RED BEAR: ENTERPRISE ONTOLOGY MANAGEMENT · SUITE DATA SHEET · FORMERLY RED BEAR: ENTERPRISE ONTOLOGY TOOLKIT

One essence, many entities

A continuously maintained architectural intelligence system for governments, groups and alliances: what each entity is responsible for, how it is implemented, where they overlap, and what a change will affect. **Every fact with its evidence.**

KEY POINTS

- One shared essence for the federation, one implementation per entity, and cross-entity views computed from the facts rather than stored as a merged model
- Overlap and duplicated spend are found by function, not by name: two entities that realise the same transaction kind are doing the same job, whatever they call it
- Every fact carries its source, its exact location and a confidence level; AI extraction starts at *inferred* and only a named person can confirm
- Intake from questionnaires and documents needs two reviewers; the System Administrator runs the tenant and never sees its content
- The platform proposes and never decides: a named person confirms every fact, accepts every recommendation and approves every change

What it is

Few organisations of consequence are single enterprises. A government is a federation of legally distinct departments, authorities and funds; a corporate group is a set of subsidiaries and divisions; both share partners, vendors and, increasingly, platforms. Each entity documents itself in its own vocabulary, to its own depth, at its own time, and the questions that matter at group level go unanswered.

Red Bear: Enterprise Ontology Management gives a federated group one place to discover, capture, check and maintain the ontology and architecture of every entity it manages, partners with or measures itself against. An enterprise ontology is the understanding of an enterprise abstracted from its implementation (Dietz and Mulder, glossary): what it is responsible for and who is accountable for each result. Held beside each entity's TOGAF landscape, it is the unit of comparison that makes entities comparable by what they do.

The pain point

Which of our entities issue licences, and with how many different systems? Which vendor is in all of them? If this shared platform changes, who is affected? Which entity is the reference for a capability, and which should converge on it? An architecture repository cannot answer these, because it describes each entity in terms of its own systems, job titles and process diagrams, and two entities doing the same job look different.

The cost is paid three times: in **duplicated investment**, the same outcome funded in two or three places because nobody could see it was the same outcome; in **fragile change**, a modification in one entity travelling along dependencies to others that were not in the room; and in **unaccountable decisions**, whose reasons were never recorded against the evidence. The trigger is a condition, not an industry: **more entities than any one architect can hold in their head.**

How it works

Five pillars, each useful on its own and considerably more useful together, on the Red Bear Platform's ontology, gathering and research engines.

- **Knowledge and evidence** (EOM-DS-01). Every source registered with its kind, authority and licence; every fact an assertion with evidence and a confidence level; conflicts raised, never dropped; the tenant run by an administrator who sees no content; usefulness measured on a fixed evaluation set before it is claimed.
- **Understanding the enterprise** (EOM-DS-02). Managed, Partner and Baseline entities discovered and reviewed; five understanding phases per entity; the TOGAF landscape and the DEMO essence; questionnaires, workshops, Microsoft 365 and Oracle Fusion connectors, all with provenance and two reviewers.
- **Analysis** (EOM-DS-03). Overlap by function with the duplicated spend behind it; impact analysis and the dependency trace with a rule and evidence per step; the principle register; maturity against chosen baselines; the completeness score for every entity.
- **Decisions and recommendations** (EOM-DS-04). Patterns proposed with their reasons and trade-offs; questions answered with citations or refused; decisions classified twice, routed by impact, recorded, estimated on five dimensions and tracked to outcome.

- **Roadmap and governance** (EOM-DS-05). Transition timelines that respect dependencies; consolidation candidates across entities; partners onboarded, compared and monitored; four roles, approvers per decision class, and an audit of every read and refusal.

One essence, many implementations. Every result an organisation produces is created by a transaction: one actor role requests it, one executes it (Dietz and Mulder, Ch 4 §4.4.4). A transaction kind defined once is usable by any number of entities without being copied (FR-LO-07), and each entity links it to the capabilities, processes and applications that realise it (FR-LO-06). Every element belongs to exactly one entity, and every cross-entity view is computed from the layers, never stored (FR-LO-01). No entity's truth is overwritten by another's. The essence is built progressively, where overlap analysis needs it, on top of a landscape that is useful on its own.

Provenance, then authority, then a person. Every fact is an assertion with at least one piece of evidence naming its source and exact location: file, sheet and row; section and page; or URL (FR-KE-02). Facts extracted by AI start at *inferred by AI* and rise to *confirmed by accountable person* only when a named Evidence Owner confirms them (FR-KE-03). When two sources disagree a confirmed authority order decides, and the rule used is cited; when they carry equal authority an architecture decision record is opened for the owning board instead of a silent choice being made (FR-KE-08). No more than fifteen consecutive words of any external source are ever stored (FR-KE-06).

How you adopt it

A first quarter of establishment, then a steady state, with the roles most groups already have: a Group Enterprise Architect, an Entity Enterprise Architect per Managed entity, a Group Enterprise Manager, consultants, boards, and a System Administrator.

- **Connect** (weeks 1 to 2). Tenant set up, roles assigned, Microsoft 365 and Oracle Fusion connectors configured under admin consent, read-only; the group's own standards loaded as the highest authority. *You have* the organisation, people, suppliers and spend of each entity, with provenance.
- **Discover** (weeks 2 to 4). Managed entities registered; partners and baselines found by search and reviewed; every entity starts at Awareness. *You have* an honest statement of how much is known about each.
- **Understand** (month 2). Questionnaires, submissions and workshops; AI extractions reviewed by two people; the landscape fills. *You have* entities at High-Level and Low-Level Assessment, with a completeness score each.
- **Compare** (months 2 to 3). The shared essence modelled capability by capability; realisations recorded per entity. *You have* the first overlap and duplicated-spend reports, by function.
- **Govern** (month 3 onward). Cross-entity review; reference entities proposed; changes classified and routed under the two-label rule; source monitors, quarterly workshops and the annual review keep it current. *You have* decisions recorded with their impact and, later, their outcome.

RETURN ON INVESTMENT

The benefits are the ones the product itself measures, so each is reported from the product rather than asserted. **None carries a target in this document**; targets come from the first engagement's baseline.

- **Overlap made visible.** Entities realising the same transaction kind, with their systems and vendors; the overlap count and the duplicated annual contract value, confirmed and potential shown apart (FR-AN-03, FR-AN-04).
- **Duplicated investment stopped.** Consolidation candidates with the contracts behind them; the duplicated-spend trend by quarter (FR-RM-04).
- **Change made safe.** Impact and dependency traces run before a change is approved; the share of governance decisions with an impact estimate attached (FR-AN-06, FR-DI-04).
- **Knowledge that stays current.** The completeness score per entity, 0 to 100 in six levels, and its trend; the stale share; reassessment tasks closed within their period (FR-MA-06).
- **Decisions that can be defended.** Decision records with approver, authority, estimate and outcome review; estimate against observed (FR-DI-05).
- **Answers that beat a general AI.** The evaluation set scores the product against a general AI on six criteria; the bar is met only when the mean is at least equal on every criterion (FR-TQ-05).

Platform engines

Built on the Red Bear Platform, whose FSD specifies the Ontology Mapping, Organisation Information Gathering, Research, Decision Intelligence, Decision Learning and Roadmap Harmonisation engines and the shared identity, connector, provenance, AI gateway, privacy, audit and hosting services this product consumes.