



# RDM Workgroup Plus Edition 11

*RDM Workgroup is a high-performance, database management system optimized for the operating systems used within the desktop and server environments. It is ideal for standalone and networked applications such as those commonly found on today's desktop and laptop computers. Multiple APIs provide developers a multitude of programming options and functionality. In addition RDM Workgroup includes an ADO.NET 4.0 provider, JDBC 4 Type 4 driver and an ODBC 3.5 driver. Using these drivers managed code applications can easily be created using popular languages like C# and Java.*

## Overview

RDM's ACID compliant database engine supports B-tree and hash indexes; the B-tree indices can support simple and/or compound keys. Additionally the database engine has been developed to fully utilize multi-core processors, run within minimal memory, and support both in-memory and on-disk storage. Implemented as a linkable library the database is allowed to become an embedded part of your applications. Unique to the Plus package are the RDM replication and mirroring functionality that enables greater availability and scalability.

With nearly 30 years of development history, it is estimated that this embedded database management system has been used by more than 20,000 developers and deployed in over 20 million installations in all the major industries; including Aerospace & Defense, Automotive, Business Automation, Financial, Government, Industrial Automation, Medical, and Telecommunication.

---

## Key Benefits:

- Multi-Core Scalability
  - Distributed Architecture
  - Extraordinary Performance
  - Proven Reliability
  - Excellent Support
- 

## Performance Driven Features

- **Multi-Core Support** - Efficiently allocate transaction processing to take advantage of multi-core systems for optimal speed.
- **Multi-Versioning Concurrency Control (MVCC)** - Implement read-only transactions where a virtual snapshot of your embedded database is readable until the read-only transaction is terminated by the task, even if it is being concurrently updated. Avoid read locks to improve multi-user performance.
- **Pure and Hybrid In-Memory Database Operation** - Configure your database to run completely on-disk, completely in-memory, or a hybrid of both; combining the speed of an in-memory database and the stability of on-disk in a single system.
- **Multiple Indexing Methods** - Use B-Trees or Hash Indexes on tables. Hashing on large volumes often provides faster access to data than b-tree indexing methods. Hashing enhances speed by using buckets to store the index information.
- **Better Performance through Scaling Out** - Easily achieve true horizontal scaling across clustered or distributed systems without the need to re-write of your application.
- **True Global Queries** - Connect any application to one or more databases and query them as if it is a single instance. Perform global, locally or across a network, to multiple database instances with no regard for where the data is located.
- **Master-Slave Database Replication** - Create applications that replicate data between different databases on different systems, on the same system, in memory and on disk. Asynchronously replicate your data between different database instances, schemas, and systems. This allows developers to create a heavily indexed persistent query database while capturing data through a non-indexed in-memory database.
- **3rd Party Replication** - The automated database schema translator seamlessly transforms data from native RDM to SQL enabling effortless data flow into the enterprise. The result is the ability to combine near real-time critical operational data with other data sources to provide the complete information to make more intelligent business decisions at any level of your organization.
- **Master/Slave Mirroring** - Data redundancy without sacrificing performance! Master-Slave mirroring introduces data redundancy by automatically mirroring the application database locally or across a network.
- **Synchronous or Asynchronous Mirroring** - Designed for carrier grade systems the new mirroring engine is designed to remove any single point of a failure in your fault-tolerant system.

## Interoperability Features

- **ODBC 3.51 Driver** - Allows new client applications to be written and third party client applications like Microsoft Access, Crystal Reports, MS Excel and others to interface with RDM databases.
- **ADO.NET 4.0 Driver** - A fully managed driver written in 100% pure C# that enables developers to easily create .NET applications with RDM databases.
- **JDBC 4 Type 4 Driver** - A call-level API for SQL-based database access to RDM databases. JDBC technology allows you to use the Java programming language to exploit "Write Once, Run Anywhere" capabilities for applications that require access to data on RDM databases now matter where they are.

## Multiple APIs for Enhanced Usability

- **Navigational C API** - For well over 25 years developers have been using RDM's low-level C API of over 200 intuitive easy to use functions provides application developers with ultimate control of the database.
- **Comprehensive SQL API** - RDM's SQL is accessed internally through an easy-to-use API designed by Raima. This non-standard API is simpler than ODBC. In fact, our ODBC API is based on this one.
- **Standards Based ODBC API** - Following the ODBC standards Raima developed the ODBC API to provide developers with a familiar way to utilize the power of the RDM database engine.
- **Object Oriented C++ API** - Was designed with ease of use as its primary requirement while still providing developers with full access and control to both RDM's network and relational functionality.
- **Objective C API** - The Objective-C 2.0 interface to RDM is designed to augment the RDM navigation on Mac OS X and iOS platforms by creating custom objects to represent records and interface with the database. This combines the performance of RDM with the intuitiveness of an object oriented interface, making it easy to integrate with Cocoa applications.

## Database Specifications

- **Max. Databases Open Simultaneously:** No Limit
- **Maximum Records Per Database:** No Limit
- **Maximum Size of Database File:** Limited only by file system
- **Maximum Tables Per Database:** No Limit
- **Maximum Records Per Table:** No Limit
- **Maximum Keys Per Database:** No Limit
- **Max. Record Size:** 32K (excluding BLOB or VARCHAR)
- **Maximum Fields Per Table:** No Limit
- **Maximum Size of Keys:** 242Bits
- **RAM Requirements:** Minimum 50K, User configurable
- **Code Footprint:** Starting at ~270K depending on OS and database features.

## Supported Platforms:

- Windows
- Linux
- MAC OS
- Solaris
- HP UX

## Modes of Operation

- Standalone
- Client/ Server
- Application Linked
- Mixed Execution Mode

## Want to know more?

Please call us to discuss your database needs, or email us at [info@raima.com](mailto:info@raima.com). You may also visit our website for the latest news, product downloads and documentation:

[www.raima.com](http://www.raima.com)

## Try it!

Download a trial version:

[raima.com/downloads](http://raima.com/downloads)

## Industry Partners:

