Intelligence Warehousing

A three-layer knowledge architecture
for enterprise decision agents

98%	4 wk	3
Decision Accuracy	Time to Go-Live	Knowledge Layers

Questt AI · February 2026

Contents

Executive Summary

> Enterprise AI agents fail not because models are incapable, but because they lack the structured business knowledge — formulas, decision rules, exception logic, persona authority — needed to produce grounded decisions.
>
> The Intelligence Warehouse is a three-layer knowledge architecture (business ontology, metrics with formulas, decision rules) that sits between enterprise data systems and AI agents. In production deployments across FMCG and retail, this architecture has delivered 98% decision accuracy with a 4-week implementation timeline.
>
> This paper describes the architecture, explains why it works, and compares it against four common alternatives.

Background

Every enterprise deploying LLMs wants the same thing: agents that make real business decisions. The models are capable. The bottleneck is what they know at inference time.

In deploying decision agents across FMCG, retail, and consumer goods companies, agents failed not because of model limitations, but because of knowledge architecture limitations.

⚠ Data without logic

Agents connected to DW retrieved accurate numbers but lacked business rules to interpret them. Saw DOI = 140 but didn't know 120 was the trigger.

⚠ Dashboards without structure

Agents reading from BI had pre-computed metrics but couldn't recompute at different scopes or trace decision authority.

The gap

Enterprise Data(DW / ERP / BI) → **IntelligenceWarehouse** → **GroundedAI Decisions**

> *Most architectures try to go directly from data to decisions. The Intelligence Warehouse is the missing middle layer.*

The Intelligence Warehouse

A three-layer knowledge architecture, implemented as a knowledge graph, that provides AI agents with everything they need to make grounded decisions — not just data, but the logic, formulas, and rules that turn data into action.

L3 RULES	**How the business acts on what it knows** Each decision rule specifies: trigger conditions, conditional action logic (e.g., if margin > 35% then markdown 15%; if 25–35% then 10%), persona authority (who proposes, who approves), escalation paths, and exceptions. Rules cross-reference each other — a markdown rule may invoke redistribution, or a new-launch rule may override it.
L2 METRICS	**How the business measures itself** Each metric carries: an explicit formula (Sell-Through = Secondary Sales Qty ÷ Opening Stock × 100), an applicable scope (SKU × Distributor × Month), threshold bands (green/amber/red), and frequency. Composite metrics encode their component weights as graph edges.
L1 ONTOLOGY	**What the business is** Product hierarchy (Company ⟶ BU ⟶ Category ⟶ Brand ⟶ SKU), geography, sales organization, channels. Each node carries typed properties. Edges encode relationships: which RSM manages which region, which SKU belongs to which brand.

Without Layer 1: agents can't resolve multi-hop queries or scope metrics correctly.

Without Layer 2: agents approximate formulas, producing inconsistent computations.

Without Layer 3: agents produce plausible recommendations that don't match institutional policy.

How a query traverses the layers

"Should we markdown GlowMax Men 50ml in North zone?"

	Step	What the agent retrieves
1	**Resolve entities**	GM-M-50ml ⟶ GlowMax ⟶ Skin Care. North ⟶ 3 regions ⟶ RSM assignments.
2	**Compute metrics**	DOI = 140d. Sell-through = 22%. Margin = 30%. All via explicit formulas.
3	**Apply decision rule**	D1: DOI > 120 ✓ AND ST < 30% ✓ . Margin in 25–35% tier ⟶ markdown 10%.
4	**Check exceptions**	Stock > 100 ✓ . Not new launch ✓ . Not strategic ✓ . Proceed.
5	**Identify authority**	RSM proposes ⟶ ZSM approves. RSM for North = [from org graph].

No single layer is sufficient. The answer requires all three.

Why 98% accuracy in 4 weeks

On accuracy

Most enterprise AI failures are context failures, not model failures. When agents lack rules, they improvise. Here's what changes:

Error type	Without IW	With IW
Hallucinated entities	✗ Frequent	✓ Rare
Wrong formula	✗ Frequent	✓ Rare
Missing decision rule	✗ Near-certain	✓ Rare
Persona / authority error	✗ Near-certain	✓ Rare
Exception not checked	✗ Near-certain	– Occasional

> When the agent has the right entities, formulas, and rules in context, the remaining errors are edge cases in exception logic or genuinely novel scenarios. That's how accuracy reaches 98%.

On speed to go-live

The three layers are built in parallel. Layers 1 and 2 draw from existing sources. Layer 3 requires SME interviews but its scope is bounded.

	Layer	Input	Method
Wk 1–2	**L1: Ontology**	ERP, org charts	Extract + curate
Wk 1–2	**L2: Metrics**	KPI definitions, BI	Codify + validate
Wk 2–4	**L3: Decision Rules**	SME interviews, policy	Extract + iterate
Wk 3–4	**Integration**	Layers + connectors	Wire + test

A typical FMCG business has 10–25 core decision rules. This is a weeks-scale project, not a months-scale one.

Two properties worth noting

01 Curated knowledge compounds

Auto-constructed graphs reflect textual co-occurrence, not business truth. A curated graph encodes precise relationships — and each addition improves answers across the entire graph. A new exception added to one rule immediately affects every related query for every SKU in every region.

02 Knowledge separated from model

When business logic lives in prompts, it's fragile. Moving it into the graph makes rules versioned, auditable, queryable. Change one policy node and every agent reflects it instantly. Swap the underlying LLM without touching the knowledge layer.

This is work that cannot be automated or shortcut — which is precisely what makes it durable.

Model-agnostic by design. Future-proof by architecture.

How alternatives compare

Test question: *"Should we markdown GlowMax Men 50ml in North zone? DOI = 140, sell-through = 22%, margin = 30%."*

A correct answer requires: identifying the relevant rule, checking trigger conditions, applying margin-tier logic, verifying exceptions, and naming the approval chain.

LLM + Data Warehouse

✓ Accurate data retrieval. Gets the right numbers when schema is well-documented.

✗ No rules, thresholds, or personas. Agent confidently recommends 15–20% markdown instead of the correct 10%.

LLM + BI Tool

✓ Pre-computed metrics with correct formulas. Numbers are reliable.

✗ Fixed granularity, no recomputation. No conditional logic. Agent has metrics but must improvise the action.

LLM + Prompt Engineering

✓ Effective for small, contained rule sets. Rules are followed when in context.

✗ Doesn't scale. Rules compete for context window. Not version-controlled or auditable. Fragile in production.

LLM + Ontology Only

✓ Excellent entity resolution. Correctly maps SKU → Brand → Category → RSM.

✗ Knows what exists but not how to measure it or what to do about it. Improvises formulas and decision logic.

Scorecard

	+ DW	+ BI	+ Prompt	+ Ontology	IW
Data retrieval	✓	✓	–	✓	✓
Entity resolution	–	–	–	✓	✓
Metric formulas	✗	✓	–	✗	✓
Decision rules	✗	✗	–	✗	✓
Exception handling	✗	✗	–	✗	✓

Persona / authority	✗	✗	–	–	✓
Auditable & versioned	✓	✓	✗	✓	✓
Model-agnostic	✓	✓	✗	✓	✓

✓ Full – Partial ✗ Absent

The bottom line

Data warehouses tell you what happened. BI tools tell you why. The Intelligence Warehouse tells AI agents what to do about it, who should do it, and what exceptions to watch for.

The result: agents that don't just sound knowledgeable, but actually are.

About Questt

Questt builds AI decision intelligence for enterprises. Our platform converts siloed enterprise data into LLM-ready knowledge graphs, enabling autonomous agents that unify data, automate workflows, and deliver measurable ROI across FMCG, retail, financial services, and manufacturing.

questt.com · contact@questt.com