RED BEAR: ENTERPRISE ONTOLOGY MANAGEMENT · DATA SHEET · 2 · UNDERSTANDING THE ENTERPRISE

# How much we know about each entity, and how it stays true

Discover the entities, gather what is known about each from people, documents and connected systems, model it once in a shared vocabulary, and keep it current.

KEY POINTS

- Three entity types, Managed, Partner and Baseline, found by search and reviewed by a person; nothing from discovery is treated as verified
- One honest phase per entity, Awareness to Confirmed, proposed by the product when a criterion is met and confirmed by an analyst
- One TOGAF landscape per entity and one DEMO essence shared across all of them, linked by what realises what
- Questionnaires, submissions, workshops, Microsoft 365 and Oracle Fusion each carry provenance; two reviewers before intake content becomes fact
- Sources are re-read on a schedule, significant events reopen an assessment, and stale facts are marked on every screen

## What it is

Understanding the enterprise is the set of functions by which an organisation becomes known to Red Bear: Enterprise Ontology Management, and stays known. It covers the entity from the day it is a name in a search result to the day an accountable person signs off its model, and every quarter after that.

Each entity's implementation is modelled on the TOGAF 10 content metamodel, extended with vendors and contracts, with baseline, transition and target states (TOGAF Standard, Architecture Content §2.4 to §2.6). What the entity is responsible for is modelled on the DEMO essence layer: transaction kinds, each with one product kind and one executor role, shared across the federation and realised differently in each entity (Dietz and Mulder, Ch 4 §4.4.4). The two are linked, and the link is what makes entities comparable by function.

## The pain point

Knowledge of an entity is spread across contracts, org charts, ERP configuration, SharePoint sites, people's heads and the public web, and partners and peers publish only fragments. Models start from whoever answered the survey; entities that never replied are simply absent. A model that is right today is stale in a year unless something watches the sources, and most repositories become archaeology sites on which decisions are taken anyway.

The honest question a group cannot answer today is not *what is this entity's architecture* but *how much do we actually know about it, and when was that last true*. Without a stated phase and currency, a two-year-old questionnaire and last month's ERP extract carry the same weight.

## How it works

**Discover and type.** Every entity is Managed, Partner or Baseline, with the type changeable by a Group Enterprise Architect and the history kept. Candidates are found from a description, an exact name or keywords through the configured search providers, and each arrives as a card: identity with a source and date per field, existing mentions in the tenant's ontology, a suggested type with reasoning, a duplicate check and an overall confidence. Facts from discovery enter no higher than *documented in external source*, an accepted candidate starts at Awareness, and a rejected one is not proposed again unless a newer source appears.

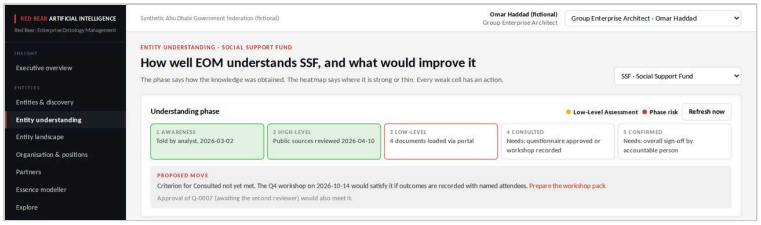
**Gather with provenance and two reviewers.** The standard questionnaire is derived from the metamodel, one section per ontology area, and is completed by a consultant in an interview, by an entity contact through an invitation link, or extracted from an uploaded PDF or DOCX with a confidence per answer. Submissions are validated, scored 0 to 100 on completeness and confidence, and reviewed by the Entity EA and a Group EA before anything becomes an assertion; nobody approves their own answers. Read-only connectors under admin consent read the directory, named SharePoint sites, Viva topics and opted-in Teams channels from Microsoft 365, and people, positions, suppliers, contracts, spend and projects from Oracle Fusion, each fact citing the object and read date. Job descriptions are drafted with a citation on every sentence and held until approved.

**Keep it true.** Each Managed entity carries one of five understanding phases, Awareness, High-Level Assessment, Low-Level Assessment, Consulted and Confirmed; the product proposes a move when the entry criterion is met and an analyst confirms it. A heatmap shows every ontology area as Known, Missing, Stale or Uncertain, with a suggested action per weak cell that becomes a task in one click. Sources are re-read on a schedule, and new facts enter as evidence for review, never overwriting confirmed ones. A leadership change, restructuring or major initiative opens a reassessment task and drops a Confirmed entity back to Consulted. Quarterly workshops and an annual review are scheduled per entity, each with a pack, and outcomes recorded with a named attendee become evidence.

## What you get

| Capability | What it does | Requirements |
|---|---|---|
| **Entity registry and discovery** | Managed, Partner and Baseline with type history and role-scoped visibility; candidates from description, name or keywords, each a card with source, mentions, suggested type, duplicate check and confidence; rejections remembered | FR-LO-01, FR-ET-01 to FR-ET-07 |
| **Landscape and essence** | The TOGAF content metamodel with vendors, contracts and valid-time states; transaction kinds with one executor, shared across entities and linked to what realises them; insight to evidence one click per level | FR-LO-02 to FR-LO-08 |
| **Phases, heatmap and refresh** | Five phases proposed by the product and confirmed by an analyst; the phase risk rule; a heatmap per entity with an action per weak cell; scheduled re-reads and significant events; manual refresh | FR-UN-01 to FR-UN-07 |
| **Workshops and reviews** | Quarterly workshops and an annual review per entity, packs in the Red Bear template, outcomes as evidence, requests from an Entity EA | FR-UN-08 to FR-UN-11 |
| **Questionnaires and submissions** | One metamodel-derived questionnaire, three completion routes, validation, scoring and two reviewers; documents, links and notes submitted with provenance | FR-QI-01 to FR-QI-07 |
| **Connected sources** | Microsoft 365 and Oracle Fusion, read-only under admin consent, scoped per entity; the shared Gov ERP mapped to entities; positions and cited job descriptions | FR-CN-01 to FR-CN-12 |

## Screens




**S-14 Entity understanding**: Entity understanding: the phase panel with the product's proposed move awaiting an analyst, the heatmap of ontology areas by coverage, confidence and currency, and the completeness score. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Two measures carry this pillar, both reported per entity by the product and neither carrying a target in this document: **the understanding phase**, how many Managed entities have reached Consulted or Confirmed and how long each took, with a risk raised for any Managed entity still below Consulted twelve months after creation (FR-UN-02); and **the completeness score**, 0 to 100 from coverage, confidence mix and currency, capped by phase and shown as one of six levels with its trend (FR-MA-06). What the pillar makes possible is a group that knows which of its entities it understands and which it only knows of; a partner or peer described from public sources at a stated confidence rather than from hearsay; intake that cannot enter the ontology on one person's word; and a model whose currency is visible on the screen rather than assumed in the meeting.

## Platform engines

Understanding the enterprise consumes the Red Bear Platform's Ontology Mapping engine (RBP FR-OM-01 to RBP FR-OM-03, RBP FR-OM-07), the Organisation Information Gathering engine (RBP FR-OG-01 to RBP FR-OG-06) and the Connectors service (RBP FR-CN-01 to RBP FR-CN-03, RBP FR-CN-05). Platform behaviour is specified in the platform functional specification, not here.