

Talk-to-Data


Governance-first semantic analytics for enterprise quotation workflows

The challenge

More than 30 salespeople could depend on internal experts to answer routine quotation questions about products, discounts, and customer agreements. Manual lookup and email or chat coordination introduced an estimated **5–30 minute delay per question**.

The approach

Build a proof of concept around a curated semantic layer instead of exposing raw enterprise tables to an LLM.



Architecture decisions

Decision	Why
Snowflake + Cortex	Fast POC path with data, AI, and governance in one platform.
dbt layers	Reproducible SQL, lineage, testing, and versioned business logic.
Semantic views	One business contract for dashboards and natural-language queries.
Daily refresh	Sufficient for quotation research; avoids premature real-time complexity.
Guardrails in data layer	Redaction, access control, freshness checks, and exception escalation are enforceable—not only prompt instructions.

Estimated outcome

Signal	Estimate
Routine answer latency	5–30 min → <1 min
Sales users in scope	30+
Data freshness	Daily batch
Main benefit	Less coordination overhead and faster quote responses.

These are **estimates based on the human workflow**, not measured production KPIs. The next delivery step is a controlled pilot that captures actual latency, adoption, escalation rate, and answer quality.

Delivery lesson

The reusable asset was not just the chatbot. It was the delivery pattern: establish the semantic contract, encode business rules, apply governance, and then expose the smallest useful interface.

Stack: SAP source system · Snowflake · dbt · Snowflake semantic views · Cortex Analytics · internal chatbot workspace

All examples are anonymized and contain no proprietary source data.