Cosmos Semantic

OPEN SOURCE · SEMANTIC LAYER FOR AI

Your AI agent just gave you a number.

Can you prove it's right?

The same number on every database and platform, with the proof in every answer.

Open source · Apache 2.0

Swipe →

THE PROBLEM

“Revenue” has three definitions. The agent picks one. Nobody knows which.

- Agents that write SQL straight over your tables, without your business rules, return numbers that look right and that nobody can explain.
- “Active customer” means 60 days to one team and 90 to another.
- When the number doesn’t match the usual report, the debate becomes opinion against opinion.
- With regulated data, no personal data may ever reach an agent.

THE SOLUTION

Define once.
Serve everyone.
Prove it every time.

1 **Define**

Metrics, business rules and concepts in one versioned model, with owners and tests.

2 **Serve**

Agents request the governed metric over MCP instead of guessing SQL. Apps use the REST API.

3 **Prove**

Every answer carries its evidence: formula, rules, filters, reference date and hashes.

MULTI-DATABASE, MULTI-PLATFORM

One model.
The same number on all of them.

DATA PLATFORMS

Databricks Snowflake

DATABASES

Oracle SQL Server PostgreSQL MySQL

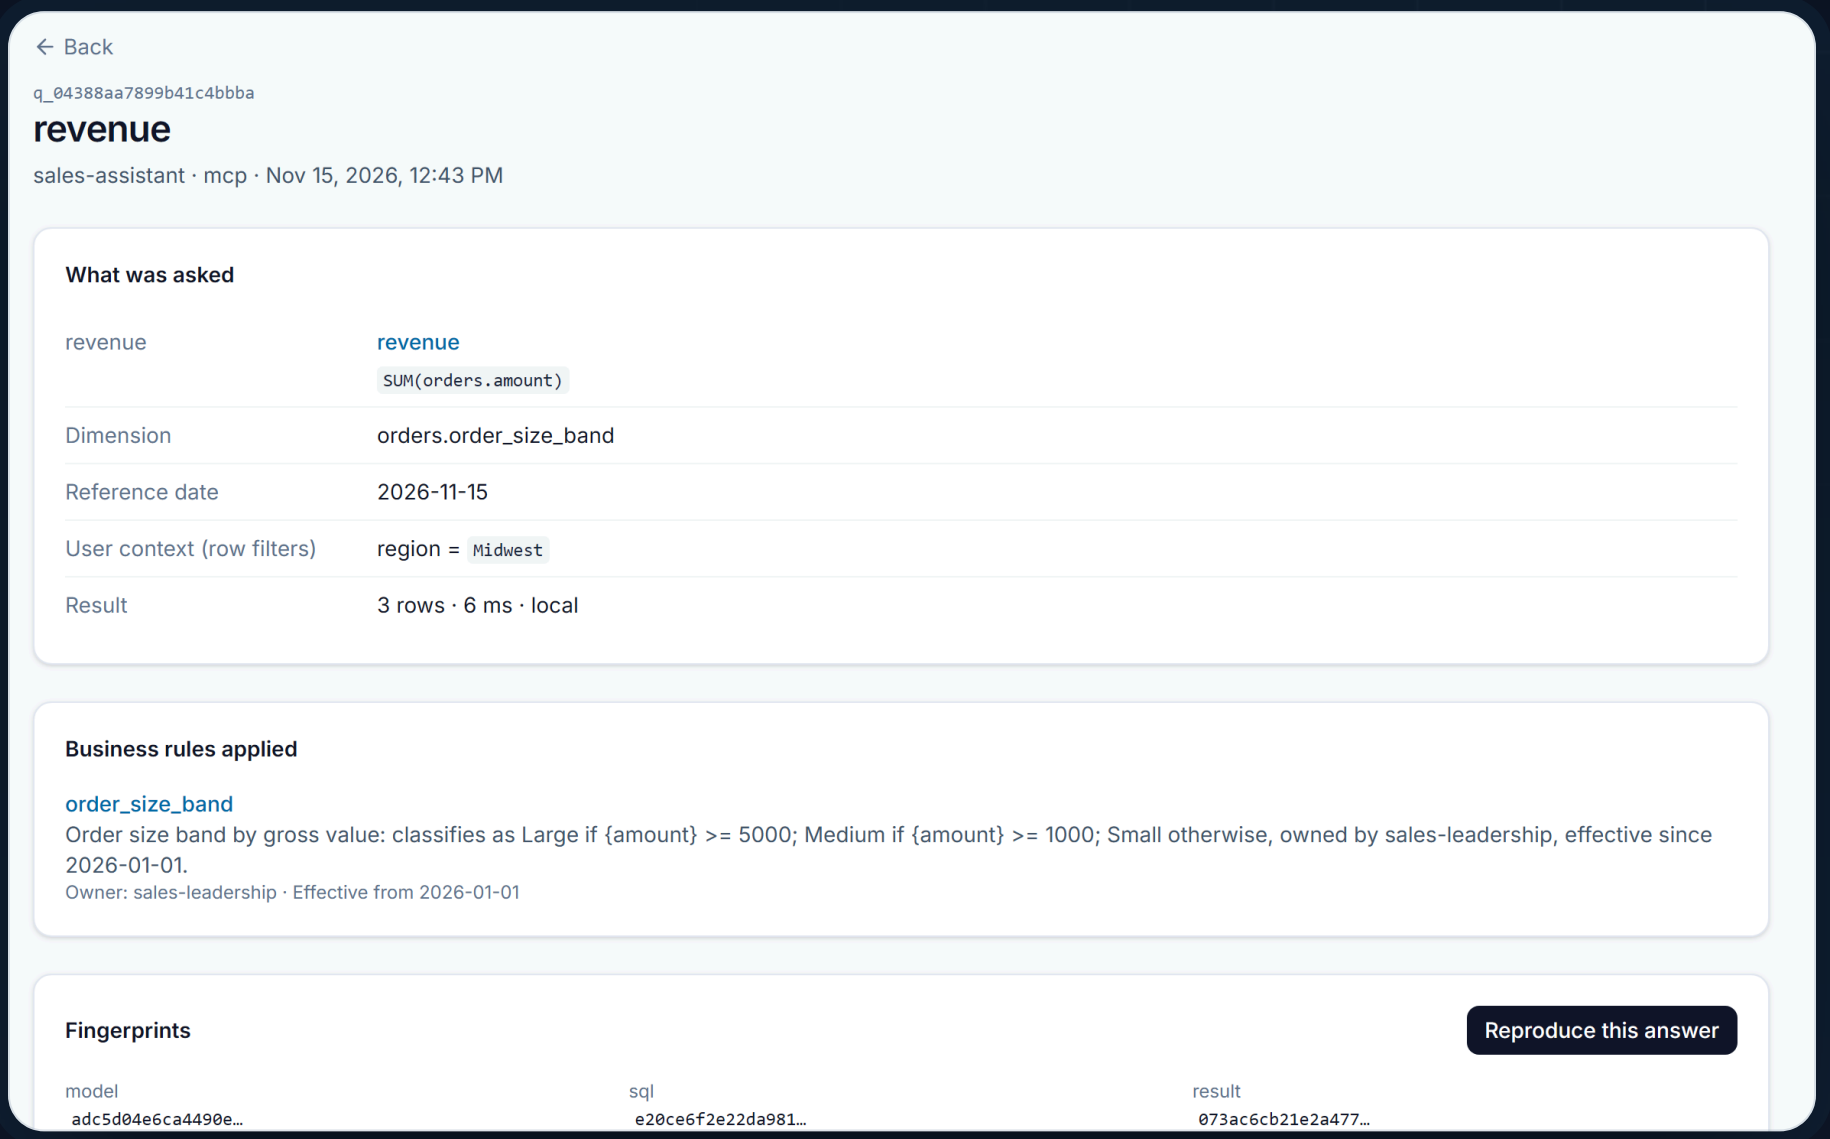
MariaDB DuckDB

The engine generates deterministic SQL for each dialect. The same equivalence suite runs on all of them:

- reference queries checked against hand-written SQL;
- a “zoo” of classic divergences;
- a randomized test that compares every engine against the others.

VERIFIABLE NUMBERS

Every answer comes with the proof.



`cosmos replay` reproduces the answer and tells you whether the number changed, and why: the model, the engine or the data.

RECONCILIATION

Doesn’t match the legacy report? Cosmos tells you why.

It compares the governed metric with the SQL your team already uses and names the likely cause of each divergence:

- Rows multiplied by a join
- Day cut in another time zone
- Integer division
- Counted rows, not distinct
- Rows lost in an inner join
- Different filter
- Blank group dropped
- Rounding

> The conversation moves from “who is right” to “which rule do we want”.

ONTOLOGY AT THE CORE

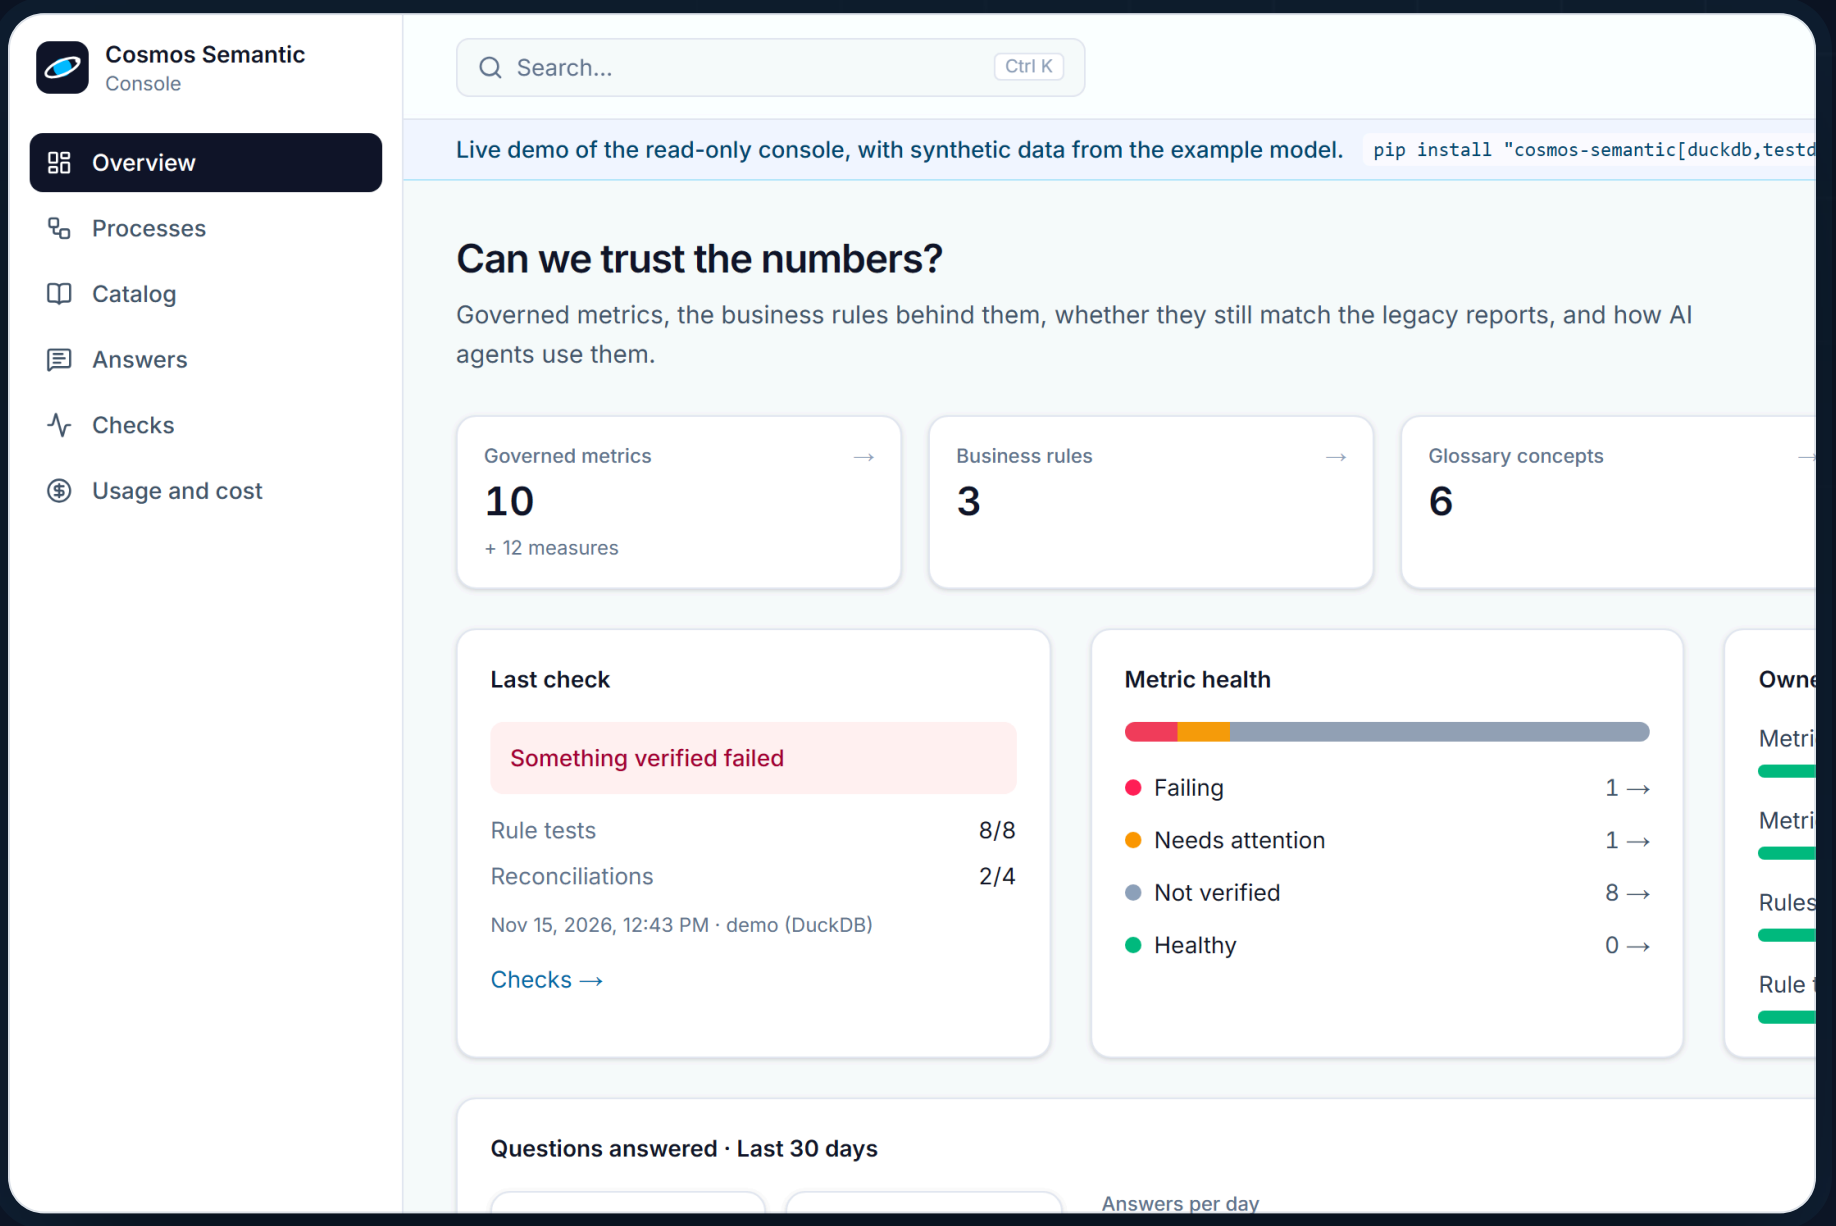
Business context, not just SQL.

- A bilingual glossary (SKOS) with synonyms: the agent knows “sales” and “revenue” are the same concept.
- Versioned rules with owners and tests; past numbers under the rule in force at the time.
- A graph of teams, systems and processes: who is affected by each change.
- Per-user row filters, and personal data blocked on the agent channel.



CONSOLE

A console
for the people who decide.



- Health and owner of every metric, processes as a graph, and every answer audited.
- Usage and cost per agent. In English and Portuguese. Live demo, no install.

Cosmos Semantic

Open.
No lock-in.

- Apache 2.0, with its own engine.
- Imports models from dbt, Power BI, Cube and Apache Ossie.
- MCP server listed in the official MCP Registry.

```
$ pip install cosmos-semantic
```

Live demo
wiliamrosa.github.io/cosmos-semantic-platform

Code, and issues to contribute
github.com/WiliamRosa/cosmos-semantic-platform

Contributions welcome