenance and it remains the strongest technical basis available for exposing decision pipelines in service of oversight, audit and redress. Nothing here competes with it; everything here assumes it.
- **Layer 2 :: reviewability.** Whether the decision, as a socio-technical process rather than a model output, can actually be examined by someone with standing to examine it. This layer is the subject of the next section, because it is the one most likely to be mistaken for custody.
- **Layer 3 :: custody.** Who is obliged to hold layers 0 to 2, for how long, against what future demand, and under whose signature. This layer is almost universally absent, because it is not a system property and therefore has no natural owner in a technology organisation.
- **Layer 4 :: covenant.** Why the institution exists and what it undertook to the people it serves. Custody without covenant is archiving for its own sake; covenant without custody is a promise with no receipt.

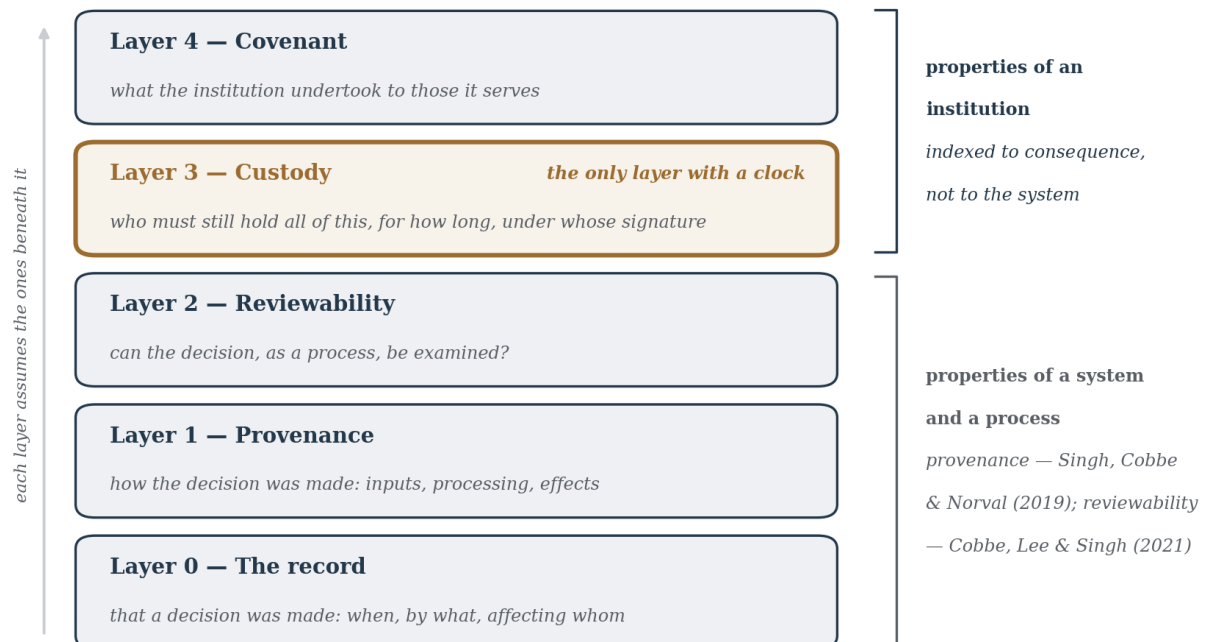


Figure 1 — The custody stack. Provenance and reviewability are properties of a system and a process; custody and covenant belong to the institution. Custody is the only layer indexed to time.

Provenance and reviewability are properties of a system and a process. Custody is a property of an institution, and it is the only one of the three with a clock in it.

That distinction is not academic. An organisation can hold a complete, signed, tamper-evident provenance record of a decision, structured for review by a competent regulator — and still have no custody of it.

This is because the pipeline belongs to a supplier whose contract runs three years against a consequence that runs thirty; because the log retention was set by a platform team optimising storage cost; because the model version identifier resolves to nothing; or because the reference data that made the record interpretable was refreshed in place and the prior state was never kept. Every one of those is a custody failure occurring in the presence of good provenance and adequate reviewability.

5. What reviewability asks, and what it does not

Readers who know the accountability field will have a question ready: has this not already been done, under the name reviewability? It deserves a direct answer, because it is the strongest objection available — and because the difference comes down to one word: time.

In 2021 Jennifer Cobbe, Michelle Seng Ah Lee and Jatinder Singh proposed reviewability as a framework for the accountability of automated decision-making. They made the right move and they made it first. They treated automated decision-making as a socio-technical process involving both human and technical elements, beginning before a decision is made and extending beyond the decision itself. They broke that process into its technical and organisational parts in order to determine what record-keeping each part contextually requires. And they drew the method from administrative law's approach to reviewing human decision-making, which is exactly the right place to have looked, because administrative law has been answering "was this decision properly made" for a century longer than computer science has existed.

Anyone building the record layer this paper depends on should start with that work rather than with this. I am not proposing an alternative to it. I want to be exact about what remains, because it is one thing and not a list.

Reviewability asks whether a decision can be examined. Custody asks how long it will remain examinable — and who is obliged to keep it so.

The gap is temporal, and it runs through the whole field rather than through one paper. The literature on algorithmic accountability is, almost without exception, written in the present tense. A framework is satisfied at the moment of assessment: the records are adequate, the process is decomposable, the review can proceed. Nothing in that assessment tells you whether the same will be true in year twelve, after two supplier generations, a platform migration, a reference-data refresh, and the retirement of everyone who was in the room. Reviewability has no clock, and it does not claim to. Custody is the clock.

Three consequences follow, and each is a place where an institution with excellent reviewability still fails.

- **Duration is unowned.** Reviewability determines what to record. It does not determine for how long, because "how long" is not a property of the process — it is a property of the consequence, which sits outside the system entirely and is usually known only to the general counsel.
- **Decay is invisible.** A record that was reviewable on the day can cease to be reviewable without anything being deleted. Reference data refreshed in place, a schema retired in a migration, a configuration treated as settings rather than as evidence — none of these is a records incident, and all of them end reviewability silently. The preservation field has a name for the property being lost:

ISO 14721 requires preserved information to remain “independently understandable” to its designated community, without the help of those who produced it — a test most decision records would fail within two supplier generations.

- **Answerability has no successor.** A review needs someone to review. Over a long enough horizon> the reviewing party, the accountable party, and the institution itself will change. Reviewability assumes a standing forum; custody has to survive its dissolution.

Custody, then, is not a rival framework. It is the question you ask after reviewability has been answered well: for how long does that answer hold, and whose name is against it.


6. The first instrument: the custody gap

The first instrument turns custody from a principle into a number.

The custody gap is the difference between the consequence horizon of a decision and the evidence horizon of the system that produced it.

The consequence horizon is how long the decision can still be relied upon, challenged, or cause harm. It is a legal and social question, not a technical one, and it is usually answerable by a general counsel in an afternoon.

The evidence horizon is subtler, and this is where most assessments go wrong. It is not the log retention period. It is the shortest of every dependency the record has: log retention, contract term, supplier solvency, model or version availability, reference and lookup data as it stood, configuration versioning, key and signature validity, format and schema readability, and the institutional memory of the people who could interpret any of it. The evidence horizon is a minimum function over that list — whichever dependency expires first sets it — and an institution that measures only the first term is measuring the one it happens to control.



illustrative — the licensing case of §8

Figure 2 — The custody gap. The evidence horizon is a minimum over every dependency the record has — and the shortest is often the one nobody is watching.

Set the two side by side, per decision class — Figure 2 shows the shape for a single decision; the table below shows the portfolio view — and the exposure becomes visible without any new instrumentation.

Decision class	Consequence horizon	Evidence horizon	Custody gap
Registration of a legal interest in land	Indefinite †	AI Act floor: 6 months †	Unbounded
Credit decision, regulated lender	25–30 years	5–7 years (MiFID II, SEC 17a-4) †	~20 years
Professional licensing or accreditation	A working life	AI Act floor: 6 months †	Decades
Entitlement or benefits determination	7–15 years	6–24 months	~6–14 years
Technical documentation (any high-risk system)	—	10 years after withdrawal (AI Act Art. 18) †	n/a — not a decision record
Content moderation action	6–24 months	6 months †	~0–18 months

† Sourced. The six-month figure is the deployer retention floor under the EU AI Act (Articles 19 and 26(6)); the ten-year figure is the technical-documentation obligation under Article 18, which attaches to documentation about the system rather than to the record of any decision — a distinction easily lost in a compliance mapping and expensively lost in a hearing. Financial-services retention reflects sector obligations under MiFID II and SEC Rule 17a-4. Unmarked figures are indicative and should be replaced with your own.

Every organisation’s own version of this table can be built in a workshop, and the workshop usually does more governance work than the artefact it produces, because it is the first time the legal horizon and the retention setting have been spoken aloud in the same room.

A custody gap is not a compliance finding. It is an unfunded liability with no line on the balance sheet.

Two properties make the gap useful as a management instrument. It is unavoidably quantitative, so it cannot be closed by assertion. And it is asymmetric: closing it is cheap at design time and ruinously expensive at challenge time, which is exactly the profile boards are equipped to reason about once it is put in front of them in those terms.

7. The second instrument: the four tests

The second instrument turns the definition into an examination. Each verb becomes a test with a stated failure condition, and the tests escalate — an institution failing an early one cannot pass a later one, whatever its documentation claims.

Test 1 — Reproduction

Can this decision be re-derived from the record alone, without the system that made it? Failure condition: reproduction requires a model version, an API, a live reference table, or a supplier that the institution does not control.

Test 2 — Explanation

Can it be explained to the person it affected, in their language, without reference to the model? Failure condition: the explanation is an architecture diagram, a confidence score, or a sentence containing the word “algorithm”.

Test 3 — Defence

Would this record stand before a body with power over us, with the supplier absent from the room? Failure condition: the defence depends on a third party’s cooperation, continued existence, or willingness to disclose material it considers proprietary.

Test 4 — Answerability

Is there a named office — an office, not a person — that will still exist and still answer at the end of the consequence horizon? Failure condition: the accountable party is an individual, a team name, a project, or a supplier.

The fourth test is the one that fails everywhere, and it fails for a structural reason rather than a negligent one. Individuals retire, teams reorganise, projects close, and suppliers are acquired; only offices are built to outlive their occupants. Institutions that have carried decisions across centuries — registries, courts, archives, notarial systems — did it by vesting answerability in an office and then keeping the office. That competence still exists inside most public bodies. It has simply never been asked of the software estate.

8. A constructed case: the licence that could not be explained

The following case is constructed. The failure pattern it describes — reference data drifting through a migration, a rule in force that was never versioned, the people who could interpret it gone, and no office defined to answer — is one I have met in more than one form across public-record, financial services and government programme environments. The institution and the particulars are invented, so that the reasoning can be set out without identifying anyone.

An authority licenses practitioners in a regulated profession. Its determinations are consequential in the ordinary sense — a refusal shapes a career, and both grants and refusals are relied upon by employers, insurers and courts for as long as the person practises. The consequence horizon is a working lifetime.

Application volumes rise. The authority introduces automated pre-assessment: applications are scored against published eligibility criteria and routed into clear-grant, clear-refuse and manual-review streams. The programme is governed to a standard most organisations would envy. There is a model card. Every refusal is confirmed by a human assessor. Logging is designed deliberately and retained well beyond the regulatory minimum. An internal audit finds nothing material.

Part of the scoring depends on an external equivalency reference — a table mapping foreign qualifications to domestic ones, maintained by a third party and refreshed on a published schedule. This is a sensible choice. Maintaining it in-house would be duplicative and worse.

Roughly two years later, an applicant who was refused writes to ask why. There is no lawsuit. She has since been licensed in another jurisdiction and simply wants to understand what happened. The authority sets out to reconstruct the determination and finds the following, in this order.

- **The record exists**, and it is good. The application, the timestamp, the inputs, the score, the routing decision and the assessor who confirmed the refusal are all present and intact.
- **The inputs no longer mean what they meant.** The equivalency reference has been refreshed eleven times since. The supplier publishes current state; it does not archive prior versions, and nobody asked it to. The qualification codes in the record are still legible, but what they mapped to on the day is gone. The authority can see exactly what was fed in and cannot say what it signified.
- **The threshold in force cannot be established.** Scoring configuration was adjusted twice during a platform migration eighteen months ago. Both changes were raised, tested, approved and deployed correctly. Neither was versioned in a way that binds a configuration to a date, because configuration was treated as settings rather than as evidence.
- **The assessor has retired.** Her successor can describe the process as it operates now. He cannot attest to how it operated then, and would be wrong to try.

- **No office owns the answer.** The authority has a registrar function, competent and long-standing — but its remit covers the licence, not the assessment that precedes it. The question of who answers for a pre-assessment made two years ago was never asked, because until someone asked it, it had no owner.

While assembling this, someone notices that the same conditions apply to every determination made in a fourteen-month window: roughly four thousand of them. That discovery, not the original letter, is the expensive part.

In the absence of evidence, the resolution is what resolutions look like. The authority reviews the affected cohort manually at significant cost, reverses a small number of determinations, reports the matter to its oversight body, and absorbs the finding. There is no misconduct anywhere in the story. Every control performed as designed. What failed was the assumption that a record of a decision is the same thing as the ability to account for it.

Note where the cause actually sits. Nothing in the failure originates in the pre-assessment capability, which was competently built and competently governed. It originates in two architecture decisions taken earlier and elsewhere: to consume reference data live rather than to snapshot it with the decision, and to treat scoring configuration as deployment settings rather than as evidence. Both were reasonable at the time by every criterion then applied. Neither was ever asked how long the decisions downstream of it would need to be defensible — which is the whole of the argument in one sentence.

What would have closed it was neither exotic nor expensive at the time. A consequence horizon stated for the decision class at design, which would immediately have shown a retention setting to be irrelevant to a career-length exposure. A requirement that external reference data in force be snapshotted with the decision rather than referenced live. Configuration versioned as evidence rather than as settings. And last) an office named, with its remit extended to cover the assessment and not only the licence. Four choices at design time, none of them costly, in place of four thousand manual reviews.

9. Why compliance will not produce custody

It is tempting to expect regulation to close this, and it will not, for three reasons that are structural rather than a matter of the rules being immature.

First, regulation is written in system time. The European framework, whose high-risk obligations began applying in August 2026, requires automatic logging across a system's lifetime and sets a deployer retention floor of six months. That is a coherent period for post-market monitoring and substantial-modification detection, which is what the provision is for. It is not, and does not claim to be, a period appropriate to answering a citizen in year twelve.

Second, accountability does not travel with the procurement, and the supply chain makes this worse rather than better. Buying a system from a supplier does not transfer the obligation to the supplier; the deploying institution remains answerable for retention, accessibility and the consequences of the output regardless of who built it. Cobbe, Veale and Singh have shown why this is now structurally hard: algorithmic systems are produced, deployed and used within supply chains of multiple actors tied together by flows of data, so that outcomes are driven by the working together of the chain rather than by any one party. Accountability frameworks built for a bilateral relationship between an actor and a forum do not capture that. Custody therefore always weighs heavier on the buyer than on the seller, and no contract written on standard commercial terms will correct that imbalance unless someone specifically asks it to.

Third, and most consequentially, compliance is assessed at a point in time and custody is tested at a point of stress. Nothing in an annual compliance sign-off reveals whether a record will still be interpretable after two supplier generations, a platform migration and a reorganisation. The only

honest test is the adversarial one, which is why the four tests are framed as questions an outsider would ask rather than controls an insider would evidence>.

10. What this asks of a board

Custody is a board matter rather than a technology matter, because it concerns liabilities that mature after the current management is gone. The questions below are deliberately answerable — each has a number, a name, or a date as its answer, and an evasive response is itself the finding.

FIVE QUESTIONS FOR A BOARD

- 1) For our three highest-consequence decision classes, what is the consequence horizon, stated in years?
- 2) What is the evidence horizon for each, and which of its components — retention, contract, supplier, model, reference data, configuration, format, or memory — is the shortest?
- 3) If our principal model supplier ceased to exist on Monday, which decisions made this year could we still defend on Tuesday?
- 4) Which named office answers for a decision made today, in year fifteen — and does that office appear on an organisation chart we control?
- last)** What custody gap are we choosing to accept, and whose signature is against that choice?

The final question is the one that changes behaviour, because it converts an omission into a decision. Gaps consciously accepted and signed for are a legitimate posture; an authority may reasonably decide that a low-value routine classification does not warrant career-length custody. Gaps that exist because nobody computed them are something else, and they have a way of surfacing in the same quarter as everything else that was never computed.

11. The proof of the promise

Institutions have always known how to do this. Land registries carried determinations across regime change and fire. Courts produced records that could be read and relied upon by people separated from the judge by generations. Notarial systems solved answerability across borders and lifetimes with nothing but paper, a seal and a well-designed office. The competence is not new and it was never lost. It has simply not yet been asked of the systems now making the decisions.

What is new is the rate at which the machinery underneath turns over, and the ease with which an institution can mistake a supplier's memory for its own. Every model migration, every reference-data refresh, every configuration change is a small transfer of custody, and most of them are executed as change requests>. Enough of them, executed that way, and an organisation arrives at the position of being unable to account for its own conduct — not through misconduct, but through a series of individually reasonable technical decisions, each correctly documented, none of them asked the only question that mattered.

A service covenant states what an institution exists to do and what it undertook to the people it serves. Custody is what makes that undertaking checkable by someone who was not there and has no reason to take it on trust. One without the other is either a slogan or a filing cabinet.

The covenant is the promise; custody is the proof.

References

1. Cobbe, J., Lee, M. S. A., & Singh, J., “Reviewable Automated Decision-Making: A Framework for Accountable Algorithmic Systems”, Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’21), 598–609. doi:10.1145/3442188.3445921
2. Cobbe, J., Veale, M., & Singh, J., “Understanding Accountability in Algorithmic Supply Chains”, Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’23), 1186–1197. doi:10.1145/3593013.3594073
3. Cook, T., “Electronic Records, Paper Minds: The Revolution in Information Management and Archives in the Post-Custodial and Post-Modernist Era”, Archives and Manuscripts, 1994; reprinted Archives & Social Studies 1 (2007), 399–443.
4. European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 12, 18, 19 and 26(6). High-risk obligations applicable from 2 August 2026.
5. Ham, F. G., “Archival Strategies for the Post-Custodial Era”, The American Archivist 44, no. 3 (Summer 1981), 207–216.
6. International Organization for Standardization, ISO 14721 — Open Archival Information System (OAIS) Reference Model. First published 2003; current edition 2025.
7. Jenkinson, H., A Manual of Archive Administration, Oxford University Press, 1922.
8. Kaminski, M. E., & Urban, J. M., “The Right to Contest AI”, Columbia Law Review 121, no. 7 (2021).
9. Singh, J., Cobbe, J., & Norval, C., “Decision Provenance: Harnessing Data Flow for Accountable Systems”, IEEE Access 7 (2019), 6562–6574. doi:10.1109/ACCESS.2018.2887201
10. Society of American Archivists, Dictionary of Archives Terminology, entry: postcustodial.
11. Upward, F., “Structuring the Records Continuum, Part One: Post-Custodial Principles and Properties”, Archives & Manuscripts 24, no. 2 (Nov 1996), 268–285.
12. Wei, P., & Shu, W., Preserving Decision Sovereignty in Military AI: A Trade-Secret-Safe Architectural Framework for Model Replaceability, Human Authority, and State Control, arXiv:2604.20867, 2026.
13. Kumar, S., The Intent Bottleneck: When AI Compresses Delivery, Strategy Becomes the Constraint, Zenodo, 2026. doi:10.5281/zenodo.21943398

AI-assistance disclosure > This paper was drafted and refined with AI assistance. The arguments, the definitions, the two instruments and the constructed case are the author’s own, and all sources cited above were verified against their primary records before publication.