

**AT A GLANCE**

The vector database that enables the building of Generative AI applications with contextual search at scale

**KDB.AI ADVANTAGE**

- **Multi-modal search** with unlimited context windows, no slow down.
- **Time aware** - detect trends, identify change, and compare moments in time
- **Filterable** to improve search performance and accuracy
- **Efficient** indexing & fast ingestion
- **More data compression** and reduced memory usage, without GPU
- **Scale to billion** vector search across enterprise data
- **Enterprise scale**, secure & real-time processing

# The Smarter Database for AI

As the world's volume and variety of data increases exponentially, businesses must store, analyze, and activate data in a smarter way.

Vector databases enable the processing of data through intelligent similarity search and by extending the capabilities of large language models (LLMs), allowing them to search and reason about data in real-time, with scalability.

KDB.AI is the world's leading multi-modal vector database, according to [DB-Engines](#). It helps AI engineers and data scientists develop AI capabilities and applications with flexibility, working with common frameworks, and integrating with popular LLMs for public and private enterprise data.

Unique among vector databases, KDB.AI enables developers to bring temporal and semantic relevancy to their AI-powered applications. Built to handle high-speed, time-based, multi-modal query data processing, business users can query real-time data using natural language search with semantic relevance.

## How Industries Use KDB.AI

By leveraging generative AI with unstructured data and similarity search, businesses are making use of this data for more impactful decision making:



**Financial Services:** Trade trillions per day while managing risk intelligently and providing efficient market access to clients around the world.




**Telecommunications:** Optimize capacity and coverage, deliver control at the edge, and maintain critical services with self-organizing networks.



**Advanced Manufacturing:** Quickly spot real-time predictive maintenance anomalies for smarter, more efficient plants.



**Aerospace and Defense:** Analysis of operational data for correlation of intelligence, improving command decision making.



**Healthcare:** Model disease and interventions in virtual clinical trials to predict patient responses and required treatment dosages.

# ➤ The Smarter Database for AI

## KDB.AI EDITIONS


- **KDB.AI Cloud** (Starter Edition) is a free-to-use, managed cloud-based instance designed for developers: 4GB memory per instance; 30GB data storage; get started quickly with sample projects.
- **KDB.AI Server** (Standard Edition) is a self-managed server-client instance with single container deployment; scale to your requirements; customize to your dev environment.
- **Azure Marketplace Offering:** KDB.AI can be integrated with Azure ML and OpenAI, with KDB.AI samples already loaded to start building AI-powered apps.

## How KDB.AI Supports Multi-Modal RAG Use Case

Multi-Modal RAG is an advanced AI technique that combines the capabilities of Large Language Models with the ability to retrieve and utilize information from various data sources (Audio, Text, Video, etc.)

This approach is particularly beneficial for businesses as it allows AI systems to provide more accurate, detailed and contextually relevant responses from public and private datasets.

KDB.AI fully supports multi-modal embeddings, including ImageBind, Clip, visualBERT, and CLAP to enable enterprise customers with seamless Hybrid Search capability and wholistic responses to end user queries.



KDB.AI allows you to set-up a knowledge-based vector database and search engine in a few simple steps.

With KDB.AI you can:

- Create an index of vectors (Flat, IVF, IVFPQ, or HNSW)
- Append vectors to an index
- Perform fast vector similarity search with optional metadata filtering
- Persist an index to disk
- Load an index from disk
- Integrate with preferred Gen AI tools such as:



LangChain



LlamaIndex



OpenAI

KDB.AI is the only vector database with temporal similarity search on both transformed and non-transformed data.

Transformed temporal similarity search will reduce time series windows by about 99% whilst maintaining the original data's shape, resulting in incredibly fast vector search capability.

In contrast non-transformed temporal similarity search negates the need to embed, extract and store vectors, resulting in near real-time similarity search with extreme memory efficiency.

## Get started with KDB.AI

For more information please contact [sales@kx.com](mailto:sales@kx.com) or visit [www.kx.com](http://www.kx.com)