

RED BEAR ARTIFICIAL INTELLIGENCE

ABU DHABI, UNITED ARAB EMIRATES

Enterprise Ontology for Federated Organisations

One essence, many entities: governments, groups and alliances that must act as one

Why a federation cannot see itself in its architecture repositories, what an enterprise ontology adds, and how Red Bear: Enterprise Ontology Management builds and keeps one shared essence across every entity, partner and peer.

Empower. Trust. Verify.

White paper EOM-WP-01 · Red Bear white paper series · Version 0.1

26 September 2026

Summary

A federation is a set of legally distinct entities that must nevertheless act as one. 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: who does the same job, with how many systems, from how many vendors, and what a change in one place does to the others. This paper explains what an enterprise ontology is, why it is the only description that survives reorganisation, and how Red Bear EOM discovers, captures, checks and maintains one shared essence with per-entity implementations, so that overlap is found by function rather than by name and every fact carries its evidence.

WHO SHOULD READ THIS

- Chief architects and heads of digital in governments, groups and alliances
- Entity architects who are asked for a view they cannot produce alone
- Transformation and consolidation programme leads
- Board members who approve cross-entity change

WHAT YOU WILL BE ABLE TO DO AFTERWARDS

- Explain the difference between an architecture and an ontology to a board
- Recognise the six failure modes of federated architecture work
- Describe how a shared essence with per-entity realisations is built and kept
- State what the AI is allowed to do and what a person must confirm

1 What it is

An **enterprise** is any collection of organisations that share common goals: a government agency, a department, a group of governments, an alliance or a supply chain, including its partners, suppliers and customers (TOGAF Standard, Introduction and Core Concepts §1.1). Most enterprises of consequence are therefore federations. A government is a set of departments, authorities and funds, each with its own legal identity, budget and systems. A corporate group is a set of subsidiaries and divisions. Both share partners, vendors and, increasingly, platforms: one ERP, one identity service, one payment gateway.

Enterprise architecture describes such an enterprise in terms of its implementation: capabilities, processes, applications, data, technology and the suppliers behind them, with a baseline, a target and the principles that constrain the journey between them (TOGAF Standard, Architecture Development Method §1.5.3; Greefhorst and Proper, Ch 2 §2.5.4). It is indispensable and it is not enough, because two entities that do the same job with different systems, different job titles and different process diagrams look different in that description. Their overlap is invisible.

Enterprise ontology is the understanding of an enterprise *abstracted from its implementation* (Dietz and Mulder, glossary). It asks what the enterprise is responsible for and who is accountable for each result, and ignores how that is currently arranged. In the DEMO methodology every result is produced by a transaction: one actor role requests it, one actor role executes it, and the two follow a universal pattern of request, promise, declare and accept (Dietz and Mulder, Ch 4 §4.4.4). A transaction kind with its executor role is the molecule of an organisation. A whole organisation is a modest number of these, arranged in trees and joined by three kinds of dependency: who may request what, who must wait for what, and who may read what (Dietz and Mulder, Ch 4 §4.4.6). This is the **essential model**. DEMO claims it is more than 90% smaller than a conventional process model of the same organisation (Dietz and Mulder, Ch 9 §9.2), and, more importantly for a federation, it does not change when systems and departments are reorganised.

A federated enterprise ontology is therefore two things held together: **one shared essence**, the transaction kinds the federation as a whole is responsible for, and **many implementations**, each entity's realisation of those transaction kinds through its own capabilities, processes, systems and vendors. Dietz and Mulder show that a single Cooperation Model can serve an entire alliance because it abstracts from each member's hierarchy; SkyTeam used one to implement lounge access across its airlines and to nudge other programmes toward a common design (Dietz and Mulder, Ch 12 §12.2.3–12.2.4). Red Bear EOM is built on the same idea, with one addition: cross-entity views are always computed from the per-entity facts and never stored as a merged model, so no entity's truth is overwritten by another's.

Key points

- An architecture says how each entity is implemented; an ontology says what it is responsible for
- The essential model is small and stable; it is the only description that survives reorganisation
- A federation needs one shared essence, per-entity realisations, and computed cross-entity views

2 Why it matters

The cost of a federation without a joined-up view is paid three times. It is paid in **duplicated investment**: the same outcome funded in two or three entities because nobody could see that they were the same outcome. It is paid in **fragile change**: a modification in one entity travels along dependencies to others that were not in the room, and breaks them. And it is paid in **unaccountable decisions**: architecture choices whose reasons were never recorded against the evidence, so they cannot be reconstructed, defended or learned from.

TOGAF recognises the scale of the problem directly. Very large federated enterprises are to be integrated through an integration framework and a publish-and-subscribe governance model with compliance assessment and dispensation, and governments are named as a typical case (TOGAF Standard, Architecture Development Method §1.5.1; Applying the ADM §4.3). What TOGAF does not supply is the unit of comparison. Its business architecture is expressed in functions, processes and services that each entity names differently, and the standard itself concedes that its artefacts are largely implementation-oriented (TOGAF Standard, Architecture Content §3.6.6–3.6.7, Red Bear ingestion assessment). The transaction kind supplies that unit. Two entities that realise the same transaction kind are doing the same job, whatever they call it.

For a government the stakes are civic as well as financial. Dietz and Mulder's cases show what happens when the essence is not understood: the British Post Office treated system-derived facts as original facts and prosecuted hundreds of its own subpostmasters on them (Dietz and Mulder, Ch 10); the Dutch childcare allowance scheme made the decline, reject and revoke states of its transactions practically unreachable, and families paid for it (Dietz and Mulder, Ch 12 §12.3). Both are failures of ontology before they are failures of technology. A federation that knows what each entity is responsible for, and who is accountable, is better placed to avoid them.

Question a group asks	Why the repository cannot answer it	What the ontology adds
Which entities issue licences, and with how many systems?	Each entity names the capability differently and models it to a different depth	One transaction kind, *trade licence is issued*, realised in each entity by its own systems
Which vendor is in all of them?	Vendors are recorded per contract, per entity, if at all	Vendors linked to the systems that realise each transaction kind, across entities
If the shared payment gateway changes, who is affected?	The dependency is a wait on a fee confirmation, not a system interface	Wait links between transaction kinds carry the dependency across entity boundaries

Question a group asks	Why the repository cannot answer it	What the ontology adds
Which entity should the others converge on?	There is no common measure of maturity	Maturity and completeness against chosen baselines (EOM-WP-03)
Who decided, and on what authority?	Decisions live in minutes, disconnected from the model	Every change to the essence is a governance record with its evidence (EOM-WP-05)

3 The challenges

Six things make federated ontology work hard, and each has a failure mode that is visible in most large organisations today.

Challenge	Why it is hard	What goes wrong without a method
Discovering what exists	Knowledge is spread across contracts, org charts, ERP configuration, SharePoint sites, people's heads and the public web. Partners and peers publish only fragments	The model starts from whoever answered the survey; entities that never replied are simply absent
Capturing it faithfully	Every fact has a source, a date and a level of confidence. AI can extract quickly but cannot vouch; people can vouch but not at scale	Derived or inferred statements are treated as facts (Dietz and Mulder, Ch 10)
Maintaining it	Organisations reorganise, vendors are replaced, people move on. A model that is right today is stale in a year unless something watches the sources	The repository becomes an archaeology site; decisions are taken on it anyway
Understanding the impact of change	A change travels along dependencies: a delayed transaction blocks the one waiting for it; a changed data bank misinforms every reader; a shared platform change lands in every entity	Impact is guessed from memory; the entities not in the room are the ones that break
Driving standardisation	Standardising the wrong thing is worse than not standardising: DEMO warns against standardising social acts that need judgement (Dietz and Mulder, Ch 6 §6.3); Greefhorst and Proper against mandating storage when only exchange need be common (Greefhorst and Proper, Ch 6 §6.6.2.2)	Convergence is announced but never measured; overlap is funded again
Knowing who decides	A change to what an entity is responsible for is governance; a change to how it is implemented is management. They need different authority and evidence (Dietz and Mulder, Ch 9 §9.6.2–9.6.3)	Boards decide operational detail while essential changes slip through as IT projects

A seventh challenge is specific to federations: **sovereignty**. Entities are legally distinct and will not accept a model in which another entity, or a central team, can overwrite their facts. Any method that depends on a single merged repository fails at the first dispute over whose version is right.

4 How Red Bear EOM helps

Red Bear: Enterprise Ontology Management is a hosted application, with a Claude-based analysis layer, that 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. Against the challenges above it does the following.

Challenge	What EOM does	Module
Discovering	Three entity types, Managed, Partner and Baseline, each discoverable by description, exact name or keyword search across Google, Claude-assisted search and the UAE business registers; every result is a candidate until an analyst accepts it, with duplicate checks against the ontology	ET Entities
Capturing	Five confidence levels from *inferred by AI* to *confirmed by accountable person*; questionnaires, workshops, document submissions and connectors each carry provenance; two reviewers before intake content becomes fact	KE Knowledge, QI Intake
Maintaining	Source monitors on a schedule per entity type; significant events raise reassessment tasks; stale evidence is marked on every screen; the annual ontology review is a built-in cycle	UN Understanding, PN Partners
Impact of change	Impact analysis over the landscape and, since v1.1, the dependency trace over the essence: wait, access and initiator links with a rule and evidence per step	AN Analysis, VZ Visualisation
Standardisation	Principles with conformance scoring, reference entities chosen per capability, patterns with their reasons, and gaps and roadmaps per entity that can be harmonised (EOM-WP-04)	PL Principles, RP Recommendations
Who decides	The two-label rule on every change; a delegated-authority table per tenant; boards or the Group EA as approver per decision class; an audit trail of who decided	GV Governance, RA Roles
Sovereignty	Every element belongs to exactly one entity; cross-entity views are computed, never stored; Partner and Baseline knowledge is private to the tenant that gathered it; role-based entity scoping	RB-MM §2, RA Roles

Standing rule

AI proposes. A named person decides. EOM can find, compare, draft and recommend. It cannot confirm a fact, accept a recommendation or approve a change. Every one of those carries the name and authority of a person.

5 How it is used

The working pattern is a first quarter of establishment followed by a steady state. The roles are the ones most groups already have: a **Group Enterprise Architect** who owns the shared essence and the cross-entity view, an **Entity Enterprise Architect** for each Managed entity, a **Group Enterprise Manager** who reads the results and takes them into decisions, **consultants** who gather and model under review, **boards** at team, entity and federation level, and a **System Administrator** who runs the tenant without seeing its content.

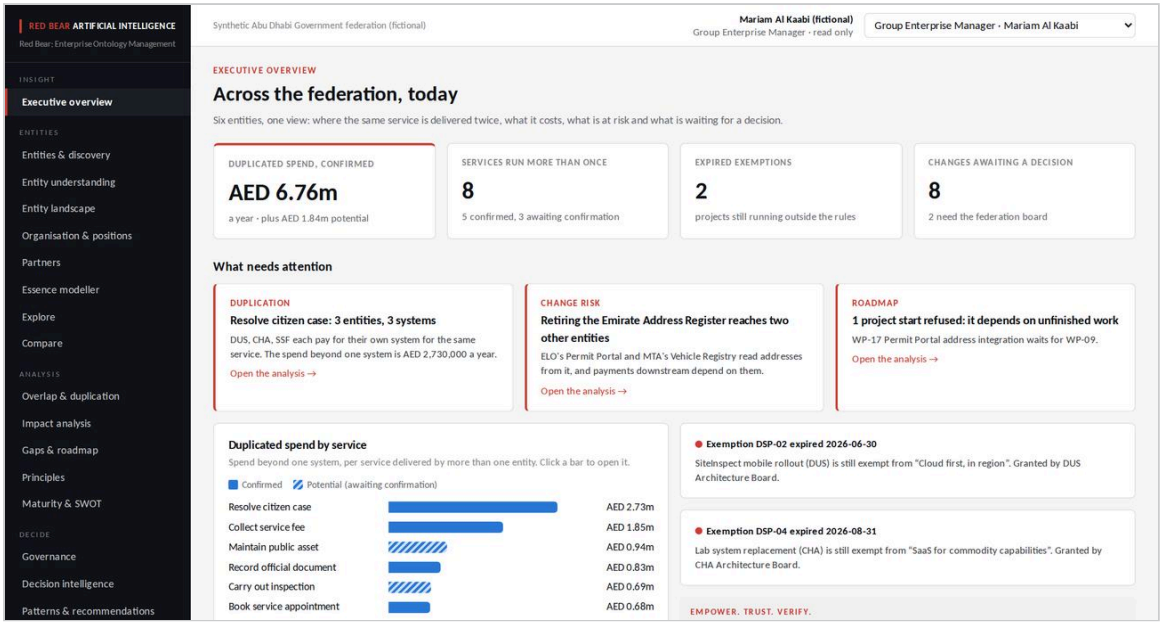
When	Who	What happens
Week 1–2	SA, GEA	Tenant set up; roles assigned; Microsoft 365 and Oracle Fusion connectors configured with admin consent, read-only; the group's existing standards loaded as customer sources with the highest authority
Week 2–4	GEA, CON	Managed entities registered; discovery run for partners and baselines; candidates reviewed; each entity starts at the Awareness phase
Month 2	EEA, CON	Questionnaires sent, documents submitted, workshops held; AI extractions reviewed; entities move to High-Level and Low-Level Assessment; the landscape (L3) fills
Month 2–3	GEA, EEA	The shared essence (L1) is modelled where overlap analysis needs it, capability by capability; realisation links recorded per entity; first overlap and duplicated-spend reports

When	Who	What happens
Month 3	GEM, boards	First cross-entity review: overlaps, findings, proposed reference entities; governance-class changes routed under the two-label rule
Steady state	All	Source monitors run; reassessment tasks close; quarterly workshops; the annual ontology review; every decision recorded with its impact and, later, its outcome

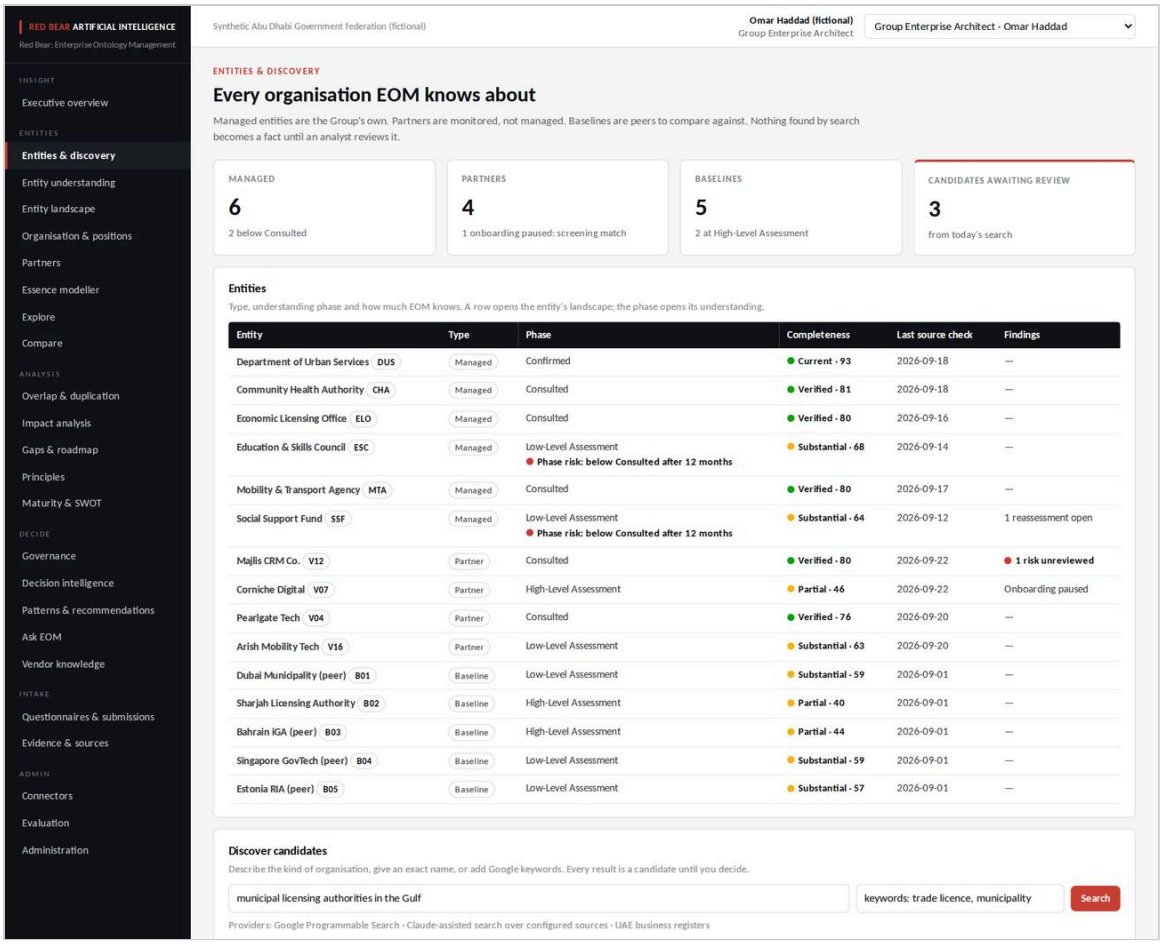
The essence is built **progressively**, not up front. The TOGAF landscape is the entry point and the DEMO layer is added where overlap detection needs it (design decision MC-1). This matters in a federation, where modelling skill is uneven: an entity can be useful to the group at the landscape level while its essence is still being modelled.

6 The user experience

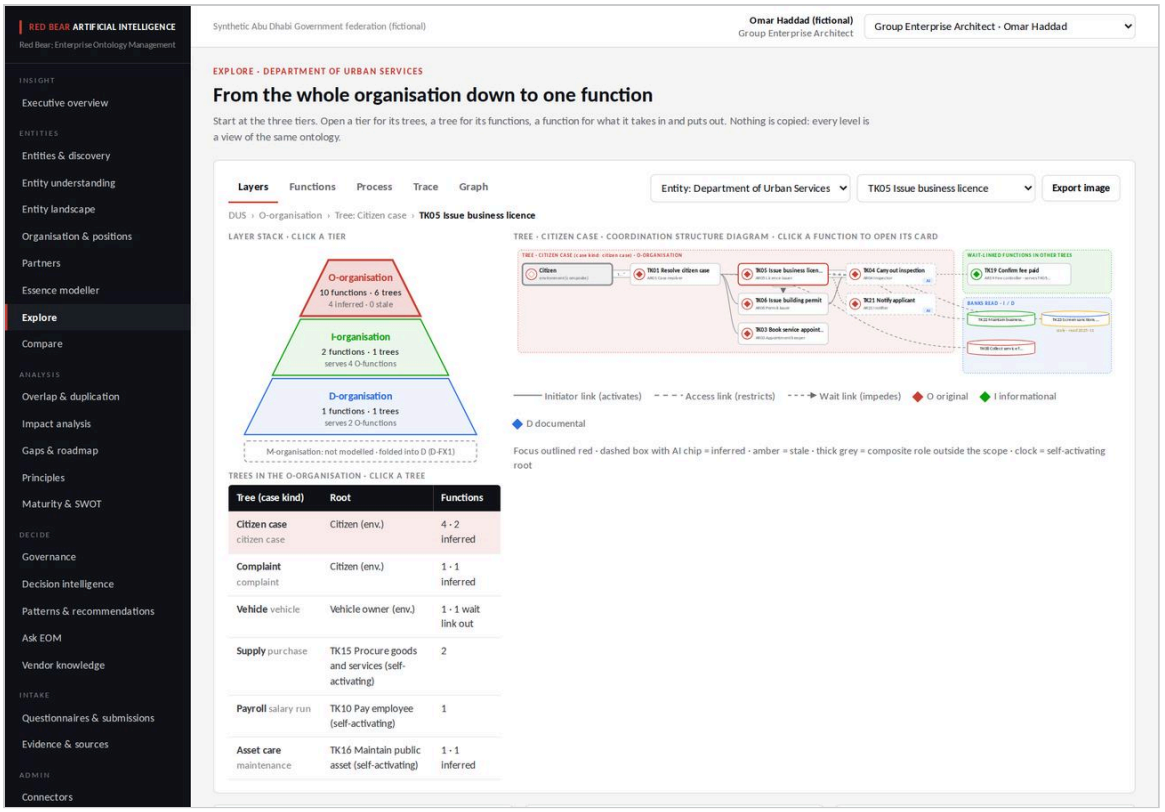
Users move progressively from insight to evidence: **Executive insight → Analysis → Architecture object → Relationship → Evidence**. Ontology vocabulary stays hidden until the user asks for it, and every insight can open the evidence behind it, labelled Red Bear, Customer or External.



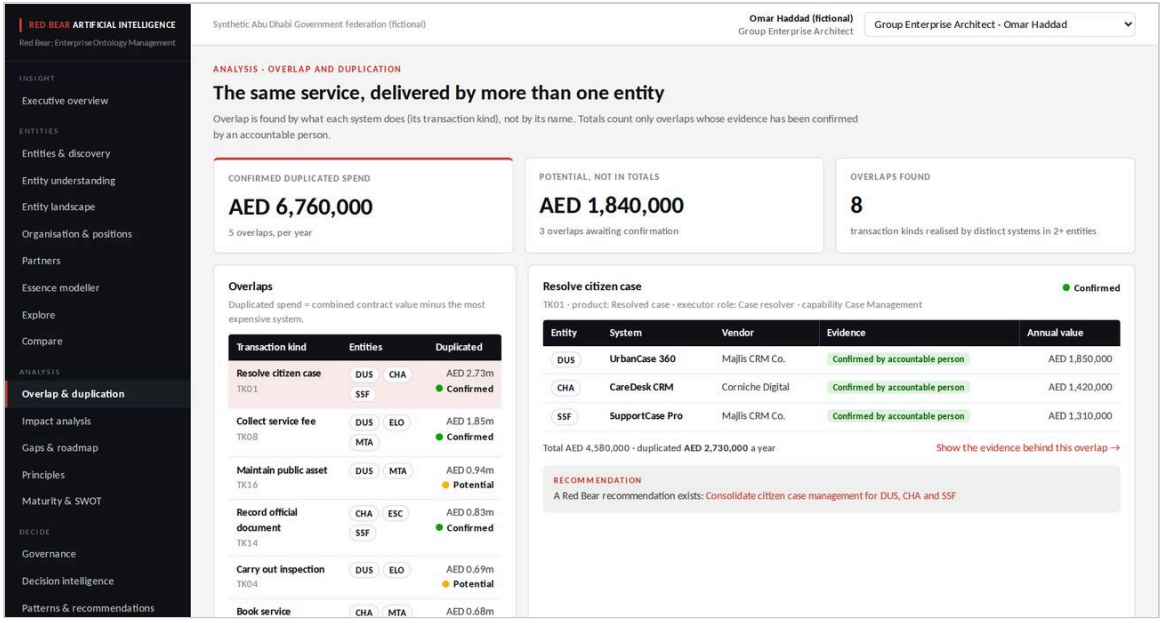
overview: S-01 Executive overview. The federation at a glance: duplicated value, overlaps, risks and decisions waiting, each opening the analysis behind it. Prototype render on synthetic data for fictional entities.



S-13 Entities & discovery: S-13 Entities & discovery. Every organisation EOM knows about, by type and understanding phase, with the completeness score and the last source check. Below, discovery candidates from a search, each with its duplicate check, waiting for an analyst to accept or reject. Prototype render on synthetic data for fictional entities.



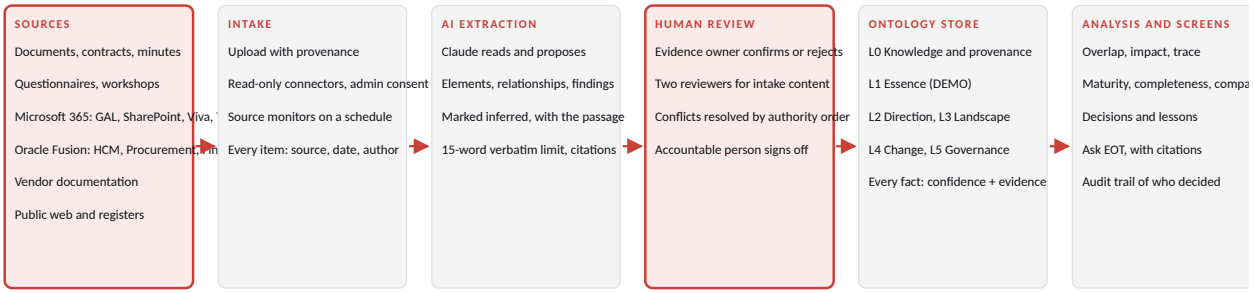
UC-33: S-19 Explore, Layers mode. The layered organisation as a pyramid, the interaction trees in the chosen tier, and one tree drawn as a Coordination Structure Diagram in DEMO notation. Inferred links are dashed with an AI chip; stale evidence is outlined amber. Prototype render on synthetic data for fictional entities.



overlaps: S-03 Overlap & duplication. Entities that realise the same transaction kind, the systems and vendors behind each, and the duplicated spend. Detected by function, not by name. Prototype render on synthetic data for fictional entities.

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

Nothing enters the ontology without a source, and nothing becomes a fact without a person. The pipeline below is the same for every entity type; what differs is which sources are available and who may confirm.



The data pipeline. Sources on the left are read through uploads and read-only connectors; AI extraction proposes elements and relationships with the passage they came from; a person confirms or rejects; only then does the fact enter the store, carrying its confidence and evidence into every analysis and screen.

!Sources by entity type

Entity type	Typical sources	Who may confirm
Managed	The entity's own documents, questionnaires, workshops, Microsoft 365 (GAL, SharePoint, Viva, Teams opt-in) and Oracle Fusion (HCM, Procurement, Financials, PPM) through admin-consented read-only connectors	The entity's accountable person; the Entity EA for landscape facts
Partner	Contracts and deliverables held by the group, the partner's public materials and documentation, procurement records in Fusion, screening sources	The Group EA; the partner is a subject, not an author
Baseline	Public web, annual reports, published strategies and registers; nothing is requested from the peer	The Group EA; confidence is capped at *documented in external source*

How it is checked



The five confidence levels. Every fact carries one; every screen shows it. A fact only rises to *confirmed* when a named accountable person signs it off.

- **Provenance on everything.** Each assertion has at least one piece of evidence, and each piece of evidence names its source, date and author. A screen that shows a number can open the passage it came from.
- **AI is marked.** Whatever Claude extracts or infers is labelled inferred, drawn dashed on diagrams, and cannot be cited as a fact in a report until reviewed. Quotations are limited to fifteen words and always carry a citation.
- **Two reviewers for intake.** Questionnaire answers and submitted documents are reviewed by the Entity EA and a Group EA before anything they contain enters the ontology.
- **Authority order for conflicts.** When sources disagree, a fixed order decides: regulatory, then the customer's own standard, then Red Bear's, then vendor material, then external sources, frameworks, cookbooks, academic work, historical records and advisory notes. Where two sources carry equal authority an architecture decision record is opened rather than a silent choice being made.
- **Staleness is visible.** Each source has a period; a fact whose source has not been seen within it is marked stale on every screen, and a reassessment task is raised.
- **Connectors are least-privilege and read-only.** Admin consent, application permissions limited to read, no write-back to any source, and a System Administrator who configures connections but never sees the content they bring

in.

8 Benefits realised

The benefits are the ones the tool itself measures, so each can be reported from EOM rather than asserted.

Benefit	How it shows up	Measured by
Overlap made visible	Entities realising the same transaction kind listed with their systems and vendors	S-03 overlap count and duplicated spend (AN-01, AN-02)
Duplicated investment stopped	Consolidation candidates with the contracts and exit costs behind them	Duplicated-spend trend by quarter
Change made safe	Impact and dependency traces run before a change is approved	Share of governance decisions with an impact estimate attached
Knowledge that stays current	Stale facts fall; reassessment tasks close within their period	Completeness KPI per entity; stale share
Decisions that can be defended	Every essential change carries its evidence, approver and, later, its outcome	Decision records with outcome reviews (EOM-WP-05)
Faster answers	Ask EOM answers group-level questions with citations in seconds	Evaluation set: EOM vs general AI on six criteria (TQ-04)
Sovereignty respected	No entity's facts overwritten; cross-entity views computed	Audit trail of refused actions and cross-entity reads

9 Risks managed

An ontology tool introduces risks of its own. The controls below are part of the specification, not afterthoughts.

Risk	Control in Red Bear EOM
Over-trust in AI inference	Inferred facts are labelled and drawn dashed; they cannot be cited as fact, and confidence is lowered through every inferred link in a trace. Claude proposes; a person confirms
Derived facts mistaken for original facts (the Post Office failure)	Provenance distinguishes human-declared original facts from system-derived outputs; recommendations are never presented as facts
Standardising the wrong thing	Principles are traceable to goals and scored for conformance; DEMO's warning against standardising social acts and Greefhorst and Proper's *exchange not storage* are recorded as patterns the recommender applies
One entity overwriting another	Per-element ownership; computed cross-entity views; entity-scoped roles; Partner and Baseline knowledge private to the tenant
Governance bypassed as an IT change	The two-label rule on every change; the delegated-authority table; approver per decision class; audit trail
Data residency and AI processing	Application and stores in UAE Azure regions. Claude runs on the Anthropic API in development and in Microsoft Foundry inside the Azure tenancy for live; neither is UAE-resident today, which is stated as NFR-12 and mitigated by data minimisation and no training on customer data

Risk	Control in Red Bear EOM
Connector over-reach	Admin consent, least privilege, read-only, no write-back; contact details verified against a live account before use
Stale model used for decisions	Source periods, staleness marking, reassessment tasks, and the completeness KPI shown beside every entity
Claims the tool cannot support	The evaluation set: EOM's answers are scored against a general AI on a fixed rubric before any claim of usefulness is made

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
Entity	A legally or organisationally distinct member of the federation: Managed (the group's own), Partner (monitored) or Baseline (a peer to compare against)
Enterprise ontology	The understanding of an enterprise abstracted from its implementation (Dietz and Mulder, glossary)
Essential model	The ontology of the O-organisation: transaction kinds, actor roles and their coordination structures, independent of systems and org charts
Transaction kind	A kind of result the enterprise is responsible for, with exactly one product kind and one executor actor role
Realised by	The link from a shared transaction kind to the capability, process and systems that implement it in one entity (RB-MM bridge)
Scope of Interest	The part of the federation being modelled: one entity, or a cross-entity scope
Provenance	The source, date, author and confidence attached to every fact
Authority order	The fixed precedence used when sources disagree: regulatory, customer standard, Red Bear standard, vendor, external, framework, cookbook, academic, historical, advisory
Two-label rule	Every change carries whether it alters the essential model (governance) or only its implementation (management), and its TOGAF change class

Term	Meaning in this paper
Completeness KPI	How much EOM knows about an entity, computed from the heatmap of known, missing, stale and uncertain cells

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 (this paper)
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?
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