

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Enterprise Ontology Maturity and Baselining

How much do you actually know? Measuring completeness and maturity against the entities you choose to be judged by

Two numbers a federation should never be without: how much it knows about each of its entities, and how each compares with the peers it has chosen as baselines. This paper explains both, the phases an understanding passes through, and how Red Bear EOM computes, evidences and keeps them honest.

Empower. Trust. Verify.

White paper EOM-WP-03 · Red Bear white paper series · Version 0.1 DRAFT

26 September 2026

Summary

A model that does not say how complete it is will be trusted exactly as much as it should not be. Most architecture repositories present every entity with the same confidence, whether it was modelled from a two-day workshop or from a confirmed, source-backed understanding built over a quarter. This paper introduces two measures that Red Bear EOM computes rather than asserts: a completeness KPI, derived from a heatmap of what is known, missing, stale and uncertain for each entity, and a maturity assessment against Red Bear's maturity model, scored relative to Baseline entities the group has chosen to be judged by. It describes the five understanding phases an entity passes through, the SWOT and initiatives that follow an assessment, and the evaluation set that checks EOM's own answers against a general AI before any claim of usefulness is made.

WHO SHOULD READ THIS

- Heads of architecture and transformation offices
- Programme leads choosing where to invest in understanding
- Entity architects preparing for a maturity review
- Audit and assurance functions that need to know what the model is worth

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Distinguish completeness from maturity and know which question each answers
- Read an entity's heatmap and understanding phase
- Choose baselines and interpret a maturity comparison
- Explain how EOM proves its answers are worth more than a general AI's

1 What it is

Two different questions hide behind the word "maturity". The first is about the **model**: how much does the group actually know about this entity, and how current is it? The second is about the **entity**: how well does it perform on the dimensions that matter, compared with the peers it should be measured against? Red Bear EOM keeps the two apart, because conflating them is how a well-documented weak entity comes to look strong and a poorly-documented strong one comes to look weak.

Completeness: how much we know

For every entity EOM maintains a **heatmap** of knowledge areas (essence, capabilities, systems, data, vendors and contracts, organisation and positions, principles, decisions, and so on) against the entity, with each cell marked *known*, *missing*, *stale* or *uncertain* and carrying the action that would improve it. The **completeness KPI** is the weighted share of cells that are known and current. It rises when an analyst confirms something, falls when a source goes stale, and is shown beside the entity on every screen so that no reader mistakes a thin model for a full one.

Completeness is tied to an **understanding phase** that says how the knowledge was obtained, following the sequence a consultancy engagement would: *Awareness* (someone has told EOM the entity exists), *High-Level Assessment* (public sources and registers read), *Low-Level Assessment* (documents and information submitted by the entity reviewed and loaded), *Consulted* (a questionnaire completed or a workshop held),

and *Confirmed* (an accountable person at the entity has signed off the in-scope assertions). A phase is proposed by EOM when its criteria are met and moved by an analyst, never automatically.

Maturity: how well the entity does

Maturity is assessed against the **Red Bear Maturity Model**, six dimensions at five levels, each level with the evidence required to claim it. The dimensions draw on the sources: TOGAF's architecture capability and governance maturity (TOGAF Standard, EA Capability and Governance), the principle lifecycle and conformance scale of Greefhorst and Proper (Greefhorst and Proper, Ch 4–5), the agile architecture practices of the Open Group playbook (the Agile Enterprise Architect Playbook), and the organisational essence and governance separation of Dietz and Mulder (Dietz and Mulder, Ch 9 §9.6). A level is claimed only with evidence at the required confidence; the model itself, and the requirement that the confirmation labels are Red Bear's own, are marked RB-ORIGINAL.

Maturity is always **relative to baselines**. A Baseline entity is a peer the group chooses to be judged by: another emirate's licensing authority, a national digital government agency, a regulated utility. Baselines are known only from public sources and are never asked for anything, so their confidence is capped at *documented in external source* and their scores are shown with that caveat. Choosing baselines is a decision, recorded with its reasons; EOM suggests candidates and a Group EA confirms them.

Key points

- Completeness measures the model; maturity measures the entity; never blend them
- Understanding passes through five phases, moved by people when criteria are met
- Maturity is scored against chosen baselines, with the evidence required per level

2 Why it matters

A group that cannot say how much it knows will either over-trust its model or ignore it. Both are expensive. Over-trust produces consolidation decisions taken on a landscape that was last confirmed two reorganisations ago. Ignoring the model produces the parallel spreadsheet, the consultant re-survey, and the board pack assembled from memory.

A group that cannot say how mature each entity is, relative to peers, cannot direct investment. Every entity will claim to be the reference for the capabilities it cares about. Without a common model, evidence rules and chosen baselines, the argument is settled by seniority. With them, it is settled by the assessment, and the SWOT and initiatives that follow are the entity's own improvement plan rather than an imposition.

The final reason is credibility. EOM uses AI to read, extract, compare and draft. A tool that does that must be able to show that its answers are at least as useful as the general AI answer a manager could get in a browser, or it has no claim on the group's trust. Red Bear's position is that the claim is made only when the evaluation set says so, and the evaluation set, its rubric and the rule that judging is currently by AI only are published with the specification.

Without the measures	With them
Every entity presented with the same confidence	Completeness KPI and phase beside every entity; stale share visible

Without the measures	With them
Maturity claimed by assertion	Levels claimed only with evidence at the required confidence, against chosen baselines
Improvement plans written from opinion	SWOT and initiatives derived from the assessment, each with its evidence and the dimension it moves
AI trusted or distrusted wholesale	Evaluation set scored on six criteria; the tasks EOM lost are listed

3 The challenges

Challenge	Why it is hard	What goes wrong without a method
Discovering what is missing	Absence is invisible in a repository; nobody records what was never asked	The heatmap is all green because there are no red cells to draw
Capturing at a stated confidence	A workshop note, a signed questionnaire and a public annual report are not equally reliable; treating them alike is the norm	A partner's website becomes a confirmed fact
Maintaining currency	Sources age at different rates; a contract register from last month is current, a strategy from three years ago is not	Staleness is discovered when a decision goes wrong
Understanding the impact of change on the measures	A reorganisation, a new vendor or a leadership change can invalidate a whole knowledge area	The KPI stays high while the truth has moved
Driving standardisation of the assessment	Assessors differ; the same evidence scores differently on different days	Maturity becomes a negotiation
Knowing who decides	Who may move a phase, confirm a level or choose a baseline is rarely stated	The entity assesses itself; the group discounts it

4 How Red Bear EOM helps

Challenge	What EOM does	Module
Discovering what is missing	The heatmap is defined by the metamodel, not by what happens to be loaded: every knowledge area exists for every entity from the day it is registered, and starts *missing*	UN Understanding (FR-UN-02)
Capturing at a confidence	Five levels from *inferred by AI* to *confirmed by accountable person*; questionnaire and submission content needs two reviewers; Baseline knowledge capped at *external*	KE Knowledge, QI Intake
Maintaining currency	Each source has a period; facts go stale when their source is not seen within it; source monitors per entity type; significant events raise reassessment tasks	UN, PN
Impact of change on the measures	Reassessment tasks name the knowledge areas invalidated; the KPI and the phase can be lowered by an analyst, never only raised	UN (FR-UN-04)

Challenge	What EOM does	Module
Standardising the assessment	Evidence required per level, written into the maturity model; Claude proposes a level from the evidence, an assessor confirms; the reasons are stored	MA Maturity
Who decides	Phase moves, level confirmations and baseline choices are role-gated and audited; the accountable person at the entity signs the Confirmed phase	RA Roles

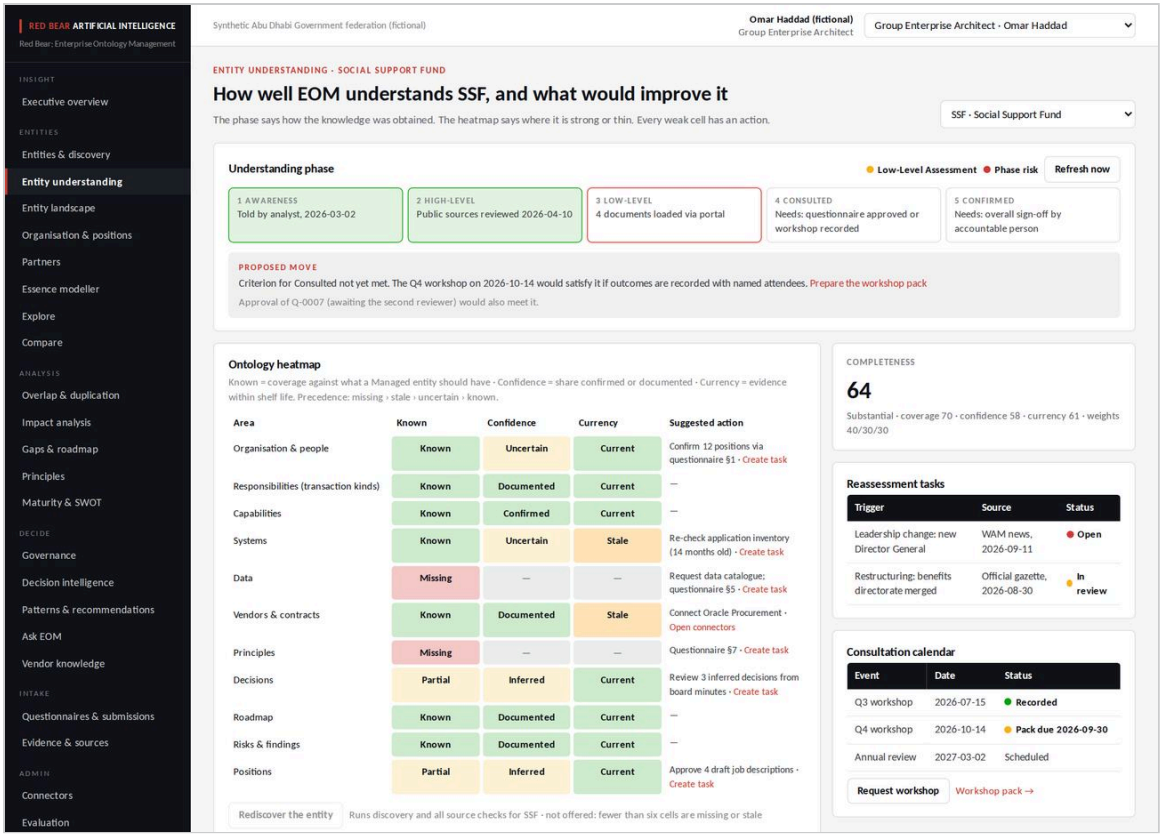
The quality bar

Before EOM is described as useful, its answers to a fixed set of tasks are scored against a general AI answer to the same question, with the same inputs, by an AI judge on six criteria. The tasks it lost are listed. The claim is made only when the aggregate condition holds, and the report states that judging was by AI only.

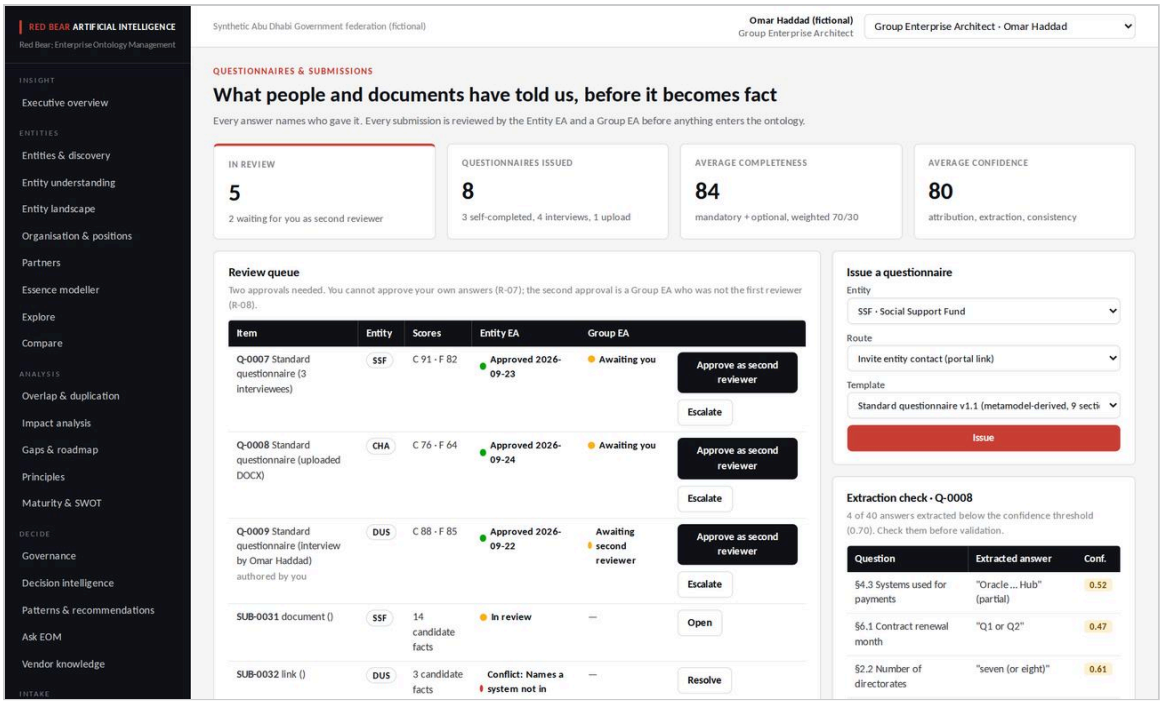
5 How it is used

Step	Who	What happens
1 Register and choose	GEA	Managed entities registered; EOM suggests Baseline candidates per capability from discovery; the Group EA confirms the baselines with reasons
2 Work the heatmap	EEA, CON	Each red or amber cell carries an action: send a questionnaire, request a document, run a connector, hold a workshop; the Entity EA works them in priority order
3 Move the phases	EEA, accountable person	EOM proposes a phase move when the criteria are met; the Entity EA moves it; the accountable person signs the Confirmed phase with the in-scope assertions listed
4 Assess maturity	GEA, EEA	Claude proposes a level per dimension from the evidence; the assessor confirms or overrides with a reason; the comparison against baselines is drawn
5 Derive the SWOT and initiatives	EEA, GEM	Strengths and weaknesses from the comparison, opportunities and threats from findings and partner monitoring; initiatives proposed with owner and cost basis; taken into the roadmap (EOM-WP-04)
6 Review quarterly, re-baseline annually	GEA, boards	Quarterly workshop packs generated from the heatmap and findings; the annual ontology review re-confirms scope, baselines and phases
7 Run the evaluation set	GEA	On each release, the evaluation tasks run; the report is published with the tasks EOM lost

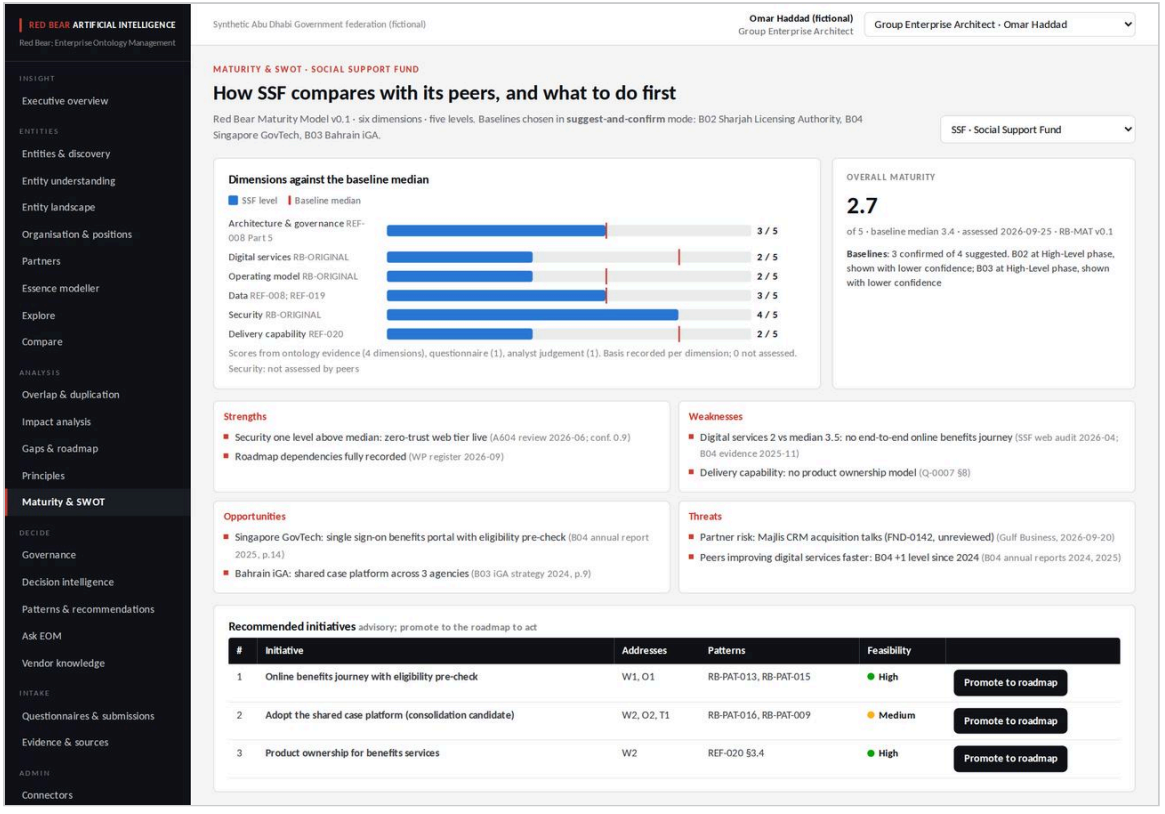
6 The user experience



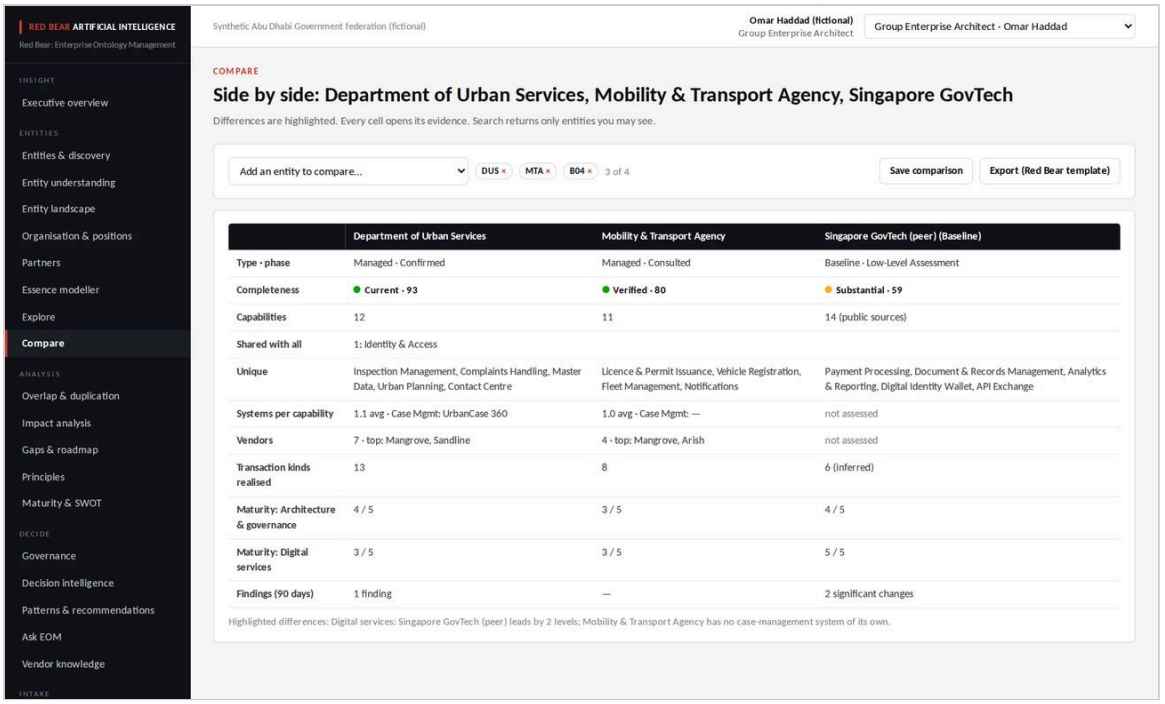
S-14 Entity understanding: S-14 Entity understanding. The five phases as a stepper, the completeness KPI, and the heatmap of knowledge areas: known, missing, stale or uncertain, each cell with the action that would improve it. The next phase is proposed when its criteria are met and moved by a person. Prototype render on synthetic data for fictional entities.



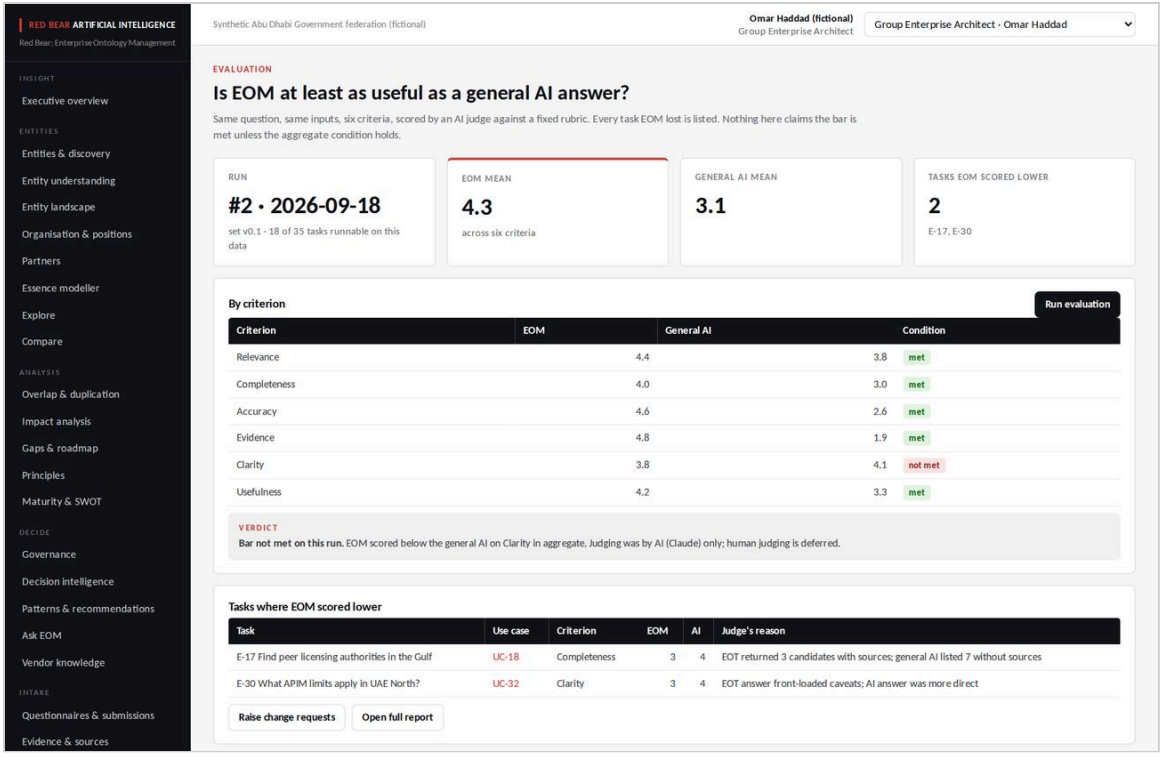
S-15 Questionnaires & submissions: S-15 Questionnaires & submissions. What people and documents have told EOM before it becomes fact: who answered, what was submitted, and the two reviewers each item needs before anything enters the ontology. Prototype render on synthetic data for fictional entities.



S-17 Maturity & SWOT: S-17 Maturity & SWOT. Six dimensions at five levels for one entity against its chosen baselines, the evidence behind each level, the SWOT derived from the comparison and the initiatives proposed from it. Prototype render on synthetic data for fictional entities.

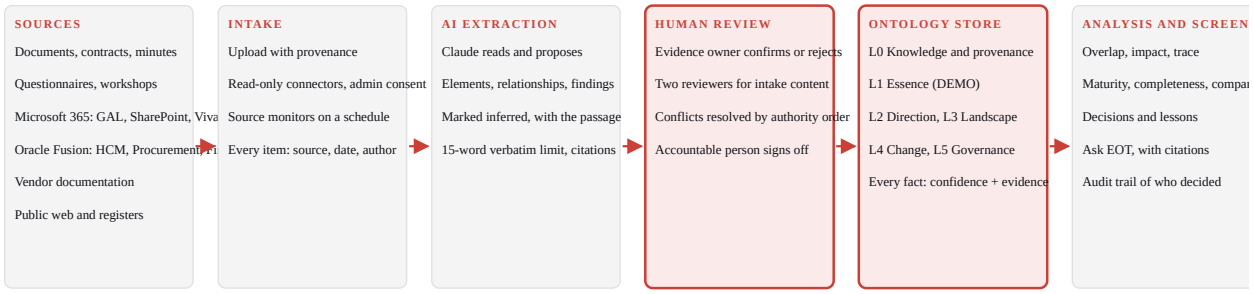


S-20 Compare: S-20 Compare. Up to four entities side by side: capability coverage, systems and vendors per capability, transaction kinds, maturity where assessed, completeness KPI and findings, with differences highlighted and every cell opening its evidence. Prototype render on synthetic data for fictional entities.



S-22 Evaluation: S-22 Evaluation. Is EOM at least as useful as a general AI answer? The same question, the same inputs, six criteria, an AI judge against a fixed rubric; every task EOM lost is listed. Prototype render on synthetic data for fictional entities.

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



The pipeline, with review and the store highlighted: the completeness KPI and maturity levels are computed from what is in the store at its recorded confidence, never from what was merely proposed.

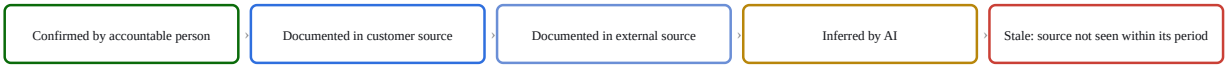
What feeds each measure

Measure	Inputs	Rule
Heatmap cell	Assertions in the knowledge area with their confidence and source date	Known if at least one confirmed or documented assertion is current; stale if the source period has passed; uncertain if only inferred; missing otherwise
Completeness KPI	Heatmap cells, weighted by knowledge area	Weighted share of known cells; shown with the stale share beside it
Understanding phase	Phase criteria: sources read, submissions reviewed, questionnaire or workshop, accountable sign-off	Proposed by EOM when criteria are met; moved by an analyst; lowered on a significant event
Maturity level	Evidence required per level per dimension	Proposed by Claude from the evidence; confirmed by an assessor with reasons; Baseline levels capped at external confidence
SWOT item	Maturity comparison, findings, partner monitoring	Each item cites the comparison cell or finding it came from
Evaluation score	The evaluation set, the data pack, the rubric	Both answers judged by Claude on six criteria 1–5; aggregate condition; tasks lost listed; AI-only judging stated

How it is checked

- **No cell is green by default.** Knowledge areas come from the metamodel and start missing; only evidence turns them.
- **Two reviewers for intake.** Questionnaire answers and submitted documents pass the Entity EA and a Group EA before anything in them counts toward completeness.
- **Confirmed means signed.** The Confirmed phase requires a named accountable person at the entity to sign off the listed in-scope assertions; the signature, date and list are stored.
- **Levels need their evidence.** A maturity level cannot be confirmed without the evidence the model requires for it; an override is allowed only with a reason, and the reason is shown.
- **Baselines are caveated.** Every Baseline score carries *documented in external source* and the date the source was read.

- **The KPI can go down.** Staleness and significant events lower it; a rising-only measure would be a vanity metric.
- **EOM grades itself in public.** The evaluation report lists the tasks EOM lost; the rule that judging is by AI only is stated on the report until human judging is added.



The confidence levels that every measure inherits from the facts beneath it.

8 Benefits realised

Benefit	How it shows up	Measured by
Investment in understanding directed where it counts	Heatmap actions worked in priority order; entities at Awareness identified	Completeness KPI trend per entity; cells closed per quarter
A model whose worth is known	Every screen shows the KPI, phase and stale share	Stale share; share of assertions at confirmed
Maturity that is defensible	Levels with their evidence, against chosen baselines	Dimensions assessed; overrides with reasons
Improvement plans the entity owns	SWOT and initiatives derived from the assessment, taken into the roadmap	Initiatives adopted; dimension movement year on year
Group-level comparison without argument	Up to four entities compared on the same evidence rules	Comparisons saved and cited in decisions
Trust in the tool that is earned	Evaluation report per release	Aggregate condition met; tasks lost, trending

9 Risks managed

Risk	Control in Red Bear EOM
Completeness mistaken for maturity	Two measures, two screens, two vocabularies; the compare view shows both side by side and never blends them
Vanity KPI	The KPI falls on staleness and significant events; the stale share is shown beside it
Assessor bias	Evidence required per level; Claude proposes from the evidence; overrides need a reason and are audited
Baseline knowledge over-trusted	Capped at external confidence; source and read date on every score; peers never contacted
Entity self-assessment discounted by the group	The Confirmed phase and maturity confirmations are role-gated to named people; the group sees who signed and when
AI judging its own work	The rule is stated on every evaluation report; the rubric and tasks are published; human judging is the planned next step
Privacy of questionnaire respondents	Answers name the respondent to reviewers only; published facts cite the questionnaire, not the person, unless they are the accountable signatory

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
Completeness KPI	How much EOM knows about an entity, computed from its heatmap: the weighted share of cells that are known and current
Heatmap	The grid of knowledge areas by entity showing known, missing, stale or uncertain, each cell with an action
Understanding phase	Awareness, High-Level Assessment, Low-Level Assessment, Consulted, Confirmed: how the knowledge was obtained
Maturity model	Red Bear Maturity Model v0.1: six dimensions, five levels, with evidence required per level
Baseline entity	A peer the group chooses to compare against; known only from public sources; never asked for anything
SWOT	Strengths, weaknesses, opportunities and threats derived from the maturity comparison, each with its evidence
Initiative	A proposed action from the SWOT, with owner, cost basis and the dimension it moves
Evaluation set	A fixed list of tasks on which EOM's answers and a general AI's are scored by an AI judge against a rubric
Stale	A fact whose source has not been seen within the period set for that source

Term	Meaning in this paper
Reassessment task	Work raised when a source changes or a fact goes stale

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? (this paper)
EOM-WP-04	Enterprise Architecture Roadmap Harmonisation	Roadmaps that agree
EOM-WP-05	Enterprise Architecture Decision Intelligence	Decisions with a memory

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