RED BEAR: ENTERPRISE ONTOLOGY MANAGEMENT · DATA SHEET · 1 · KNOWLEDGE AND EVIDENCE

Where every fact comes from, and how sure we are

Nothing enters the ontology without a source, and nothing becomes a fact without a person. This pillar is the floor every other one stands on.

KEY POINTS

- Every source is registered with its knowledge kind, authority class, licence and version; every fact is an assertion with evidence naming its exact location
- Confidence is graded, not assumed: AI extraction starts at *inferred* and only a named accountable person can raise it to *confirmed*
- Conflicts between sources are never dropped silently: a fixed authority order decides and cites its rule, or an architecture decision record is opened
- Red Bear, customer and external knowledge are kept in separate stores; one customer never sees another's elements, evidence or file names
- The System Administrator runs the tenant, its backups and its subscription without seeing its content; usefulness is measured on a fixed evaluation set before it is claimed

What it is

An architecture repository holds statements: this entity runs this system, this vendor holds this contract, this capability is realised here. It rarely holds where each statement came from, when it was true, or who would stand behind it. When two statements disagree, one of them quietly wins. When an AI drafts the model, its inferences look exactly like facts a person checked.

Knowledge and evidence is the layer beneath every element in Red Bear: Enterprise Ontology Management. Every fact is an assertion carrying at least one piece of evidence and a confidence level, and every piece of evidence names a registered source and the exact place in it. Whatever an executive sees on a screen can be opened down to the row, page or passage it came from, labelled Red Bear, Customer or External.

The pain point

Dietz and Mulder describe what happens when system-derived outputs are mistaken for original facts: the British Post Office prosecuted hundreds of its own subpostmasters on them (Dietz and Mulder, Ch 10). The failure is one of provenance before it is one of technology. A federation's architecture model faces the same risk at lower stakes every day: derived or inferred statements are treated as facts, a stale survey answer outranks last month's contract register, and nobody can say how sure anyone is.

AI makes this sharper, not easier. Claude can read a contract register or an interview transcript and propose elements and relationships quickly; it cannot vouch for them. People can vouch, but not at scale. The layer exists to hold the two apart: extraction that is fast and labelled, confirmation that is slow and named, and a record of which is which on every screen.

How it works

Register, then assert, then grade. Every source enters as a Knowledge Source with its knowledge kind (Red Bear, Customer or External), authority class, licence, version and ingestion date. Every ingested fact is stored as an assertion with at least one piece of evidence naming the source and its location: file, sheet and row; section and page; or URL. Facts extracted by AI start at *inferred by AI*; they move to *confirmed by accountable person* only when an Evidence Owner confirms them, and the record shows who confirmed and when. A source flagged licence-restricted cannot be used in derivation until a permission record exists, and no more than fifteen consecutive words of any external source are stored.

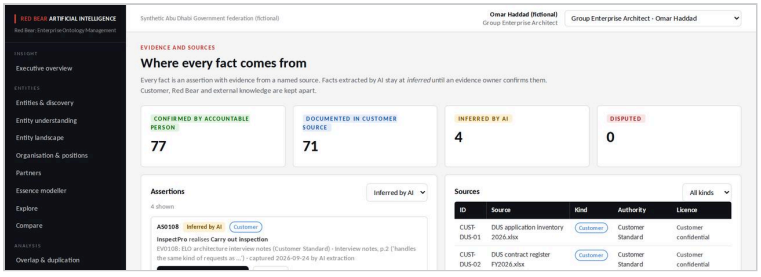
Resolve by authority, or escalate. Whenever two assertions cannot both hold, a Conflict is raised and neither is dropped. Between sources of different authority the confirmed order decides, regulatory first, then the customer's own standard, then Red Bear's, then vendor material and the rest, and the rule used is cited on the resolution. Between sources of equal authority the product does not choose: it opens an architecture decision record and assigns it to the owning board. Three stores, Red Bear, Customer and External, are kept apart, so every answer can say which kinds of knowledge it used and no customer's elements, evidence or file names are ever visible to another.

Run the tenant without reading it. Tenant setup, Entra ID sign-in, role mapping, session policy, backup status, point-in-time restore requests, plan limits and support requests are the System Administrator's, and none of them exposes entity content. A restore needs a Group Enterprise Architect's acknowledgement. And before the product's usefulness is claimed it is measured: a fixed evaluation set of approved tasks is put to the product and to a general AI with the same inputs, an AI judge scores both on six criteria with its reasons stored, and the report lists every task the product lost. It states that the bar is met only when the aggregate condition holds.

What you get

Capability	What it does	Requirements
Source register	Every source with knowledge kind, authority class, licence, version and ingestion date; licence-restricted sources held back from derivation until permitted	FR-KE-01, FR-KE-05
Assertions with evidence	Every fact an assertion with evidence naming the source and its exact location; no more than fifteen words of external text stored	FR-KE-02, FR-KE-06
Graded confidence	From *inferred by AI* to *confirmed by accountable person*; who confirmed, and when, on the record	FR-KE-03
Separate stores, isolated tenants	Red Bear, Customer and External knowledge apart; Partner and Baseline knowledge private to the tenant that gathered it	FR-KE-04
Conflict detection and resolution	Conflicts raised and never dropped; resolved by the authority order with the rule cited, or escalated as an ADR to the owning board	FR-KE-07, FR-KE-08
Tenant administration	Setup, SSO, role mapping, session policy, backup and restore, subscription and support, all without entity content	FR-TQ-01, FR-TQ-02, FR-TQ-03
Evaluation set	Approved tasks scored for the product and a general AI on six criteria by an AI judge; every task lost is listed; no claim unless the aggregate condition holds	FR-TQ-04, FR-TQ-05

Screens



S-10 Evidence & sources: Evidence and sources: the confidence mix across the ontology, assertions awaiting an Evidence Owner, and the registered sources with their kind, authority and licence. Prototype render on synthetic data for fictional entities.

RETURN ON INVESTMENT

Three measures carry this pillar, each reported by the product itself and none carrying a target in this document: **the confidence mix**, the share of assertions at each level and how fast inferred facts are confirmed; **the conflict count and its resolution**, how many are open, how many were settled by the authority rule and how many needed a board; and **the evaluation result**, the product's mean against a general AI's on relevance, completeness, accuracy, evidence, clarity and usefulness (FR-TQ-05). What the layer makes possible is a number on an executive screen that can be traced to its source rows in no more than four clicks (NFR-05); a recommendation whose customer facts, patterns and external evidence are each labelled; and a model that can be handed to an auditor, a board or a successor without anyone having to remember where it came from.

Platform engines

Knowledge and evidence consumes the Red Bear Platform's Provenance and evidence service (assertion, evidence and confidence, RBP FR-PV-01; promotion by a person only, RBP FR-PV-02; the three kinds of knowledge, RBP FR-PV-04), the Organisation Information Gathering engine's conflict handling (RBP FR-OG-07), the Research engine's evaluation harness and quotation limit (RBP FR-RS-05, RBP FR-RS-07), Identity and tenancy (RBP FR-ID-03) and the Audit service (RBP FR-AU-01). Platform behaviour is specified in the platform functional specification, not here.