

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Enterprise Architecture Decision Intelligence

Decisions with a memory: impact before you decide, lessons after

Architecture decisions are taken, forgotten and repeated. This paper explains what decision intelligence means for enterprise architecture, why a decision record is worth little without its impact estimate and its outcome, and how Red Bear EOM discovers, estimates, routes, records and learns from every decision that changes what a federation is or how it works.

Empower. Trust. Verify.

White paper EOM-WP-05 · Red Bear white paper series · Version 0.1

26 September 2026

Summary

A decision is only as good as what was known when it was taken, and only as useful as what was learned after. Most architecture decisions leave a trace in minutes and a consequence in the landscape, with nothing connecting the two. Decision intelligence connects them. It means knowing what a proposed change will affect before it is approved, on cost, risk, governance, compliance and scope; routing it to the person or board with the authority to decide it; recording the options, the choice and the reasons against the evidence; and, later, reviewing the outcome and turning it into a lesson the next decision can use. This paper describes the decision lifecycle Red Bear EOM implements, the impact estimate on five dimensions, the dependency trace that grounds it in the ontology, the alignment questions an Entity EA can ask before submitting anything, and the standing rule that AI proposes and a named person decides.

WHO SHOULD READ THIS

- Board members and chief architects who approve architecture change
- Entity architects who submit changes and want them approved first time
- Programme and portfolio offices that need decisions to stick
- Assurance and audit functions

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Describe the decision lifecycle from discovery to lesson
- Read an impact estimate and know what each dimension rests on
- Ask an alignment question and interpret the verdict
- Explain why an AI-drafted recommendation cannot become a decision on its own

1 What it is

Decision intelligence for enterprise architecture is the discipline of taking architecture decisions with the impact known, the authority right, the reasons recorded and the outcome learned from. It rests on four things a repository does not usually hold: an estimate of what a change will affect before it is approved; a routing rule that sends the change to whoever has the authority to decide it; a record that ties the choice to its options, reasons and evidence; and a review, later, of what actually happened.

The sources give each of these a foundation. TOGAF supplies the architecture decision record and the governance apparatus around it: compliance reviews, dispensations, architecture boards at global, regional and local level, and change classification (TOGAF Standard, EA Capability and Governance; Architecture Content §7.5). Dietz and Mulder supply the rule that decides who decides: a change to the essential model, to what the organisation is responsible for and who is accountable, is governance; a change to how that is implemented is management; and a governance change is decided at the level its impact reaches (Dietz and Mulder, Ch 9 §9.6.2–9.6.3). Their ship-on-route metaphor puts it simply: governance sets the route, management keeps the ship on it (Dietz and Mulder, Ch 9 §9.6.1). Greefhorst and Proper supply the principle side: principles as controls that a decision either conforms to, needs a dispensation from, or changes (Greefhorst and Proper, Ch 3, Ch 5).

Red Bear EOM adds the ontology as the ground the estimate stands on. A proposed change touches elements in the landscape and functions in the essence; the impact analysis walks the landscape and the dependency trace walks the essence, and between them they produce the affected systems, entities, functions, principles, standards and goals with a rule and evidence for every step. The impact estimate is computed from that walk, on five dimensions, each with an uncertainty statement, and it is attached to the decision before anyone is asked to approve it.

The lifecycle is: **discovered** (in documents, minutes or messages, as a candidate), **recorded** (context, options, choice, reasons, evidence), **estimated** (impact on five dimensions), **routed** (by the two labels to the approver in the delegated-authority table), **approved or declined**, **tracked** (to the work packages that implement it), **reviewed** (outcome against expectation) and **learned from** (lessons attached to the vendor, pattern or capability). Every state change carries a named person.

Key points

- Impact before approval, on five dimensions with uncertainty stated
- Routing by the two-label rule: essential-model changes are governance, decided at the level their impact reaches
- A record that survives: options, reasons, evidence, approver, outcome, lessons

2 Why it matters

Architecture decisions are expensive to take and cheap to lose. A board spends an hour on a platform choice; the minutes record the choice and none of the reasoning; two years later the same board, with different members, revisits it with no memory of the options rejected or the assumptions that have since failed. Meanwhile the decisions that mattered most were never seen by any board, because they were taken as IT changes: a shared register migrated, a transaction's executor moved to a vendor, a decline path automated away.

The Post Office case is the cautionary tale for the last of these. Decisions about how system evidence would be treated, about which transaction states were reachable and who could revoke, were implementation choices that turned out to be essential ones, and no governance body saw them as such (Dietz and Mulder, Ch 10). The childcare allowance case shows the same: a design in which decline, reject and revoke were unreachable in practice was a governance failure disguised as a system configuration (Dietz and Mulder, Ch 12 §12.3). The two-label rule exists so that such changes are recognised for what they are before they are approved.

There is also a quieter cost: the decision taken well, on good evidence, that nobody can find. Its lesson is lost. The next entity facing the same vendor, the same pattern, the same consolidation, starts from nothing. Lessons attached to the things they concern, and surfaced when those things appear in a new decision, are how a federation stops paying for the same education twice.

Failure	What decision intelligence changes
Choice recorded, reasoning lost	Options, reasons and evidence stored with the record; the rejected options as visible as the chosen one
Essential change approved as an IT ticket	The two-label rule classifies it; the delegated-authority table routes it; the trace shows why
Impact discovered in delivery	Impact estimate and dependency trace attached before approval; scope-completeness dimension lists what the estimate could not cover
Same mistake in the next entity	Outcome reviews produce lessons attached to vendors, patterns and capabilities, surfaced in the next decision that touches them
AI recommendation acted on as if decided	A recommendation has no effect until a named person accepts it; the acceptance is the decision, with its record

3 The challenges

Challenge	Why it is hard	What goes wrong without a method
Discovering decisions	They are taken in meetings, threads and documents, in many entities, and rarely labelled as decisions	The record holds only what someone remembered to file
Capturing options and reasons	The winning option is written up; the losers and the reasons are not	Revisited decisions start from zero
Maintaining the link to the landscape	A decision implemented over two years drifts from what was approved	Nobody can say whether the decision was executed
Understanding the impact of change before deciding	Impact spans cost, risk, governance, compliance and scope, across entities, with uncertainty	A single-number estimate that hides what it does not know
Driving standardisation of the decision process	Each board has its own template and threshold; consistency across entities is rare	Decisions are comparable only by accident
Knowing who decides	Authority differs by decision class and by tenant; some groups delegate to the Group EA, others to boards	The wrong body decides, or none does

4 How Red Bear EOM helps

Challenge	What EOM does	Module
Discovering	Claude reads minutes, strategies and opted-in Teams channels and proposes candidate decisions with the passage; an analyst reviews each before it becomes a record	DI Decision intelligence (FR-DI-01)
Capturing	The decision record extends the ADR with options, status lifecycle, outcome review and lessons; every field cites its evidence	GV (FR-GV-06, modified in v1.1)
Maintaining	A decision is linked to the work packages that implement it and to the elements it changes; delivery updates its status; drift between approved and delivered is flagged	DI, RM Roadmap
Impact of change	Impact analysis over the landscape and the dependency trace over the essence produce the five-dimension estimate: cost with basis, risk by class, governance level with the principles touched, compliance with the standards affected, and scope completeness listing what is uncovered; each with an uncertainty statement; not marked complete until scope items are resolved	AN Analysis, DI (FR-DI-03), VZ (FR-VZ-09)
Standardisation	One record, one estimate format, one routing rule across every entity; alignment questions let an Entity EA test an initiative before submitting it	DI (FR-DI-04), RP (FR-RP-05)
Who decides	The two-label rule; the delegated-authority table configurable per tenant (board or Group EA per class); the approver named on the record; the audit trail	GV (FR-GV-02, FR-GV-03), RA Roles

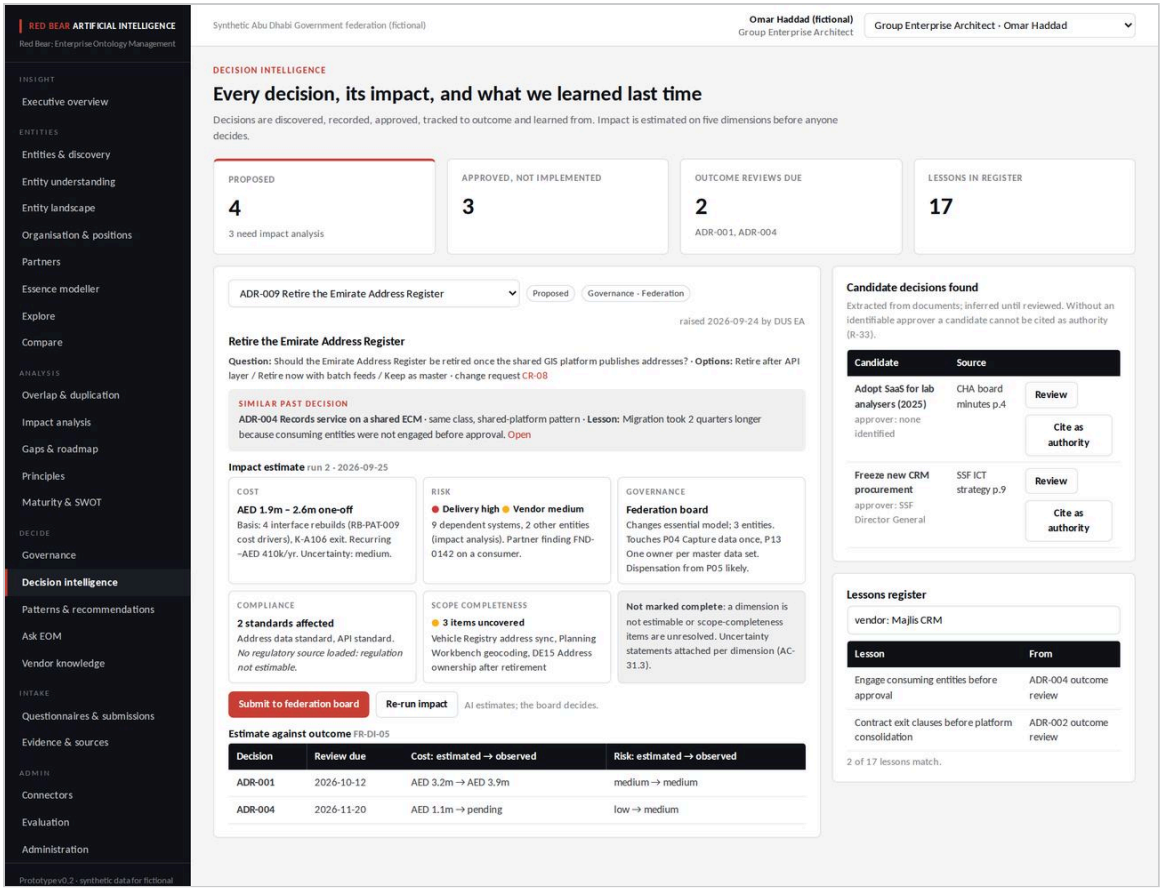
Standing rule

AI proposes. A named person decides. Claude can discover a candidate decision, draft the record, estimate the impact and recommend an option. None of that has any effect on the ontology, a decision or a pattern until a Group EA or the named approver accepts it, and the acceptance is what is recorded.

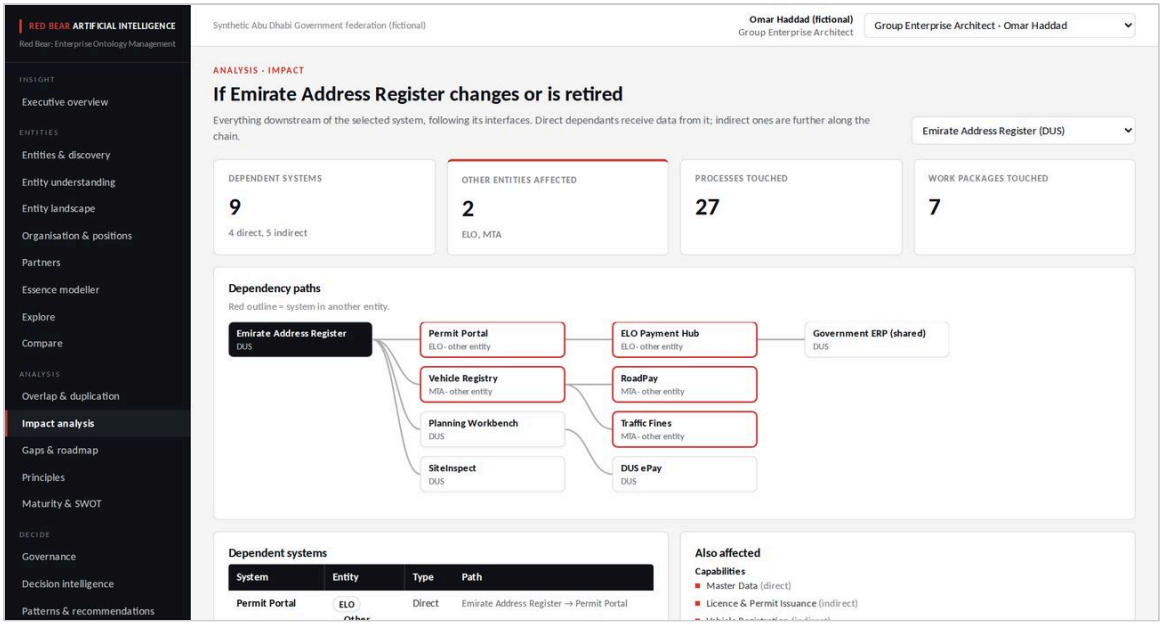
5 How it is used

Step	Who	What happens
1 Ask before submitting	EEA	On Ask EOM or the decision screen, the Entity EA asks whether an initiative, design or change aligns with the group; the verdict per dimension (principles, patterns, roadmap, overlap, standards) comes with evidence and what would need to change
2 Record	EEA, CON	The decision is recorded with context and options; Claude drafts from the submission; the Entity EA edits; evidence attached
3 Estimate	EOM, EEA	Impact analysis and the dependency trace run; the five-dimension estimate is produced with uncertainty statements; scope-completeness items listed; the Entity EA resolves or acknowledges each
4 Route	EOM	The two labels are set; the delegated-authority table names the approver; the decision enters that queue
5 Decide	Board or GEA	The approver sees the record, the estimate, the trace, the lessons from earlier decisions touching the same vendor, pattern or capability, and approves, declines or returns with questions
6 Track	EEA	Work packages linked; status moves as they deliver; drift flagged
7 Review and learn	GEA, EEA	At the review date the outcome is recorded against the estimate; lessons written and attached to the things they concern; surfaced in the next relevant decision

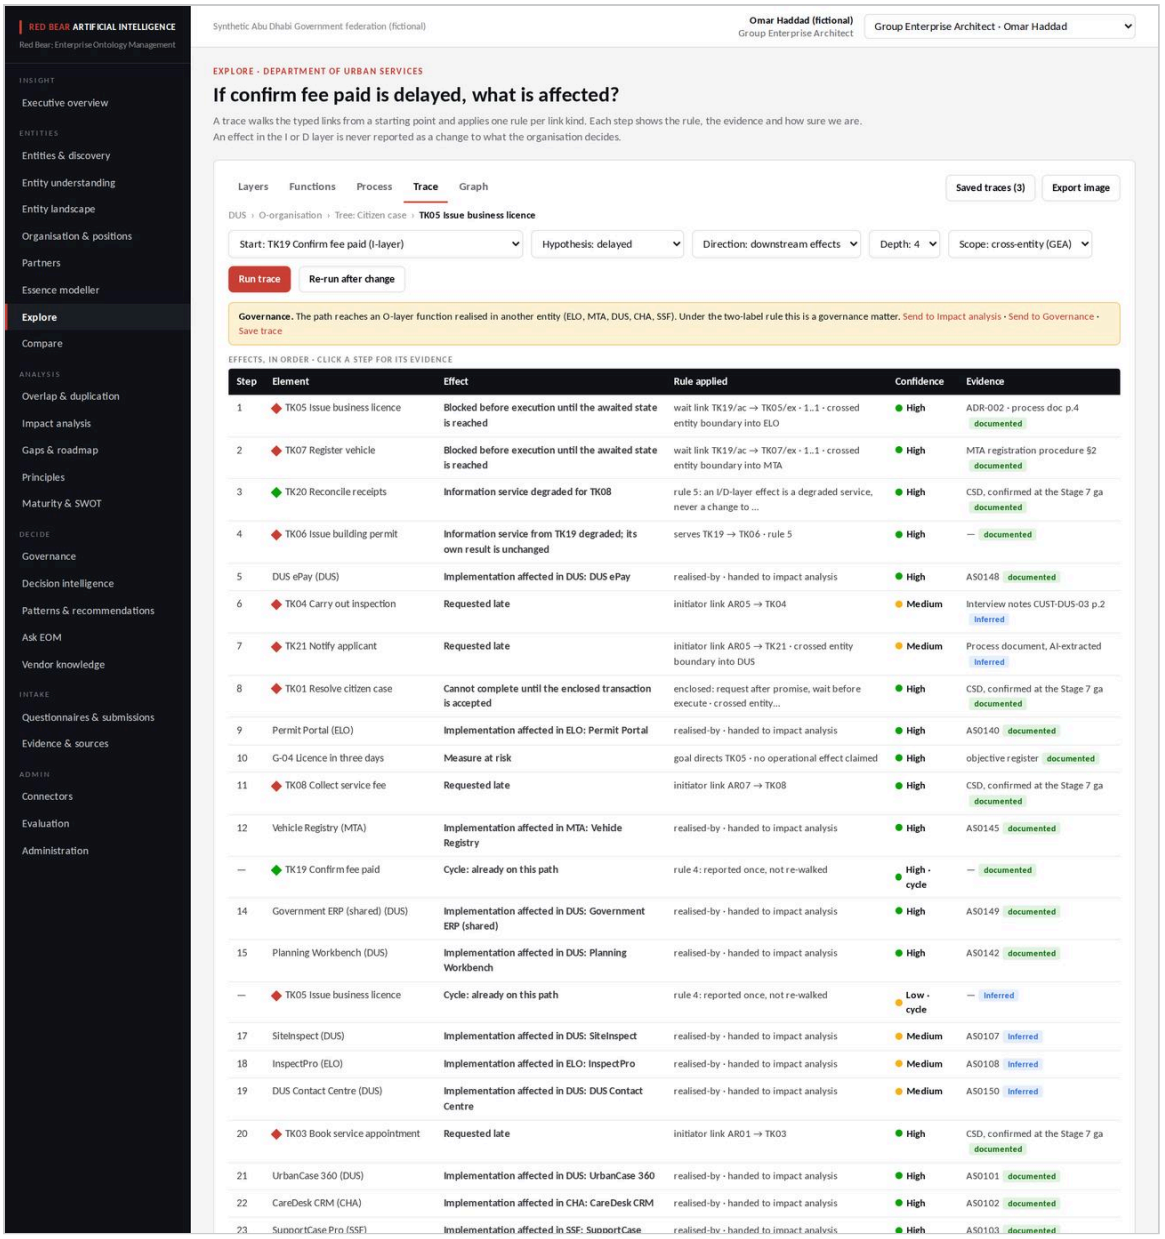
6 The user experience



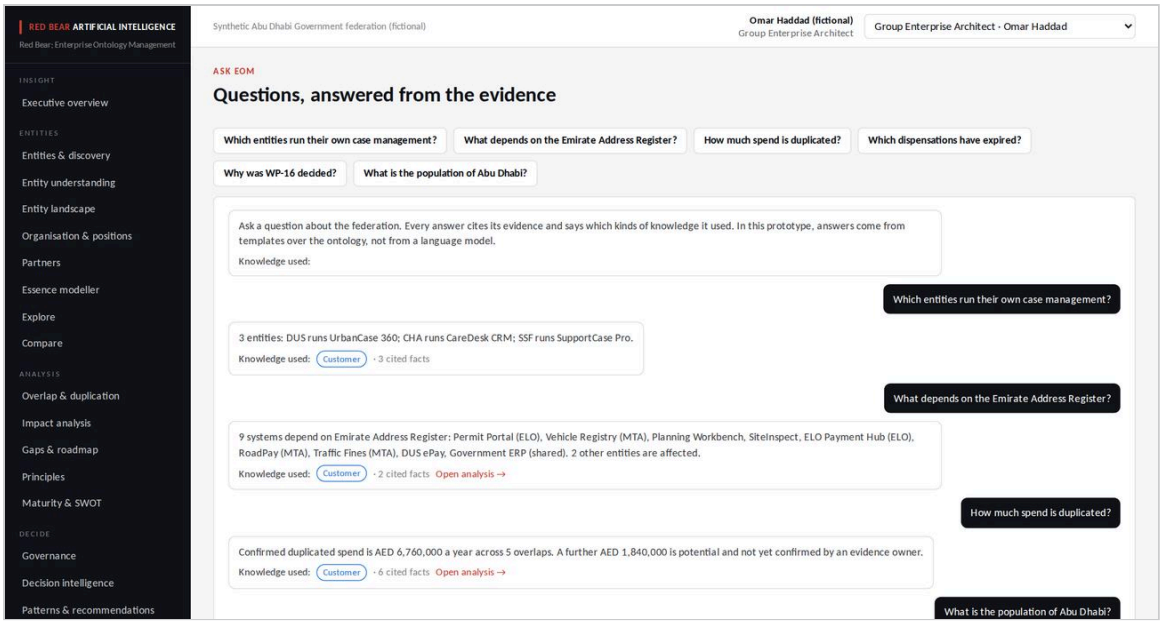
S-18 Decision intelligence: S-18 Decision intelligence. Every decision with its status, the five-dimension impact estimate with uncertainty per dimension, the scope-completeness items that keep it from being marked complete, candidate decisions found in documents awaiting review, and the lessons register. Prototype render on synthetic data for fictional entities.



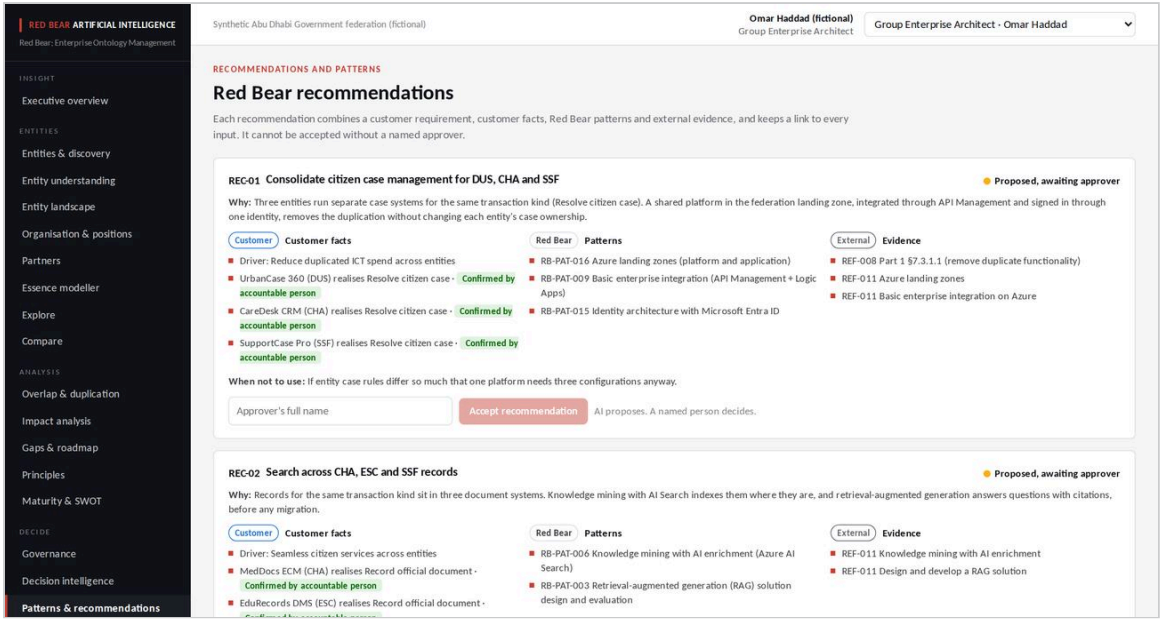
impact: S-04 Impact analysis. What a proposed change touches across the landscape: systems, interfaces, data, entities and principles, with the change class and the governance level it implies. Prototype render on synthetic data for fictional entities.



UC-41: S-19 Explore, Trace mode. The essence-level view behind the estimate: what a change blocks, delays or misinforms, step by step, with the rule and evidence for each, the confidence falling through inferred links, and the governance flag when two entities are reached. Prototype render on synthetic data for fictional entities.

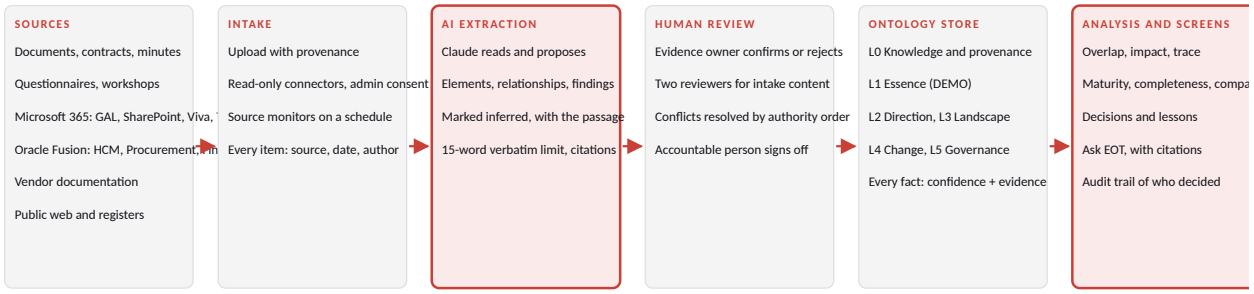


ask: S-11 Ask EOM. A natural-language question answered from the ontology with citations; the alignment scope returns a verdict per dimension. Questions outside the ontology are refused rather than guessed. Prototype render on synthetic data for fictional entities.



patterns: S-08 Patterns & recommendations. A recommendation composed from the customer requirement, the customer ontology, a Red Bear pattern and external evidence, with the link to each input; it has no effect until a Group EA accepts it. Prototype render on synthetic data for fictional entities.

7 Where the data comes from, and how it is checked



The pipeline for decisions: candidates are extracted from documents and messages; estimates and traces are computed from the confirmed ontology; every recommendation stops at a person.

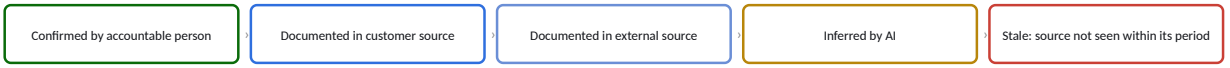
What each part of a decision rests on

Part	Comes from	Check
Candidate decision	Minutes, strategies, opted-in Teams channels, submitted documents	Inferred until an analyst reviews; the passage shown; fifteen-word quotation limit
Options and reasons	The submission; Claude drafts; the Entity EA edits	Every reason cites evidence or is marked as the author's judgement
Cost dimension	Contracts and spend from Fusion; pattern cost drivers from the catalogue; exit costs	Basis stated; uncertainty stated; figures cite their source
Risk dimension	Dependent systems and entities from impact analysis; partner findings from monitoring	Classes stated; each risk cites the element or finding
Governance dimension	The essential-model label from the trace; entities reached; principles touched	Level follows the delegated-authority table; principles cite conformance evidence
Compliance dimension	Standards affected in the landscape	States plainly when no regulatory source is loaded
Scope completeness	Elements the walk reached but could not assess	Listed; the estimate is not marked complete until each is resolved or acknowledged
Outcome and lessons	The review at the review date; delivery status	Outcome recorded against the estimate; lessons attached to vendor, pattern or capability with the decision as evidence

How it is checked

- **Nothing decides itself.** A candidate, a draft, an estimate and a recommendation are all proposals. The acceptance by a named person is the decision, and the record shows who and when.
- **Uncertainty is written down.** Each dimension carries a statement of what the estimate rests on and what it does not know; a single number is never shown without it.
- **Scope gaps block completion.** The scope-completeness dimension lists what the walk could not assess; the estimate stays *not complete* until each item is resolved.
- **Confidence follows the ontology.** An estimate built on inferred links inherits their lower confidence; the trace shows where.

- **Authority is configured, not assumed.** The delegated-authority table names the approver per class per tenant; changing the table is itself a governance decision.
- **Outcomes are compared to estimates.** The review records what happened against what was expected, which is how the estimate rules improve.
- **Refusal is an answer.** Ask EOM refuses questions outside the ontology and says so; it does not fill the gap with a general answer.



The confidence levels, inherited by every fact a decision rests on.

8 Benefits realised

Benefit	How it shows up	Measured by
Decisions approved first time	Alignment checked before submission; estimate and trace attached	Share of submissions approved without return
Impact known before approval	Five dimensions with uncertainty; scope gaps resolved	Share of decisions with a complete estimate
Right body, right decision	Two labels and the delegated-authority table	Governance-class decisions approved at the named level
Decisions that can be reconstructed	Options, reasons, evidence and approver stored	Records with all fields; time to answer "why was this decided?"
Lessons that travel	Lessons attached to vendors, patterns and capabilities; surfaced in new decisions	Lessons cited in later decisions
Drift caught	Approved vs delivered flagged	Decisions with drift flags resolved
Estimates that improve	Outcome reviews compared with estimates	Estimate error trend by dimension

9 Risks managed

Risk	Control in Red Bear EOM
AI recommendation treated as a decision	No effect until a named person accepts; the acceptance is the recorded decision; LLM outputs logged with user, sources and time
False precision in the estimate	Uncertainty statement per dimension; scope-completeness dimension; no single number without its basis
Essential change approved as management	Two-label rule; the trace marks governance when an O-layer element is reached in more than one entity or an actor role or transaction kind changes
Wrong approver	Delegated-authority table per tenant; changing it is a governance decision; audit trail
Decisions discovered from private channels	Teams channels by owner opt-in only; chats never read; candidates show their source

Risk	Control in Red Bear EOM
Lessons that blame	Lessons are attached to vendors, patterns and capabilities, not to people; outcome reviews record what happened, not who erred
Compliance claimed without a regulatory source	The compliance dimension states when no regulatory source is loaded; the gap is a known gap in the library
Estimates that never improve	Outcome reviews compare estimate and result; the estimate rules are reviewed against the error trend

Sources

Every definition and every claim attributed to a source in this paper names its source in the text, in the form *author, work, chapter or section*; the full references are below. Statements without a citation are Red Bear's own position, developed for Red Bear: Enterprise Ontology Management, and are marked RB-ORIGINAL in the product's method. Quotations are limited to fifteen words.

Cited as	Source	Authority
TOGAF Standard	The Open Group, *The TOGAF Standard*, 10th Edition, C220 Parts 0–5 (2022), used under Red Bear's licence	Industry Framework
Dietz and Mulder	J. Dietz and H. Mulder, *Enterprise Governance and Management*, Springer (2026)	Academic Reference
Greefhorst and Proper	D. Greefhorst and E. Proper, *Architecture Principles*, Springer (2011)	Academic Reference
Agile EA Playbook	The Open Group, *The Agile Enterprise Architect Playbook*, G226 (2022)	Industry Framework
Azure Architecture Center	Microsoft, *Azure Architecture Center* (19 patterns, read 25 September 2026)	Vendor Reference Architecture
Red Bear	RB-MM metamodel, RB-EAM method, EOM-FSD-01 v1.1 draft, function-explorer.md, eval-set.md	Red Bear

Glossary

Term	Meaning in this paper
Decision record	The stored decision: context, options, choice, reasons, evidence, approver, authority, status, outcome and lessons
Impact estimate	The pre-decision assessment on five dimensions: cost, risk, governance, compliance, scope completeness, each with an uncertainty statement
Dependency trace	The walk over the essence that shows what a change blocks, delays or misinforms, with a rule and evidence per step
Alignment question	A question of the form "does this initiative, design or change align with the group?", answered with a verdict per dimension and evidence
Candidate decision	A decision discovered in documents or messages, inferred until an analyst reviews it
Outcome review	The post-decision check of what happened against what was expected
Lesson	A reusable statement derived from an outcome review, attached to the vendor, pattern or capability it concerns
Two-label rule	Whether a change alters the essential model, and its TOGAF change class
Delegated-authority table	Per tenant: which board or role approves each decision class

Term	Meaning in this paper
Architecture decision record (ADR)	The TOGAF/industry form of a decision record; EOM's record extends it with impact and outcome

The series

This paper is one of five. Each stands alone and each follows the same nine sections, so a reader who knows one knows where to look in the next.

Paper	Title	Theme
EOM-WP-01	Enterprise Ontology for Federated Organisations	One essence, many entities
EOM-WP-02	Enterprise Ontology for Inter-Entity Automation	Why automation across the boundary fails without a shared ontology
EOM-WP-03	Enterprise Ontology Maturity and Baselining	How much do you actually know?
EOM-WP-04	Enterprise Architecture Roadmap Harmonisation	Roadmaps that agree
EOM-WP-05	Enterprise Architecture Decision Intelligence	Decisions with a memory (this paper)

Red Bear Artificial Intelligence · Abu Dhabi, United Arab Emirates · info@redbearai.net · www.redbearai.net