


The ODL Is Dead. Long Live the Semantic Layer.

The Operational Data Layer earned its place by unifying data from systems of record and serving it fast — it answers *what is true right now*. AI has changed the question to *why*: agents and analysts need meaning, context and governed access. The semantic layer is the next evolution of the ODL you already have — five components on one platform, MongoDB Atlas — and the shift from hunting for answers to validating ranked, cited hypotheses.

The anatomy — five components, one platform



1 Selective operational data

From the systems of record into Atlas — bring what the questions need. The document model lets each business line extend the record without a migration.

2 Embedded on write

The vector is generated as the document is written; record, vector and index live in one cluster. The AI's own outputs are persisted with a pointer to what they summarised.

3 Retrieve with meaning

Full-text and vector search in one aggregation, fused and reranked. Route by intent — and skip semantic search when the question doesn't need it.

4 Context, then governed access

Five small collections — registered sources, tool catalog, semantic model, resources, audit log — read before any data is touched. Access is through MCP tools: identity on every call, scope enforced inside the tool, never in the prompt.

5 Provenance and memory

Every answer cites collection, _id and timestamp or the warehouse row, and is checked against its evidence before a human sees it. Every tool call is logged with the scope applied.

Beside the path of record — not in it.

Enterprises want AI-ready data, not AI in the path of record. Keep the deterministic core deterministic — pipelines you can audit that behave the same at three in the morning as in the demo — and let agents reason over governed, read-only views of it. If an agent sits in the write path of your system of record, ask who signs off the audit.

Evaluation set before the code

Ground truth first. It is the hardest thing to get from a client and the only way to know you are improving.

Success is a relevance curve

Shown every week, next to the latency SLA it used to be — not a demo at the end.

Three signs it is already happening in your estate

- 1 Your analysts spend more time deciding where to look than acting on what they find.
- 2 You have a vector store, a search index and an operational database that don't share a security model.
- 3 Someone has already pointed an AI assistant at your data — and nobody can say exactly what it can see.

Two of three? You don't have an AI problem — you have a semantic-layer opportunity.

Three first steps — with the data you already have

- 1 **Pick one domain you already serve from an ODL.**
Write down the why questions nobody can answer today. That is your scope.
- 2 **Build the evaluation set before the pipeline.**
Fifty real questions and what a good answer looks like will save you months.
- 3 **Put the agent behind governed tooling from day one.**
Provenance and role-based access aren't phase two — they're the reason you'll be allowed a phase two.

Start with one domain, one evaluation set, one governed tool.