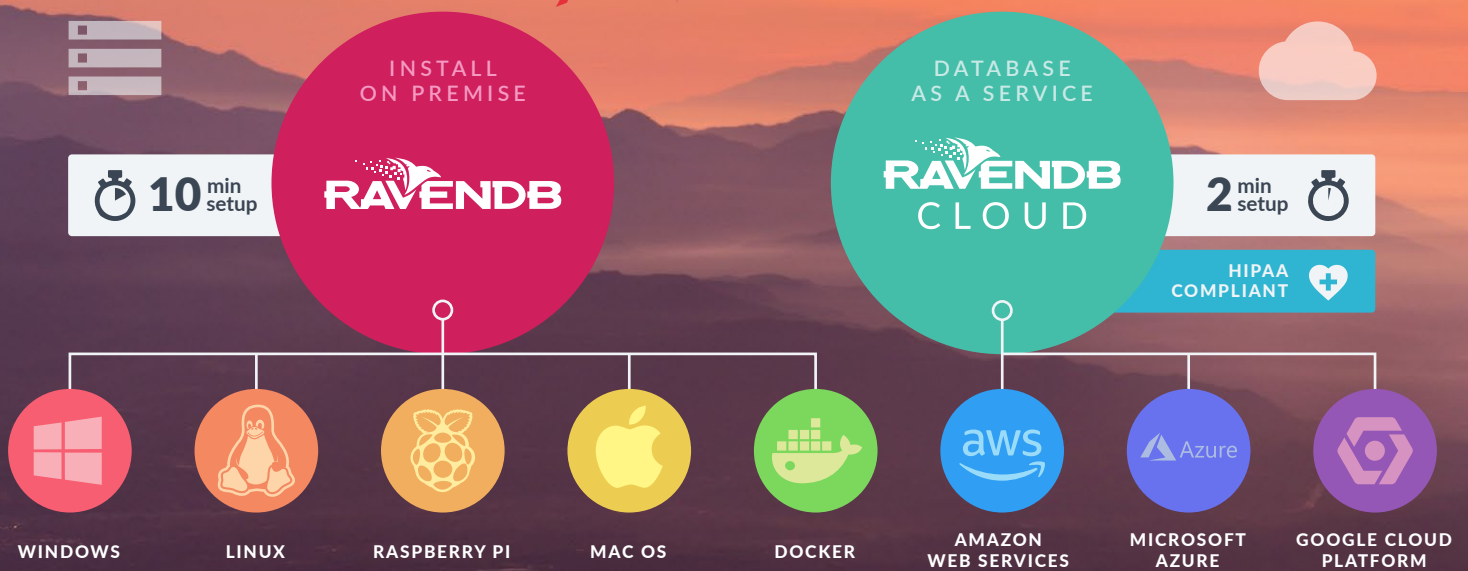




# Where does RavenDB excel?

RavenDB meets the needs of today's applications by addressing their primary areas of concern and offering features unique to our solution.

## Versatile, Quick and Easy to Setup



## Everything You Need is Already Inside



### FAST PERFORMANCE

- Transactional Merger
- Map Reduce
- Auto Indexes
- Built-in Caching

### EASY TO USE



- GUI
- API & Docs
- Dev Tech Support
- Full Text Search

- Resource Efficient
- Auto Query Optimiser
- Dynamic Querying
- Document Compression

- ACID
- Change Tracking
- Multi Master
- Built for Distributed Applications



### COST SAVINGS

- Throttling Handling For Burstable Instances

### DATA INTEGRITY



## About RavenDB

Mentioned in both Gartner and Forrester research, RavenDB is a pioneer in NoSQL database technology with over 2 million downloads and thousands of customers from Startups to Fortune 100 Large Enterprises.

Over 1,000 businesses use RavenDB for IoT, Big Data, Microservices Architecture, fast performance, a distributed data network, and everything you need to support a modern application stack for today's user.



### Are you interested in setting up a Proof of Concept?

We offer 2 FREE hours of POC support, along with reference materials to help you command the best tools for your next project.

Contact us at [info@ravendb.net](mailto:info@ravendb.net)

Take a free instance of RavenDB Cloud at

[cloud.ravendb.net](https://cloud.ravendb.net)

Or take a free RavenDB On Premise Download at

[ravendb.net](https://ravendb.net)

# A Cost Cutter

## Manhours

Intuitive GUI, human readable RQL and straightforward API reduces development time



## Compute Cost

Automatic Indexes, Dynamic Queries, Caching, Documents Compression, and Optimization Suggestions lower your cloud costs.

## Consistent Requirements

Small resource demands remove the need for frequent upgrades. New RavenDB versions strive to remain lightweight.

## Support Cost

RavenDB resolves many issues automatically or gives you tools to easily diagnose the problem yourself

Over 95% of RavenDB clients stick with RavenDB. There's no need to change databases.

## The Juicy Details

### ACID

RavenDB is fully transactional over multiple documents, providing ACID guarantees throughout your database cluster.

The first document database to offer ACID data integrity, RavenDB's dazzling performance is accompanied by the highest level of data integrity.

### Full Text Search

Perform advanced full-text search operations over one or more fields at once, with the ability to supply your own custom analyzer.

### Transactional Merger

Before sending your data from the REST endpoint to storage, RavenDB will bundle multiple transactions into a single transfer. Minimizing the number of round-trips between the client and data store reduces latency and boosts performance. On the cloud, the lower number of I/O requests is translated to lower bills.

In rare scenarios of transaction failure, RavenDB will unbundle the remaining transactions and commit them one by one.

### Map Reduce

In addition to regular map indexes, RavenDB provides a pre-computed map-reduce mechanism that continuously aggregates existing indexes in the background so executed queries will use optimized indexes. Doing so allows for large-scale aggregations, keeping latency tight as your data expands.

Map Reduce is native to RavenDB, eliminating the need for third-party aggregation add-ons.

### Machine Learning Integration

Machine-learning models can be integrated into RavenDB and used by the database for your needs.

An image analysis model, for example, can be called by your code to classify pictures attached to your documents. The model can be trained offline until its accuracy is sufficient, and when employed, its results can be indexed and queried by RavenDB's native features and locate profiles by various details of their Gravatar picture.

### RavenDB Query Language (RQL)

It's the SQL of the Document Database.

80% of it is the traditional query language and the rest is human readable, quick to learn, and simple to use with great documentation to back it up.

## RavenDB's "Data Includes" is our answer to the expensive select N + 1 issue.

The select N+1 problem taxes your servers with countless queries.

A query for *open help desk cases and customers associated with them*, for example, would approach the database once for the open cases and then again for each customer related to each case.

### GUI

Easily find and instantly resolve performance bottlenecks.

Your state-of-the-art GUI monitors the internals of your database and measures its performance while it queries, indexes, aggregates, and more.

RavenDB's Management Studio enables you to perform most functions available via our API. You can run queries, create new databases and nodes, and much more.

You do not need a special connection for this; every local node has its own studio instance embedded in the server.

### Resource Efficient

Over 1 million reads and 150k writes per second on a single node on commodity hardware.

Consistently lightweight, there will be no unpleasant surprises with your next version.



Runs on multiple Operating Systems and platforms including Windows x64/x86, Linux, macOS, Docker, Raspberry Pi, and ARM Chips.

### Auto Indexes

RavenDB is the only database to automatically index your queries on the fly, saving you the trouble of coding indexes yourself (though you are also free to create indexes manually if you wish).

We slash your latency and supercharge your performance using machine learning indexes. You can even index Microsoft Office documents like Word files.

With every query you make, RavenDB learns more about your system and operational behavior. When you execute the query, RavenDB will either use an existing index or create a new one.

Indexes that haven't been used for a while will be discarded to keep your memory fresh.

### Built-in Caching

Automatic and aggressive caching are already part of your database, so you don't have to waste your developer's time coding it.

Unchanged queried data will be fetched from memory, removing the need for an additional trip to the database server when users repeat a query.

### Change Tracking

RavenDB keeps track of the changes made in documents and updates the database with all changes in a single transactional operation.

You do not have to handle change tracking, handle change management, or figure out how to best write to the database.

### Built for Distributed Applications

We do not use *locking* to manage concurrent updates to your distributed database.

RavenDB's conflict resolution mechanisms prevent document duplication and other errors when multiple users update a document concurrently, maintaining smooth service for these users.

This conduct is vital because RavenDB is a multi-master database, allowing your users to approach any node in your database cluster independently. Built for distributed applications, RavenDB handles the biggest challenges you can face in such an environment.

### Multi-Master

Permitting clients to use any cluster node to read or write, minimizes latency, maintains high-availability for the database cluster, and allows nodes to keep working even when the network is offline.

### Document Compression

Some documents contain mainly property names. When such documents are numbered by the millions, or when the database holds time series with billions of entries, storage becomes a serious concern and its cost may be considerable, especially on the cloud.

RavenDB comes with standard and adaptive-learning document compression, that learns and compacts repeated patterns within and across documents and document extensions like time series, to save over 50% of your cloud storage costs and 10% of your total cloud database costs.

### Throttling Handling For Burstable Instances

RavenDB makes sure your database is using the most essential operations as you reach your CPU credit limit on the cloud. It quickly reassigns background tasks to operative nodes if nodes that run them get throttled, and suspends certain operations as you reach your CPU credit limit, so users won't be impacted by throttling.

### Multi-Model

Do a lot of different things with your database.

RavenDB supports the Document Model, Time Series, Graph Queries, Key-Value, and Distributed Counters, to keep overhead manageable in a microservices architecture.

That's a lot of work, which on the cloud amounts to a huge computing cost.

A RavenDB *Include* request allows you to reduce this scenario to a single roundtrip by retrieving all the data you need in a single batch.

