

SQLv2 Fact Sheet

What is SQLv2? SQLv2 is an **open standard for AI-native databases**. It extends ANSI SQL with in-engine machine learning inference, vector search, and generative functions. With SQLv2, developers can query structured data, embeddings, and AI models directly in one database — no external pipelines or APIs.

Key Features

- **Native ML Inference** — Run predictive models directly in the database engine.
 - **First-Class Vector Search** — Similarity search and vector indexes built into SQL.
 - **Generative Functions** — Inline functions like `GENERATE_TEXT`, `SUMMARIZE`, and `CLASSIFY`.
 - **Multimodal Data Types** — Store and query images, audio, video, and documents alongside rows and columns.
 - **Zero-Copy Execution** — Minimize latency by keeping data and models in place.
 - **GPU & Modern CPU Acceleration** — Designed to leverage tensor math and vectorized instructions.
-

Why SQLv2 Now?

- **In-engine inference is production-ready** — models run within the database process with predictable latency.
 - **Hardware acceleration is ubiquitous** — GPUs and advanced CPUs enable fast vector and tensor operations.
 - **Mature vector indexing** — HNSW, IVF, and quantization make large-scale vector search practical.
 - **Compliance & security demands** — inference inside the database keeps sensitive data within one trust boundary.
-

Compatibility

- Backward-compatible with **ANSI SQL:2016** — existing queries run unchanged.
 - Migration paths for **PostgreSQL, MySQL, and MongoDB** users.
 - Functions and types defined as **natural SQL extensions**.
-

Example Query

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

One query replaces multiple systems: prediction, generation, and vector search in a single engine.

Use Cases

- **E-commerce** — Real-time personalization and product recommendations.
 - **Healthcare** — Patient similarity search and clinical NLP.
 - **Financial Services** — Fraud detection blending transactions, embeddings, and models.
 - **Media & Entertainment** — Content recommendations using multimodal data.
 - **Biopharma / R&D** — Candidate selection and discovery acceleration.
 - **Defense / Security** — Real-time threat analysis and anomaly detection.
-

First Implementation

- **SynapCores** is the first production-grade database implementing SQLv2.
 - Released under **Creative Commons Attribution 4.0** — free to review, implement, and extend.
-

Quick Links

- Specification: github.com/synapcores/sqlv2
- Explainer: synapcores.com/sqlv2
- Beta Signup: synapcores.com/early-access-beta
- Press Kit: synapcores.com/press

