

Luis B. Mata Unveils SQLv2: The Open Standard for AI-Native Databases

SQLv2 unifies SQL, machine learning, vector search, and generative AI in a single query language for the AI era.

Charlotte, NC, September 25, 2025: The modern data stack is broken. AI teams juggle five or more disconnected systems. PostgreSQL or MySQL for transactions. Pinecone for vectors. Snowflake for analytics. External APIs for generation. This fragmentation drives complexity, latency, and rising costs.

Technology leader **Luis B. Mata** today announced **SQLv2**, an **open standard** that merges AI and traditional SQL in one language. SQLv2 lets developers query structured data, embeddings, and AI models directly inside the database engine. No data movement between services.

“Teams spend up to 70% of their time building plumbing between databases, vector stores, and AI services,” said Mata, creator of SQLv2. “With SQLv2, work that once required multiple systems and hundreds of lines of Python becomes a single SQL query.”

Example: From Complex Pipelines to One Query

```
SELECT
  customer_id,
  PREDICT('churn_model', customer_features) AS churn_risk,
  GENERATE_TEXT('Create retention offer for', segment) AS personalized_offer
FROM customers
WHERE embedding <=> EMBED('high-value customer behavior') > 0.85
  AND last_purchase < CURRENT_DATE - INTERVAL '30 days';
```

Before SQLv2: separate systems and complex ETL pipelines. With SQLv2: one query, one engine.

Four Technical Breakthroughs

1. **Native ML inference.** Models run inside the database process with predictable latency.
2. **First-class vector search.** Built-in indexing and similarity operators replace external vector databases.

3. **Generative AI integration.** Text generation, summarization, and classification inline with SQL queries.
4. **Multimedia data types.** Query and analyze images, video, audio, and documents alongside structured data.

Why Now Four developments make SQLv2 viable today.

- **In-engine inference is production-ready.** No external AI services required.
- **Hardware acceleration is ubiquitous.** GPUs and modern CPUs speed in-database inference and embedding generation.
- **Vector indexing and storage matured.** HNSW, IVF, product quantization, columnar storage, and zero-copy execution enable real-time search and scoring in one engine.
- **Security and compliance demand data-in-place.** Keeping inference inside the database preserves the trust boundary and simplifies governance.

First Implementation: SynapCores has built the first production-grade database implementing SQLv2. The system unifies transactional, vector, and AI workloads in one engine.

“By unifying workloads in one engine, we expect major gains in performance and development velocity,” said Mata. “One architecture replaces an entire zoo of databases and services.”

SQLv2 maintains **backward compatibility with ANSI SQL:2016**. Existing queries run as is. New AI features arrive through natural extensions.

Target Use Cases

- E-commerce: real-time personalization with unified customer data
- Financial services: fraud detection with transactions and embeddings
- Biopharma: drug discovery candidate selection and ADMET triage
- Healthcare providers: patient similarity search and clinical NLP
- Media: recommendations using text, images, and video
- Defense: real-time threat analysis and detection

Open Standard for the AI Era SQLv2 is released under the **Creative Commons Attribution 4.0** license. The specification includes full technical documentation, AI-native data types, a comprehensive function reference, and migration guides for PostgreSQL, MySQL, and MongoDB users.

“SQL democratized data access in the 1970s,” said Mata. “SQLv2 aims to democratize AI-powered applications. This is not about one company owning the future. This is a standard the industry can use and improve.”

Get Involved

- Read the specification: <https://github.com/synapcores/sqlv2>
- Join the community: <https://synapcores.com/sqlv2>
- Sign up for the SynapCores beta: <https://synapcores.com/early-access-beta>

About Luis B. Mata Luis B. Mata is a serial entrepreneur and technology leader who has built large-scale data systems for startups and enterprises. His work focuses on hard problems at the intersection of data, AI, and distributed infrastructure.