RED BEAR: ENTERPRISE DECISION INTELLIGENCE · DATA SHEET · B · RECORD AND EVIDENCE

One record of how the business actually decides

Record the judgement, not just the transaction it produced.

KEY POINTS

- One canonical, append-only entry per decision, whoever or whatever made it: question, options, choice and grounds, commitment
- The policy version, the dial and the evidence as they stood at the time, so any decision can be reconstructed rather than recalled
- A human decision and an agent decision are the same kind of object, in the same register, with the same fields
- Residency and retention tagged at ingest and never re-derived; consent recorded; what is not decision content discarded

What it is

Every large organisation has a system of record for money, for people, for customers and for inventory. None has one for judgement. The decision that moved any of those numbers has a folder of minutes, a thread of emails and somebody's memory of the call.

The central decision register is that missing system of record. It holds one canonical, append-only, versioned entry per decision, covering what was decided, between which options, on what reasoning, against which policy, at what dial setting, and what happened next. It is the product's centre of gravity: capture routes fill it, policy evaluation writes into it, every intervention appends to it, and all analysis reads from it. The register does not decide anything. It records what was decided.

The pain point

Told they need a record of decisions, most organisations reach for logging: audit trails in the ERP, sign-off histories in the workflow tool, retained mailboxes. **A log tells you what happened. A register tells you what was decided.** A log records that a user with a particular identifier changed a field to a particular value at a particular time. A register records that a named person or a named agent chose this option over those other three, for this stated reason, with this much at stake, under the policy in force that morning, and later, whether it worked. Only the second answers the question anyone asks a year afterwards, which is never "what changed?" but always "why did we do that, and who decided?".

That was tolerable while decisions were made slowly, in rooms where the reasoning was at least spoken aloud. It stops being tolerable when a pricing engine adjusts a rate overnight and an agent settles a refund at three in the morning.

How it works

What one record holds. A register entry is the decision, in the form a reviewer needs once the people involved have moved on. The question that admitted more than one answer, the options genuinely available, including the option to do nothing, the choice and the grounds for it in the decider's own words at the time, and the commitment that followed. The decision class, which sets the evidence standard, the memory standard, the context depth and the access granularity. The value at stake. Each participant named as a person, an application or an agent, with a role: Recommend, Decide, Own, Execute or Review, and Own held by exactly one named person. Every policy tested, the result, and the version in force at that moment. The dial band the decider held and the review obligation it carried. The evidence as an immutable snapshot addressed by digest, never a live link. And the outcome, once it is known, attached, never inferred.

Three properties that make reconstruction possible. *Append-only:* actions add events; nothing rewrites them. An editable record tells you what the organisation currently believes it decided; only an append-only one tells you what it decided at the time, in what order, and with what known to whom. *Policy as of the time:* the record keeps the version of each policy in force when the decision was made, not the version in force when somebody later comes to look. Without it, reconstruction is revisionism; with it, a reviewer can separate two questions: was the decision made properly under the rules that applied, and were those rules any good. *Equal footing:* a human decision and an agent decision are the same entry, the same fields, the same evaluation. The interesting questions cross that boundary: what share of a class is now settled by agents, whether they breach policy more often than the managers they relieved, and who overrode an agent, on what reasoning.

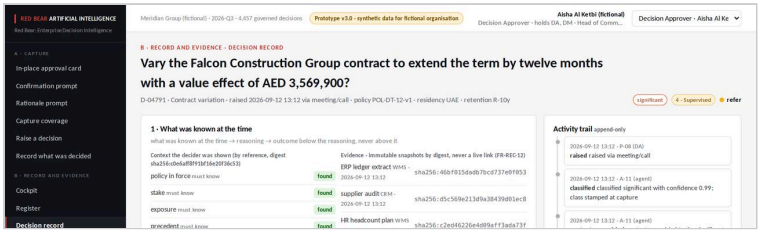
Governing a class, not a department. The unit of governance is the decision class: a recurring, consequential decision described once, with an accountable owner and its standards. A department is a thing only people belong to; a class is a thing a person, an application and an agent can each do, which is what lets one register and one policy serve all three.

Clean, respectful, provable. One person or agent is recognised across channels; one decision seen in three merges into one record with three sources; missing context is flagged rather than left blank. Each record is tagged at ingest with its residency boundary and retention class, and an aggregation that spans boundaries returns a partitioned result and says so. Retention runs by class, legal holds override deletion, e-discovery export carries a manifest whose digest matches, and controls are mapped to the EU AI Act's human-oversight requirement, ISO/IEC 42001 and the NIST AI RMF. Mapping is not certification, and Red Bear holds none.

What you get

Capability	What it does	Requirements
Taxonomy, classification and class standards	The class catalogue, extended by the customer; each candidate classified with a confidence, held for a person below the threshold; the four standards each class drives	FR-REC-01, FR-REC-06
Identity resolution and completeness	One actor across channels; one decision seen several times merged into one record; missing context flagged before release	FR-REC-02, FR-REC-03
Consent, redaction, residency and retention	The decision captured, not the conversation; tagged at ingest and never re-derived	FR-REC-04, FR-REC-05
Cockpit, register and class views	Every aggregate figure clickable, landing on a filtered register whose total equals the figure	FR-REC-07
Evidence vault and audit replay	Append-only, tamper-evident storage; any decision reconstructed as it stood, with the policy version, dial and context the decider saw	FR-REC-08, FR-REC-09, FR-REC-12
Retention, legal hold, e-discovery and regulatory mapping	Holds that override deletion; reproducible export with a manifest; producible evidence per control	FR-REC-10, FR-REC-11

Screens



S-11 Decision record: One decision record in full: reasoning, options, policy evaluation, roles, outcome and the append-only activity trail. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

The register is the first thing to switch on and the only one worth switching on alone. Through Connect and Observe it fills without enforcing anything or changing how anyone works, and an organisation that can see how a class of decisions is actually made is better placed to argue about enforcing policy than one working from recollection. **Risk and audit** get any decision reconstructed with the policy of the day, rather than an argument about what the rule used to say; the measure is time to reconstruct. **CFO and COO** see where authority actually sits, class by class, against where it was supposed to sit. **The Chief AI Officer** gets agent decisions in the same register as human ones, so autonomy can be argued from evidence. **Class owners** get one honest denominator before anyone proposes changing how a class is governed. **Legal** gets the decision and its evidence, not somebody's mailbox. None of these measures carries a target here; targets come from the pilot's own baseline.

Platform engines

Decision Intelligence for the record and classification (RBP FR-DI-01, RBP FR-DI-02, RBP FR-DI-06); the Provenance (RBP FR-PV-01), Privacy (RBP FR-PA-02, RBP FR-PA-04) and Audit (RBP FR-AU-01 to RBP FR-AU-03) services. Platform behaviour is specified in the platform FSD, not here.