RJL SOLUTIONS

FIELD GUIDE / BUSINESS MEANING FOR AI

Ontology or Semantic Layer?

What each one tells an AI, why agents need both, and which to build first.

MEANING · METRICS · OWNERSHIP

Richard White
RJL Solutions · rjlsol.com

VERSION 1.1 · PUBLISHED EDITION
NOVEMBER 2026

01 / THE ANSWER FIRST

A semantic layer says how to count. An ontology says what is being counted.

People ask the question because vendors now sell both, often under each other's names. The difference is practical. A semantic layer defines metrics, dimensions and the joins and calculations behind them, so every tool and agent computes "open claims" or "net written premium" the same way. An ontology defines the things in a business, what each one means, how they relate and which rules apply, so an agent knows that the customer on a claim is the claimant and not the employer that bought the policy.

Three findings

- They answer different questions. A semantic layer answers "how is this number calculated?" An ontology answers "what is this, how does it relate to that, and which rule applies?" Agents need both, because they count things and they act on things.
- The ontology should feed the semantic layer, not the reverse. Meaning changes when the business changes; calculations change when the data changes. Kept apart, each can be owned, approved and versioned by the right people.
- The market is converging on both at once, but maturity varies. In 2026 the main data platforms shipped semantic layers and announced ontologies, and the open semantic interchange specification, now Apache Ossie, covers both. Several ontology products are still in preview or announced without a date.

RULE OF THUMB If the question is about a number, start with the semantic layer. If it is about who, what, which or whether, start with the ontology. If an agent will act, you need the ontology first, because actions depend on roles, relationships and rules, not on aggregates.

02 / SIDE BY SIDE

Same goal, business meaning for AI. Different jobs.

Both exist to stop people and machines from interpreting the same data in different ways. They do different work, change for different reasons and belong to different owners.

	SEMANTIC LAYER	ONTOLOGY
Question it answers	How is this metric calculated, at what grain, from which tables?	What is this thing, what does it mean here, how does it relate to other things, which rules apply?
Building blocks	Measures, dimensions, entities for joining, time grains, filters	Concepts and their definitions, relationships, roles, identifiers, constraints, valid time, rules
Typical example	Open claims = count of claims meeting the reserving definition at month end	A claimant is the person who suffered the loss; a policyholder owns the policy; one person can be both or neither
What it produces	Generated SQL and consistent numbers	Shared meaning for queries, documents, workflow checks and agent tools
Main consumers	BI tools, dashboards, analytics agents	Agents, search and retrieval, workflows, integration, data products
Natural owner	Analytics engineering, with finance or the metric owner	Business owners and data stewards, with architecture
Changes when	Tables, pipelines or calculation rules change	The business, product, regulation or organization changes
Fails without it	Two dashboards, two numbers	The right number about the wrong thing, or an action on the wrong party
Open standards	Apache Ossie (incubating); dbt MetricFlow (open source)	W3C OWL and SKOS; industry models such as FIBO; Apache Ossie also covers ontology

A glossary is the start of both: agreed terms and definitions. A knowledge graph is one way to store an ontology together with the facts it describes. Neither is a substitute for deciding who owns each definition.

03 / ONE CLAIM, THREE QUESTIONS

The synthetic claim CL-9055 needs both.

Demo 1 on rjlsol.com asks three questions about one synthetic disability claim. Each answer goes wrong for a different reason, and each needs a different fix.

QUESTION	WITHOUT APPROVED MEANING	WHAT FIXES IT	WHICH LAYER
Show me the customer's prior claims for CL-9055	41 claims: every claim on the employer's group policy, because the data calls the employer the customer	The role: for an examiner, the customer is the claimant, mapped from member_id	Ontology (roles and relationships)
Which claims-handling procedure applies?	Version 5, the newest valid procedure	The rule: the procedure in force on the policy date, version 4	Ontology (rules and valid time)
How many claims are open at quarter-end?	6, the rows with status = Open	Two named measures: 8 for reserving, 7 for service reporting, each with its own definition	Semantic layer, resting on the ontology's definition of an open claim

The first two answers come from correct data and still mislead, because nothing told the model what the words mean in this business. No metric definition fixes them. The third is a calculation problem, and a semantic layer is the right place to fix it, but only once the business has agreed what "open" means for each purpose.

TRY IT Demo 1 at rjlsol.com/products/demo-ai-answers shows the three answers with and without approved meaning. Ontology Studio lets you build the ontology for the same claim and export it.

04 / HOW THEY FIT TOGETHER

Meaning first, calculation second, both before the model answers.

In a governed stack the two layers sit in sequence. The ontology is the source of meaning. The semantic layer implements the metrics that meaning calls for, on a particular platform. Agents and tools draw on both, and access is still decided by the data platform, not by either layer.

LAYER	WHAT IT HOLDS	WHO APPROVES
Ontology	Concepts, definitions, roles, relationships, rules, valid time; one version, exported to each platform	Business owner of each term, with a steward
Semantic layer	Measures and dimensions implementing the approved definitions on each platform	Metric owner, with analytics engineering
Data products	Certified tables and contracts the semantic layer reads	Data product owner
Access control	Who may see which rows and columns; enforced in the engine	Data owner and security; never the ontology
Agents and tools	Use the ontology to choose meaning and tools, and the semantic layer to compute numbers	Agent owner, within the release controls

Four design rules

- Keep meaning and logic apart. A definition says what an active customer is; a calculation says how to count active customers from a given table. Store them separately, with separate owners and versions.
- Approve once, export everywhere. When data sits on more than one platform, approve each definition in one place and generate each platform’s semantic model from it, rather than approving it again in every tool.
- Name purpose-specific measures. Where finance and operations disagree for good reasons, keep two named measures with one owner each, and state which is the default for which purpose.
- Record what was used. Every agent answer should record the definitions and metric versions it relied on, so a reviewer can reconstruct it.

05 / WHO DOES THE COUNTING

The ontology says what to count. It doesn't do the count.

An ontology is a description, not a calculator. It says an open claim for reserving is one with an outstanding reserve above zero at the reporting date, but it runs nothing. Something else has to turn that definition into a number.

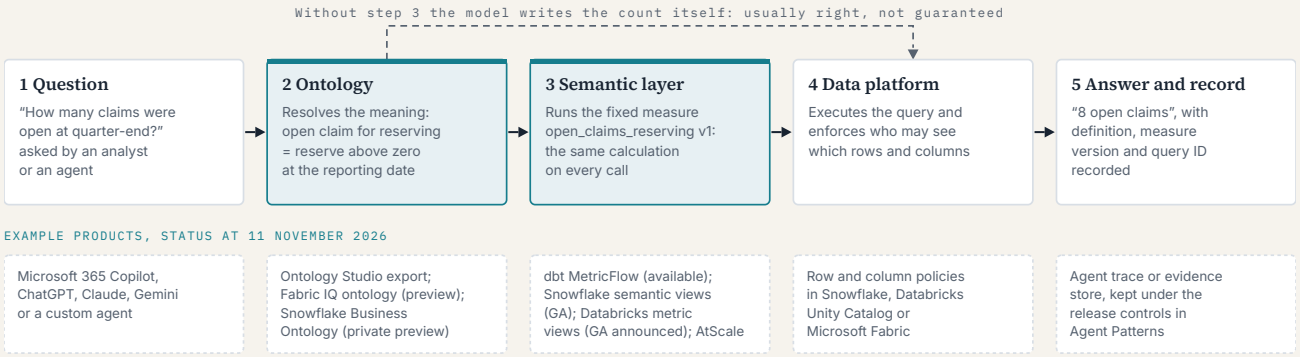


Exhibit: One question through the stack. Products are examples, not endorsements; statuses as reported by each source [4, 5, 6, 7, 10, 15].

WHO COUNTS	WHAT YOU GET	USE IT FOR
The model, from the approved definition	Usually the right answer, but worked out afresh on every call; a second run can apply the rule differently or miss a row	Exploring, drafting, and demos on small data
A fixed measure: a semantic layer metric or a tested query the agent calls as a tool	The same number on every run, from every dashboard and agent, traceable to a measure version	Any number someone will report, reconcile or act on

RULE OF THUMB

The ontology holds the meaning. The semantic layer, or the model if there isn't one, does the count. With a fixed measure behind it, the model's job is to pick the right measure, which is where the ontology helps, and to say which one it used. Demo 1 shows both: the model's answer and a fixed count of 8 and 7 from code.

06 / WHICH TO BUILD FIRST

Start from the questions and the number of platforms.

Most firms will end up with both. The order depends on what the first use case asks and how many platforms hold the data.

	DATA MOSTLY ON ONE PLATFORM	DATA ON SEVERAL PLATFORMS
QUESTIONS MOSTLY ABOUT NUMBERS	**START WITH THE PLATFORM'S SEMANTIC LAYER** Define the governed metrics in the platform you already run. Add a short glossary of the terms the metrics depend on.	**PORTABLE METRIC DEFINITIONS** Author metrics once in a portable format, such as dbt MetricFlow with Apache Ossie export, and generate each platform's model from it.
QUESTIONS ABOUT THINGS, ROLES, RULES OR ACTIONS	**ONTOLOGY ON THE PLATFORM, WITH CARE** Platform ontologies are arriving, but several are in preview. Use one for a bounded first workflow and keep an export.	**PLATFORM-NEUTRAL ONTOLOGY FIRST** Approve meaning in a neutral tool, then export it to each platform's semantic layer and to agents. This is where portability pays most, for example across a private equity portfolio.

Signs you need the ontology now

- The same word means different things to different teams, such as customer, claim, exposure or default.
- An agent will take actions, route work or choose tools, not just report numbers.
- Answers depend on roles and relationships: who is the claimant, which policy governs, which entity is the borrower.
- Rules depend on dates, such as the procedure or rate in force on a given day.

07 / THE MARKET IN NOVEMBER 2026

Both layers are now product categories. Check the status before you plan on one.

In 2026 the major data platforms shipped or announced semantic layers and ontologies for agents. Status varies by component, so ask which part is generally available. Statuses below are as each source reports them on 11 November 2026.

OFFERING	WHAT IT IS	STATUS AT 11 NOVEMBER 2026
Apache Ossie (formerly Open Semantic Interchange)	Vendor-neutral, YAML-based specification for semantic layer and ontology definitions	Accepted into the Apache Incubator in June 2026; v0.1 specification January 2026 [1, 3]
dbt Semantic Layer and MetricFlow	Governed metric definitions executed consistently; MetricFlow open source under Apache 2.0	Available; dbt Labs is a founding Ossie participant [2]
Databricks Unity Catalog metric views	Metrics defined once in YAML and registered in Unity Catalog	General availability of Unity Catalog Business Semantics announced; some cloud documentation still shows preview [4, 5]
Microsoft Fabric IQ	Business-context layer with an ontology, graph and agents over OneLake data	Fabric IQ generally available since June 2026; the ontology item still in preview [6, 7]
Snowflake Business Ontology	One approved definition per term, with an owner, drafts and an audit trail	Private preview, announced 3 November 2026 [10]
Oracle Business Semantic Layer	Semantic layer for the Autonomous AI Lakehouse	Announced 26 October 2026; expected within 12 months [8]
SAP Ontology Manager	Ontologies across SAP and other data	Announced 27 October 2026; no availability date [9]
NVIDIA Auto Ontology	Proposes starting definitions from database metadata	Open source under Apache 2.0 [11]
W3C OWL and SKOS; FIBO	Formal ontology and vocabulary standards; an industry ontology for finance	Established standards [12, 13, 14]

Vendor tools generate drafts quickly. None removes the need for a business owner to approve each definition before agents rely on it.

Meaning is a decision. Treat it like one.

The software is the easy part. What makes either layer trustworthy is that someone accountable approves each definition, drafts stay drafts, and every answer can be traced to the version it used.

Five mistakes to avoid

- Treating the semantic layer as the ontology. Metrics on top of tables do not say who a party is or which rule applies, so agents that act will still guess.
- Building the ontology as an academic project. Start from the questions a workflow must answer and the terms they depend on; a bounded, approved vocabulary beats a complete, unapproved one.
- Approving meaning separately in every platform. The copies drift. Approve once and export.
- Letting AI-suggested definitions reach production. Generated drafts are a fast start, not an approval.
- Using the ontology for access control. Meaning helps choose the right answer. Whether a user may see the data is decided by the data platform.

Before an agent relies on either layer

CHECK	EVIDENCE
Every term the agent uses has a named owner	Owner and approval date on each definition
No unapproved draft can reach the agent	Test that a draft definition is rejected
Metrics implement approved definitions	Each measure points to the definition it implements
Definitions are portable	Export with owner, version and approval history
Answers are traceable	Each answer records the definitions and metric versions used

09 / SOURCES

Sources and notes.

Prepared by Richard White. Version 1.1, 11 November 2026; adds who does the counting and the flow exhibit (section 05). Product statuses come from the public sources below and were not confirmed with the vendors; inclusion is not an endorsement. The CL-9055 example is synthetic and comes from Demo 1 on rjlsol.com.

[1] Snowflake, Apache Ossie (Incubating): the new name for Open Semantic Interchange, 8 July 2026. snowflake.com

[2] dbt Labs, What the Open Semantic Interchange (OSI) spec means for metrics, semantics, and AI, 29 January 2026. getdbt.com

[3] Dremio, Apache Ossie: the Open Semantic Interchange joins the Apache Incubator, 2026. dremio.com

[4] Databricks, Unity Catalog metric views (documentation). docs.databricks.com

[5] Databricks, Announcing general availability and open sourcing of Unity Catalog Business Semantics (via daily.dev), 2026. app.daily.dev

[6] The New Stack, Microsoft Build 2026: Fabric IQ goes GA, June 2026. thenewstack.io

[7] Atlan, What is Fabric IQ? (status of the ontology item), October 2026. atlan.com

[8] Oracle, Oracle advances Autonomous AI Lakehouse with Business Semantic Layer, 26 October 2026. oracle.com

[9] SAP, SAP empowers AI developers to build what matters, 27 October 2026. news.sap.com

[10] MarketScale, Snowflake wants every AI agent to agree on what net revenue means, November 2026. marketscale.com

[11] Emergent, NVIDIA Auto Ontology: open semantic layer for AI agent SQL, 2 November 2026. emergent.sh

[12] W3C, OWL 2 Web Ontology Language Document Overview. w3.org

[13] W3C, SKOS Simple Knowledge Organization System Reference. w3.org

[14] EDM Council, Financial Industry Business Ontology (FIBO). spec.edmcouncil.org

[15] Snowflake, Overview of semantic views (documentation); defining and querying semantic views generally available since 2025. docs.snowflake.com

ABOUT

Richard White

I was global firm chief data officer at McKinsey & Company, owning data, AI and platforms end to end, and before that chief data officer at New York Life, where I built the chief data office and reported data risk to risk committees at multiple levels. Earlier I led data at the Federal Reserve Bank of New York and enterprise data governance at Citigroup. I founded RJL Solutions in 2024. My book, *The Day After the Demo*, is free at rjlsol.com/book.

RJL Solutions connects strategy, architecture, engineering, governance and operation through Frame, Prove, Build, then Operate and transfer. Each phase is separately scoped and priced in a written proposal before you commit. The first sixty-minute working session is free.

MEANING **Context Before Query**
Giving agents approved meaning before they query.

OWNERSHIP **Data Product Certification**
The sign-off model definitions should borrow.

AGENTS **Agent Patterns**
Patterns, autonomy and the controls that prove them.

RJL SOLUTIONS

WEB rjlsol.com

EMAIL rwhite@rjlsol.com

Product statuses checked against public sources on 11 November 2026 and not confirmed with vendors. Inclusion is not an endorsement.