RED BEAR: ENTERPRISE ONTOLOGY MANAGEMENT · DATA SHEET · 3 · ANALYSIS

Where entities overlap, what a change will break, and how mature each one is

The questions an executive cannot answer today, answered from the facts, with the evidence one click away and every figure reproducible.

KEY POINTS

- Overlap is detected by function: every transaction kind or capability realised by two or more entities, with each entity's systems, vendors and contracts behind it
- Duplicated investment is quantified per overlap and in total, and only human-confirmed matches count; the rest is shown apart as *potential*
- Impact analysis walks the landscape; the dependency trace walks the essence, with one effect rule and one piece of evidence per step, and flags when the change is a governance matter
- Maturity is scored against chosen baselines with the basis of every score recorded; a dimension without evidence is *not assessed*
- Every figure is deterministic: rebuilding from the same data reproduces it exactly

What it is

Analysis is the set of computed views over the shared essence and the per-entity landscapes that Red Bear: Enterprise Ontology Management holds. None of them is stored as a model of its own; each is calculated from the facts, so it is always as current as the evidence beneath it and can be reproduced on demand.

Four families of question are covered. *Where do we overlap*, and what does the duplication cost. *What will this change affect*, directly or indirectly, and by what path. *Are we following our own principles*: the register, its quality flags and conformance. *How mature is each entity, and how well do we know it*: a maturity model against chosen baselines, a SWOT with evidence, and the completeness score.

The pain point

Overlap between entities is invisible in the repositories that are supposed to reveal it, because each entity names the same job differently. Three entities each running their own case-management system for the same licence look like three different capabilities, and the three contracts behind them are never added up. Consolidation is announced and never measured.

Impact is guessed from memory. A change in one entity travels along dependencies to others: a delayed transaction blocks the one waiting for it, a changed data bank misinforms everyone who reads it, a shared platform change lands in every entity that uses it. The entities not in the room are the ones that break. And maturity is asserted rather than assessed, on a scale nobody can trace to an observable criterion, against peers nobody has chosen.

How it works

Overlap by function, spend by confirmation. Every capability or transaction kind realised by two or more entities is listed with each entity's realising elements, and every overlap opens to the assertions and evidence behind it. Where an overlap is served by distinct components or contracts the duplicated annual contract value is shown per overlap and in total, but the total includes only matches a person has confirmed; unconfirmed matches are shown separately as potential. Decoys that share a name but differ in essence are not overlaps. Rebuilding from the same data reproduces every figure exactly.

Two ways to ask what breaks. Impact analysis takes a proposed change to a landscape element and lists every dependent capability, process, data entity, interface, entity and work package, separating direct from indirect, showing the path to each and flagging impacts that cross entities. The dependency trace works on the essence: given a start element, a change hypothesis, a direction, a depth and a scope, it walks the typed links breadth-first, applying one effect rule per edge kind, crossing trees only along wait and access links, lowering confidence through inferred edges and flagging stale ones. It returns the ordered effect list with the rule and evidence per step, the highlighted path and a plain-language explanation, and marks the trace governance when it reaches the essence of more than one entity. Both are explored on one screen with five modes over one store, with inferred links drawn dashed and stale elements outlined amber.

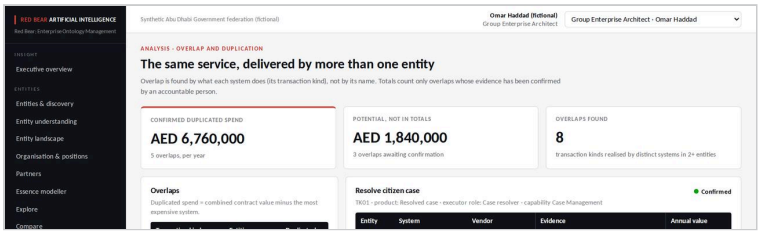
Maturity and completeness, with the basis stated. The Red Bear Maturity Model covers architecture and governance, digital services, operating model, data, security and delivery, five levels each with observable criteria and a source. Baselines are chosen in one of three ways, the product suggests and the analyst confirms, the analyst selects freely, or all same-sector Baselines, and the mode is recorded with the assessment. Each dimension is scored from ontology evidence, questionnaire answers or analyst judgement, with the basis recorded; the SWOT carries evidence, source dates and confidence, with lessons from the baselines that demonstrate each opportunity; initiatives are ranked, linked to patterns and promotable to candidate work packages that a Group EA accepts or declines. The completeness score, 0 to 100

from coverage, confidence mix and currency, capped by phase, is shown as one of six levels with its trend. The principle register flags any principle without a rationale and shows conflicting principles side by side without ranking them.

What you get

Capability	What it does	Requirements
Overlap and duplication	Overlaps by transaction kind or capability with the evidence behind each; duplicated annual contract value per overlap and in total; confirmed and potential apart; deterministic	FR-AN-01 to FR-AN-05
Impact analysis	Direct and indirect dependents of a change, the path to each, cross-entity impacts flagged; the gap matrix from baseline to target with unassigned gaps counted	FR-AN-06, FR-AN-07
Principle register	Statement, rationale, implications and conformance per principle; quality flags; conflicting pairs shown without ranking	FR-AN-08, FR-AN-09, FR-AN-10
Explore and trace	Layers, functions, process, trace and graph modes over one store in DEMO notation; the function card and derived tables; confidence and currency marked; the dependency trace with an effect rule and evidence per step, stored and re-runnable	FR-VZ-01, FR-VZ-02, FR-VZ-05 to FR-VZ-11
Compare	Up to four entities side by side on coverage, systems, vendors, transaction kinds, maturity, completeness and findings, scoped by role	FR-VZ-03, FR-VZ-04
Maturity, SWOT and completeness	The Red Bear Maturity Model against chosen baselines with the basis of every score; SWOT with evidence; ranked initiatives; the completeness score with configurable, logged weights	FR-MA-01 to FR-MA-07

Screens



S-03 Overlap & duplication: Overlap and duplication: entities realising the same transaction kind, the systems and vendors behind each, and the duplicated spend, confirmed and potential shown apart. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Three measures carry this pillar, each reported by the product and none carrying a target in this document: **the overlap count and the confirmed duplicated annual contract value**, per overlap, in total and by quarter (FR-AN-03, FR-AN-04); **the share of governance decisions with an impact estimate or trace attached** before approval (FR-AN-06, FR-VZ-09); and **maturity and completeness per entity**, each with the basis of the score on the record (FR-MA-03, FR-MA-06). What the pillar makes possible is a consolidation case built on contracts rather than on a hunch; a change reviewed with the entities it will reach named before approval rather than discovered after; and an executive who can read all of it without an ontology vocabulary, because every analysis is one click from the object and the object one click from the evidence (FR-LO-03, FR-LO-04).

Platform engines

Analysis consumes the Red Bear Platform's Ontology Mapping engine (coordination links, overlap by function and the dependency trace, RBP FR-OM-04 to RBP FR-OM-06) and the Provenance and evidence service for confidence and staleness (RBP FR-PV-01, RBP FR-PV-03). Platform behaviour is specified in the platform functional specification, not here.