

How to Choose an Open-Source Database



Rimini Street

It's time to revisit your database strategy

As digital transformation projects and other business requirements place new demands on data infrastructure, many organizations are reevaluating the database management systems (DBMSs) they have relied on for years.

Growing data volumes, real-time analytics, new applications, and other digital initiatives are all driving the need for more modern data architectures and platforms.

As a result, businesses increasingly are moving away from the traditional "legacy" database strategy, where IT teams use proprietary database software running on centralized servers in their own data centers. They are shifting to a newer generation of cloud databases that are much easier to deploy and manage and that scale automatically.

More often than not, these new databases are based on open-source software. In fact, open-source databases surpassed commercial databases in popularity in January of 2021, according to [DB-Engines](#), a website that gauges the popularity of hundreds of DBMSs.



The benefits of open source

Why are open-source databases growing in popularity? They typically are surrounded by a volunteer community and rich ecosystem of tools; development teams can customize or fine-tune the code if necessary; and they minimize the risk of “lock-in” because it’s relatively easy to move from one platform to another.

Choice... choice... choice!

There are many open-source databases from which to choose. According to Carnegie Mellon University’s “[Database of Databases](#),” more than 500 open-source databases have been developed. They include Cassandra, a NoSQL database; Redis, an in-memory database; SQLite, a SQL database that can be embedded in other programs; and HBase, a big data store that works with the Hadoop file system.

The two most popular open-source databases in the DB-Engines ranking are MySQL and PostgreSQL, both of which are available as free, downloadable “community” versions. And, like many other open-source databases, MySQL and PostgreSQL are also offered as cloud services from a variety of cloud database providers who repackage them as commercial offerings with enterprise capabilities that are not available in the basic editions. For instance, commercial versions of PostgreSQL, which is developed by the PostgreSQL Global Development Group, are available from AWS, Cockroach Labs, EDB, Google Cloud, Microsoft, Yugabyte, and others.



The emerging generation of open, cloud database systems are sometimes referred to as “cloud native” because, unlike traditional databases, they are designed from the start for deployment in the cloud. Some have been engineered to separate the database’s underlying compute and storage resources, which results in greater flexibility and control over performance and scalability.

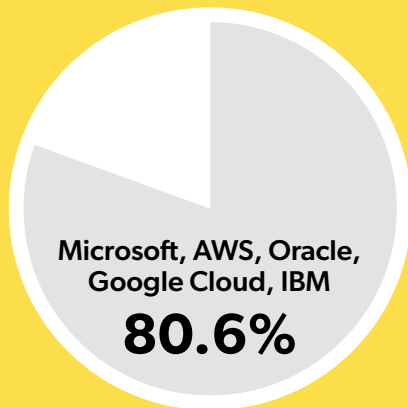
How to choose an open-source database

The market for commercial, proprietary databases has long been dominated by a few major players. The charts below show that, as per a Gartner® Market Share report,¹ Microsoft, AWS, Oracle, Google Cloud, and IBM accounted for 80.6% of the Database Management Systems Market Share revenue in 2021. The report also estimates that the DMS market share revenue grew 22.3% in 2021 to nearly \$80 billion.


But the competitive landscape is changing as startups and established vendors alike introduce new cloud database systems, many of which incorporate open-source software.

For that reason, developers and technology teams need to consider not just the underlying technology, but also the type of license associated with a particular database. Licenses determine an organization's ability to modify, use, and distribute a database, including any restrictions.

2021 Commercial Database Market Share and Growth



Market Share by Revenue¹



Market Growth by Revenue¹

Charts and graphics created by Rimini Street based on Gartner research.

With so many open-source databases and different licenses, how does an organization choose the best one? Here are a few key considerations.

What is the use case?

Open-source databases are more popular than commercial databases for certain kinds of workloads, such as spatial and time-series data, according to DB-Engines. So the data characteristics will help drive the decision.

How does it fit in the data architecture?

Capabilities such as data distribution, replication, integration with existing tools and platforms, and support for hybrid clouds will determine if a database is a good fit with an organization's IT infrastructure.

Which enterprise capabilities are required?

Commercial open-source databases may come with advanced security and governance features, 24x7 support, and other must-have capabilities.

What about database migration?

It is important to assess how easy, or difficult, it would be to import data into the open-source database, or to export data to another platform in the future if necessary.

Selecting an open-source database for ERP applications

The database choice will be largely determined by the data types and application workloads to be supported. This is especially true for enterprise resource planning (ERP) applications, which are typically optimized for use with proprietary databases from the same vendor.

There may be alternative architectures to these single-vendor application and database environments. For example, IT teams may be able to use extract, transform, and load (ETL) tools to move data from ERP applications to a data lake or data warehouse that is based on a third-party database.

In fact, ERP vendors may provide some level of support for these multi-vendor environments. For example, SAP provides data connectors for IBM, Microsoft, and Oracle databases to its Business Warehouse. And SAP's Lumira data-visualization software can connect to MongoDB's open database via a JDBC connector.

However, as these examples illustrate, there is often some degree of integration work involved to get ERP applications to work with heterogenous databases.





Considerations for moving to an open-source database

The database market is almost evenly split between commercial databases and open-source databases, but increasingly it is tilting in the direction of open source.

Of the 395 databases tracked by DB-Engines, 206 are available with open-source licenses compared to 189 with commercial licenses.

For developers and IT teams considering open-source databases, three factors guide the decision-making process:



Type of open-source license. There are more than a dozen different open source licenses, including Apache 2.0, BSD, GPL, and Server Side Public License. Read the fine print on the terms and conditions that determine usage.



In-house technical expertise. The Structured Query Language (SQL) is the “lingua franca” for relational databases, but organizations may face a learning curve with NoSQL or specialized databases, such as those for managing vectors, graphs, or spatial data. Developer hubs such as GitHub and Stack Overflow may be helpful in coming up to speed.



Deployment options. There are a range of implementation scenarios, from downloading the database software to self-managed servers, to cloud database services that are fully managed by the cloud service provider. Also, if data governance requirements in different geographies are a potential issue, the cloud service provider’s availability zones must be evaluated.

Consider getting started with open-source databases in a ‘sand box’ environment and build from there. Many cloud vendors offer their databases on a trial basis for this kind of early test and development.

Potential pitfalls to avoid

Given the potential benefits of open-source databases, what are the risks?

The security of open-source software is one consideration. For example, Log4j, an open-source program of the Apache Software Foundation, has been hit by security vulnerabilities that have forced many organizations to urgently patch affected systems.

Yet, many experts seem to agree that open source is not inherently less secure than commercial software. In fact, the open-source community is often able to respond quickly to vulnerabilities, which could be an advantage.

A bigger challenge may be the do-it-yourself (DIY) nature of open-source IT projects. Skills training, software development, database administration, and other project-related tasks require time and resources. These will all need to be addressed as managed cloud services grow in popularity and increase the opportunity for DIY solutions. Consider a [third-party support partner](#) for outsourcing or augmenting in-house resources.

The key to success with open-source databases is to select a platform that is well-matched to an organization's business needs, with strong support from the open-source community and, if desired, which is available as a cloud service. Careful planning and evaluation can go a long way in making that happen.

Additional Resources

- [Open-Source Database Market Roundup](#)
- [9 Open-Source Database Migration Best Practices](#)
- [Licensees Reducing Oracle Database Costs and Footprint with Open-Source and Cloud Alternatives](#)
- [Rimini Street Support for Open-Source Databases](#)

¹ Gartner, Market Share: All Software Markets, Worldwide, 2021, 12 April 2022.

Rimini Street®

riministreet.com

info@riministreet.com

twitter.com/riministreet

linkedin.com/company/rimini-street

About Rimini Street

Rimini Street, Inc. (Nasdaq: RMNI) is a global provider of enterprise software products and services, the leading third-party support provider for Oracle and SAP software products and a Salesforce® partner. The company offers premium, ultra-responsive and integrated application management and support services that enable enterprise software licensees to save significant costs, free up resources for innovation and achieve better business outcomes. Global Fortune 500, midmarket, public sector and other organizations from a broad range of industries rely on Rimini Street as their trusted enterprise software products and services provider.

© 2022 Rimini Street, Inc. All rights reserved. "Rimini Street" is a registered trademark of Rimini Street, Inc. in the United States and other countries, and Rimini Street, the Rimini Street logo, and combinations thereof, and other marks marked by TM are trademarks of Rimini Street, Inc. All other trademarks remain the property of their respective owners, and unless otherwise specified, Rimini Street claims no affiliation, endorsement, or association with any such trademark holder or other. This document was created by Rimini Street, Inc. ("Rimini Street") and is not sponsored by, endorsed by, or affiliated with Oracle Corporation, SAP SE or any other party. Except as otherwise expressly provided in writing, Rimini Street assumes no liability whatsoever and disclaims any express, implied or statutory warranty relating to the information presented, including, without limitation, any implied warranty of merchantability or fitness for a particular purpose. Rimini Street shall not be liable for any direct, indirect, consequential, punitive, special, or incidental damages arising out of the use or inability to use the information. Rimini Street makes no representations or warranties with respect to the accuracy or completeness of the information provided by third parties, and reserves the right to make changes to the information, services or products, at any time. LR0008013 | US-092722