

CLIENT TECHNICAL ASSESSMENT

LLMQuery Technical Maturity Assessment


Architectural depth beyond a basic LLM integration

Overall assessment

Overall assessment: an advanced specialist text-to-SQL application, with substantial engineering beyond a basic LLM integration. LLMQuery is a multi-layer AI and retrieval architecture - not an LLM wrapped around a database. It combines advanced retrieval, query decomposition, contextual reasoning, data-layer techniques, business definitions, controlled execution and local diagnostics to produce answers with a level of precision and reliability that approaches research-grade question-answering systems.

Technical maturity scale

On a four-stage scale from a basic SQL-generation assistant to an independently validated leading platform, LLMQuery belongs in stage 3: engineered specialist application.



Implemented engineering capabilities

Capability	What is implemented
Database preparation and schema grounding	Local profiling, schema inspection, relationship analysis and contextual notes help the planner work with abbreviated names, coded fields and incomplete database documentation.
Business meaning beyond column names	Firm-specific rules, reviewed definitions, clarification prompts and visible assumptions help distinguish a technically valid query from the firm's intended question.
Controlled execution	Structured model outputs feed a local execution layer with SQL guards, dialect-specific restrictions, timeouts, row limits and read-only session settings where supported.
Diagnostics beyond syntax	Local checks can flag join multiplication, unmatched filter values, empty-result conditions and changes in how repeated questions are interpreted.

Capability	What is implemented
Human-verified reuse and evaluation	Reviewed SQL can be reused for established questions, while feedback and ground-truth evaluation tools support repeatability and ongoing testing.
Broad database compatibility	Support spans SQL Server/Azure SQL, PostgreSQL, Oracle, MySQL/MariaDB, SQLite and Microsoft Access - relevant to firms with mixed or legacy systems.

Alignment with industry direction

These capabilities align with important directions in commercial text-to-SQL: richer schema context, explicit business semantics, verified query knowledge and evaluation. Google Cloud's AlloyDB approach and Snowflake's Cortex Analyst both emphasize context and semantic guidance beyond a natural-language prompt. The overlap supports LLMQuery's architectural credibility; it does not establish equivalent performance or platform breadth.

Industry references [Google Cloud AlloyDB](#) | [Snowflake Cortex Analyst](#)

Privacy boundary and local data advantage

LLMQuery's privacy-constrained architecture keeps SQL execution, retrieved database values and result composition local. The cloud planner receives the schema, the user's question and contextual metadata needed to plan the query, but it does not receive the returned result rows. This separation is a meaningful advantage for organizations that want LLM-assisted database access without transmitting operational query results to the model provider. Some notes and rules may contain coded values, so those inputs should still be governed as transmitted context.

Operational data stays local	Planner receives context, not results
Queries execute locally. Retrieved database values, result rows and result composition remain outside the cloud planning step.	Schema, questions and contextual metadata are transmitted; notes and rules may include coded values. Returned result rows are not sent.

Client conclusion

The defensible client conclusion is that LLMQuery is an advanced, purpose-built database application with layered accuracy, privacy, diagnostic and execution controls - not simply a chat interface attached to an LLM. Its local execution and result-isolation model is a meaningful differentiator for firms that need natural-language access to operational systems while limiting what reaches the cloud planner. LLMQuery's stage 3 rating reflects substantial implemented technical depth.

Assessment boundary

This assessment is grounded in documented implementation and architectural maturity. Industry certifications, market percentiles and independently verified accuracy rankings are separate forms of evidence and are not represented by the stage 3 designation.

Basis

Basis: LLMQuery 1.29.0 technical brief dated September 13, 2026, and public industry documentation reviewed September 14, 2026. This is an informed architectural assessment, not an independent code audit or competitive benchmark.