

RavenDB

The Boring Database

The high-performance multi-model database built for real-time access, high throughput, and seamless scalability.

Simplify your database management with automatic indexing, intelligent query optimization, and distributed clustering, ensuring low-latency performance, high availability, and near-zero downtime across on-premise, cloud, and hybrid environments.

RavenDB 7.0 adds Vector Search with AI, integrating with OpenAI, Azure OpenAI, Ollama, and Google AI for advanced semantic search and Retrieval-Augmented Generation (RAG) for LLMs.

Faster Queries, Lower Costs – RavenDB Thinks for You.

Capabilities



Automatic Indexes



Secured by Default



Fully ACID



Vector Search



Full-Text Search



JSON



Continuous Availability



Aggregate in Less Than 1ms



V7.0 Now with Native AI-Integration

INSTALL
RavenDB
ON PREMISE

10
min setup

RavenDB
CLOUD

2
min setup

DATABASE AS A SERVICE



Windows



Linux



Raspberry
Pi



Mac OS



Docker



Amazon
Web Service



Microsoft
Azure



Google Cloud
Platform



HIPAA
Compliant

Trusted by global industry leaders



TOYOTA



eso

trackmatic

verizon



Configit

Rakuten kobo



Luware

Medicaid



The best developer experience I've ever encountered; everyone else should take notes.

Software Engineer in the service industry

Gartner. ★★★★★



Visit us at
RAVENDB.NET