

June 14, 2006

# The Forrester Wave™: Open Source Databases, Q2 2006

by Noel Yuhanna

TECH CHOICES

FORRESTER®

Helping Business Thrive On Technology Change



June 14, 2006

## The Forrester Wave™: Open Source Databases, Q2 2006

Ingres, PostgreSQL, And MySQL Lead In Our Product Evaluation

by Noel Yuhanna

with Mike Gilpin and Megan Daniels

### EXECUTIVE SUMMARY

Many enterprises are turning to open source databases to reduce database management cost and avoid vendor lock-in. The maturity level of open source databases is at its highest level ever, with more choices, better support, and comprehensive ecosystems. To assess the state of the open source database market and see how the projects stack up against each other, Forrester evaluated the strengths and weaknesses of six leading open source database projects across approximately 90 criteria. The result: Ingres, MySQL, and PostgreSQL are the Leaders, while Derby, Firebird, and Oracle are Strong Performers. Ingres, Oracle, and PostgreSQL offer strong support for transactional processing, while Oracle and MySQL offer strong support for embedded database platforms. For data warehouses, none of the projects offer strong native data warehouse-related features, but some third-party vendors help fill the gap with their extended solution offering for open source databases.

### TABLE OF CONTENTS

#### 2 Open Source Databases Continue To Gain Momentum

The Open Source Database Market Landscape

#### 3 Open Source Databases Evaluation Overview

Evaluation Criteria Include Capabilities, Current Strength, And Future Plans

Vendors Selected For Enterprise-Class DBMS Features And Production Deployments

#### 5 Open Source Databases Offer Sufficient Features For Most Applications

#### 8 Vendor Profiles

Leaders

Strong Performers

#### 11 Supplemental Material

### NOTES & RESOURCES

Forrester conducted open source database evaluations in Q3 2005 and interviewed 18 vendors, user companies, and projects including: Derby, Firebird, FUJITSU, IBM, Ingres, MySQL, Oracle, Pervasive Software, PostgreSQL, and Sun Microsystems.

#### Related Research Documents

[“Trends 2006: Open Source Databases”](#)  
January 3, 2006, Trends

[“The Future Of DBMS Technology”](#)  
September 29, 2005, Trends

[“Open Source Databases Come Of Age”](#)  
December 29, 2004, Tech Choices

## TARGET AUDIENCE

Application development executive, chief information officer, information management executive

## OPEN SOURCE DATABASES CONTINUE TO GAIN MOMENTUM

Open source databases have come a long way in delivering strong DBMS features and reliability, and building a community that continues to grow at a fast pace. In 2005, the industry witnessed great momentum around open source databases — from product enhancements, improved customer support, and large vendors joining the bandwagon. The key factors for increased adoption of open source are:

- **The need for only basic DBMS features.** Most applications only require basic data management functionality that typically is met by most open source database projects. Forrester finds that many applications do not need thousands of DBMS features but are satisfied by a more basic list of essential data repository features.
- **The desire to avoid vendor lock-in.** Many enterprises prefer to go with open source databases because it avoids vendor lock-in. Unlike closed source databases, open source allows enterprises to contribute to the code and shape its road map.
- **The need to lower data management costs.** All enterprises are increasing the number of databases to support business applications. Open source databases help lower data management costs, largely because of lower acquisition and support costs.
- **The desire for greater ease of use.** Over the past decade, closed source databases have increased complexity by several orders of magnitude, largely because of a plethora of features. Open source databases are usually a lot easier to use, largely because they support fewer features and offer simplified user administration interfaces.

## The Open Source Database Market Landscape

Forrester estimates the current open source database market to be \$400 million, which includes support, services, and licenses, and forecasts that it will reach \$1 billion by the end of 2008. The open source database market is represented by the following categories:

- **Online transactional processing (OLTP) databases.** This category primarily focuses on applications such as enterprise relationship management (ERP), customer relationship management (CRM), supply chain management (SCM), and custom applications that require databases to be of high concurrency, performance, and scalability and to be secure. Forrester believes that transactional databases constitute 85% of the total open source database market.

- **Embedded databases.** In this category, such databases typically form an integral part of the application and/or device, often require no administration, and usually have a small memory footprint. These are often sold through value-added resellers (VARs), independent software vendors (ISVs), and OEM channels. The leading market tends to be the telecom industry, but the market continues to expand in other industries as well, including retail, technology, manufacturing, and automotive.
- **Data warehouses.** Open source databases are also being used to support data marts and data warehouses. Open source databases still lag in data warehousing features when compared to closed source databases, in areas including bitmap indexes, partitioning, and large database support. However, vendors such as Bizgres Group and Ingres continue to focus on expanding data warehouses features.

## OPEN SOURCE DATABASES EVALUATION OVERVIEW

To assess the state of the open source databases market and see how the vendors stack up against each other, Forrester evaluated the strengths and weaknesses of top open source database projects.

### Evaluation Criteria Include Capabilities, Current Strength, And Future Plans

After examining past research, user need assessments, and vendor and expert interviews, we developed a comprehensive set of evaluation criteria (see Figure 1). We evaluated vendors against approximately 90 criteria, which we grouped into three high-level buckets:

- **Current offering.** To assess projects' strengths, we evaluated each offering against 10 key groups of criteria: transaction processing, data warehousing, embedded database, programmability, data types and interfaces, database availability, security, tools, platforms, and administration. The transaction processing, data warehousing, and embedded database criteria help in establishing which project has the best support for such types of database implementations. The criteria around programmability, data types and interfaces, and administration focus on projects' interoperability, flexibility, and integration capabilities. And finally, the database availability, security, and platforms criteria focus on support for mission-critical, enterprise-class databases.
- **Strategy.** We evaluated this section based on the projects' strategy and vision; the scope and strength of any strategic alliances among system integrators, partners, and resellers; overall strategy; and cost.
- **Market presence.** To evaluate each project's market presence, we looked at its installed base, the number of paying customers, the scope and size of support and implementation organizations, and its financial viability.

**Figure 1** Evaluation Criteria

CURRENT OFFERING	
Transaction processing	What transactional support does the product offer?
Data warehousing	What data warehousing capabilities does the product offer?
Embedded database	What embedded database features are supported by the product?
Programmability	What programmability features does the product offer?
Data types and interfaces	What data types and interfaces are available for the product?
Database availability	What availability features are supported by the product to keep the database running 24x7 and with minimal issues?
Security	What database security features are offered by the product?
Tools	What are the various tools supported by the product?
Platforms	What other features are offered by the product?
Administration	What level of administration efforts and resources are required to support the product?
STRATEGY	
Product strategy	How strong is the project's strategy?
Cost	What kind of license is supported? How much does it cost to deploy and support the product?
Collaboration and control	What is the level of openness of the project?
MARKET PRESENCE	
Adoption	How many downloads does the project report?
Training	How strongly do technology partners support the product?
Services and support	How strongly do technology partners support the product?
Systems integrators and partners	How strongly do technology partners and integrators support the product?

Source: Forrester Research, Inc.

## Vendors Selected For Enterprise-Class DBMS Features And Production Deployments

Forrester included six projects in the assessment: Derby, Firebird, Ingres, MySQL, Oracle, and PostgreSQL. Each of these projects has:

- **The core database engine available as open source.** Each project has its core database engine module available as open source and is available under one of the open source licensing terms.
- **Enterprise-class DBMS features.** Each project has enterprise-class DBMS features in the areas of programmability, availability, interfaces, and security. We made this assessment by comparing the projects' features to the technical evaluation criteria.
- **Enterprises using it for critical applications.** Each project has active customers that have deployed the open source database products in production, to support mission-critical applications.

## OPEN SOURCE DATABASES OFFER SUFFICIENT FEATURES FOR MOST APPLICATIONS

The open source databases continue to narrow the gap with closed source database technology in features and functionality. Forrester estimates that 80% of all applications only require 30% of the features found in the top DBMS closed source product. Many applications only want the basic data management features, so open source database projects are a good fit because they offer this functionality. The Forrester Wave™ evaluation uncovered a market in which (see Figure 2):

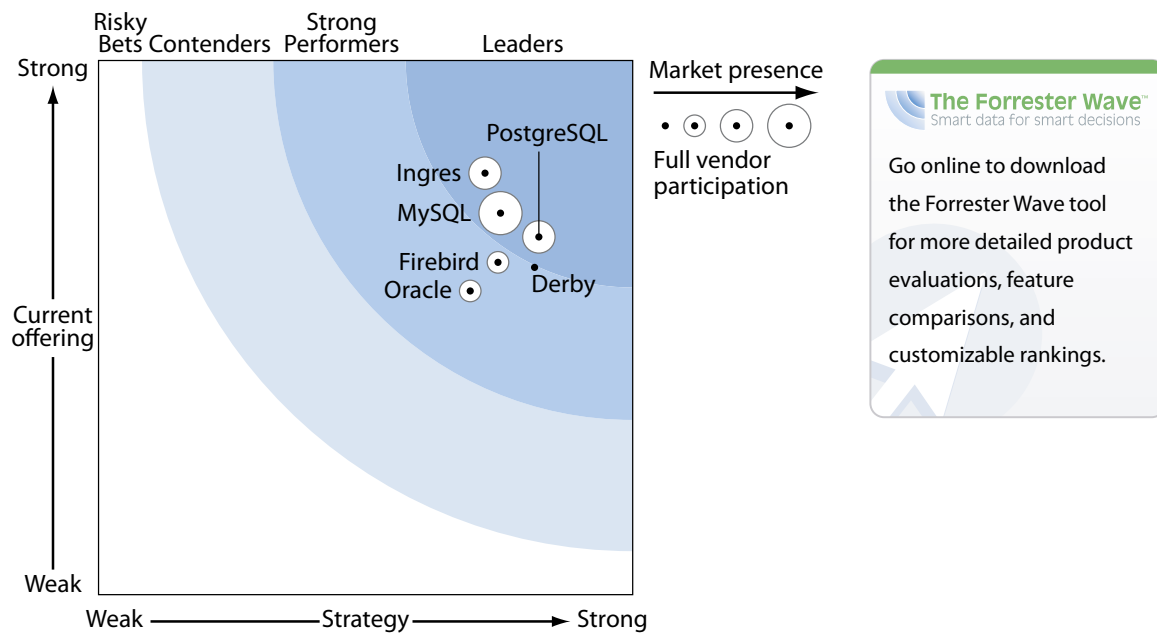
- **Ingres, MySQL, and PostgreSQL lead the pack.** Ingres has strong transactional and data warehousing support, and it scored high in programmability, availability, and security. MySQL has well-balanced support for DBMS features including strong support for embedded database applications, in-memory databases, and ease of use. PostgreSQL scored high on data warehousing, data types and interfaces, programmability options, and support on most platforms.
- **Derby, Firebird, and Oracle offer competitive options.** Derby has strong support for security, administration, and data types and interfaces, but it lacks support for data warehouses and comprehensive database tools. Firebird offers strong support for programmability and moderate support for transactional features, but it lags in security and availability features. Oracle has good transactional processing and strong support for embedded databases, but it lags in programmability, data types and interfaces, security, and tools.
- **Data warehousing features are lagging, but third-party vendors fill the gap.** None of the projects offer comprehensive support for data warehouses. Although Ingres and PostgreSQL have some basic data warehousing features, they still lag behind when compared to closed

source database offerings. Third-party vendors such as Bizgres help fill some of the gap for PostgreSQL by offering comprehensive support for data and index partitioning, parallel query, and bitmap indexes.

- **Oracle and MySQL lead the embedded databases.** Oracle Berkeley DB and MySQL offer comprehensive support for embedded databases, including support for complete in-memory databases. Both offer good programmability options, expanded support for data types and interfaces, and zero-administration capability.

This evaluation of the open source databases market is intended to be a starting point only. Readers are encouraged to view detailed product evaluations and adapt the criteria weightings to fit their individual needs through the Forrester Wave Excel-based vendor comparison tool.

**Figure 2** Forrester Wave™: Open Source Databases, Q2 '06



Source: Forrester Research, Inc.

**Figure 2** Forrester Wave™: Open Source Databases, Q2 '06 (Cont.)

	Forrester's Weighting	Derby	Firebird	Ingres	MySQL	Oracle	PostgreSQL
<b>CURRENT OFFERING</b>	50%	3.04	3.08	3.92	3.55	2.84	3.32
Transaction processing	10%	4.20	4.25	5.00	4.15	4.50	4.45
Data warehousing	10%	0.00	1.90	2.50	0.00	0.00	2.50
Embedded database	10%	2.10	1.80	2.20	4.40	4.60	0.00
Programmability	15%	3.25	4.30	4.40	3.45	0.65	3.80
Data types and interfaces	10%	4.55	4.55	4.70	4.55	3.20	4.85
Database availability	10%	2.10	2.40	4.55	3.00	3.70	2.35
Security	10%	2.80	0.90	3.00	2.60	1.50	3.10
Tools	5%	2.20	3.00	4.60	5.00	2.20	4.40
Platforms	10%	4.00	3.60	5.00	4.80	3.80	4.80
Administration	10%	4.70	3.40	3.30	4.30	5.00	3.20
<b>STRATEGY</b>	50%	4.08	3.74	3.62	3.76	3.48	4.12
Product strategy	40%	3.00	3.05	4.40	4.30	4.20	4.00
Cost	30%	4.60	3.40	2.60	3.80	2.40	3.40
Collaboration and control	30%	5.00	5.00	3.60	3.00	3.60	5.00
<b>MARKET PRESENCE</b>	0%	1.59	2.52	3.20	4.25	2.51	3.32
Adoption	40%	1.50	3.45	1.95	4.85	2.85	3.60
Training	10%	1.00	2.60	4.20	1.60	2.10	3.80
Services and support	30%	2.05	2.05	4.00	5.00	2.05	3.00
Systems integrators and partners	20%	1.35	1.30	4.00	3.25	2.70	3.00

All scores are based on a scale of 0 (weak) to 5 (strong).

Source: Forrester Research, Inc.

## VENDOR PROFILES

### Leaders

- **Ingres.** Ingres remains the most comprehensive open source database project with broad focus on transaction processing, availability, security, programmability, and data warehousing, but it lags in complete in-memory database support and ease of use. Ingres has the longest history among all open source database projects, spanning more than two decades, but it has largely been a commercial database project. CA open sourced Ingres two years ago and then spun it off last year, forming Ingres as an independent entity. Forrester believes that being an independent entity will help Ingres revive its focus and develop the much-needed ecosystem to compete against MySQL. Ingres best suits customers that are looking for databases to meet high-volume transactions for OLTP applications and want to build small to midsize data warehouses.<sup>1</sup>

“We have now been using Ingres for more than five years to support several critical applications. We found Ingres easy to use and scalable, with good DBMS features, so we quickly switched from SQL Server to Ingres, to save money. Although using open source is still risky, with Ingres, I feel confident that we get a reliable product and premium support. Our Ingres database has been running uninterrupted for the past 15 months; what else can you expect from a low-cost open source database?” (Nonprofit organization)

- **MySQL.** MySQL continues to have the largest mindshare in the industry, and with the release of version 5 last year, it has narrowed the gap with Ingres and PostgreSQL, projects that offer the most comprehensive database features. With MySQL becoming more aggressive in churning out new releases, unless Ingres focuses on innovation, MySQL could potentially overtake Ingres to claim the top spot in the feature and functionality race. MySQL has broad support for transactional applications and embedded databases, especially around in-memory databases, but it lags in data warehousing, lightweight directory access protocol (LDAP) integration, and programmability features.<sup>2</sup>

“We have approximately 25 MySQL databases supporting our eCommerce application, each of which is 20 GB in size. We chose an open source database because of cost. MySQL is easier to use than Oracle or SQL Server, partly because it’s a simpler product. MySQL does not have all of the bells and whistles that closed source databases have, but it was able to meet our requirement and deliver excellent reliability and performance.” (eCommerce company)

- **PostgreSQL.** PostgreSQL is a Leader in the open source database category, with good support for OLTP and data warehouse applications. It offers comprehensive support for transactions, programmability, data types and interfaces, security, and platforms, but it lags in embedded database support, high-availability features, and distributed queries. Last year, PostgreSQL got a boost when FUJITSU, Pervasive Software, and Sun Microsystems announced 24x7 technical

support and services to the PostgreSQL community. In addition, new vendors joined the race to extend PostgreSQL functionality, including EnterpriseDB and GreenPlum. PostgreSQL is best suited for customers that want comprehensive DBMS features for transactional and data warehousing applications, or those that are looking for advanced database security features.<sup>3</sup>

“We started using PostgreSQL two years ago, doing a complete migration of a critical internal application running on an Oracle database. Compared to Oracle, PostgreSQL isn’t perfect, but it works very well for us. Our largest database on PostgreSQL is 150 GB and growing 20% annually. We have three internal critical applications that run on PostgreSQL and support hundreds of users. Overall, we are very pleased with PostgreSQL, especially now that Sun Microsystems and Pervasive have joined the race to support PostgreSQL.”  
(Services company)

### Strong Performers

- **Derby.** Overall, Derby offers good DBMS features and functionality with comprehensive support for embedded Java databases in the areas of programmability, interfaces, data types, and administration. With IBM and, more recently, Sun Microsystems supporting Derby under their own brand names, the adoption of Derby has seen some traction, but it still lags behind other open source projects. Derby is best suited for customers that want to use a data repository for Java applications, or require a small footprint with zero-administration functionality.<sup>4</sup>

“Derby provides strong support for SQL, and the fact that it was an embedded database for Java made it an excellent fit for our solution. Our requirement was that the database repository be tightly integrated into our product, and that it required no administration support. Besides that, it had to be reliable and cost effective. Derby was a good fit. Also, with IBM and Sun promoting support so heavily, we could rely on their services and support. I would say Derby in particular is excellent for small to medium-scale databases.”  
(Technology company)

- **Firebird.** Overall, Firebird offers good support for developing Web-based applications that require comprehensive features around programmability, data types, and industry-standard interfaces. Although Firebird has good DBMS technology, it remains leaderless, with no vendor driving the project. Firebird has good penetration in the developer community, especially in Europe and Australia, but it lags in mission-critical production deployments, largely because of lack of comprehensive technical support available from vendors. Firebird is best suited for customers that want an easy-to-use database to support small transactional and Web-based applications.<sup>5</sup>

“We have more than 25 databases that run on Firebird for various internal and Web-based applications. Overall, we are very pleased with Firebird DBMS. It offers good features and functionality to support any kind of application at a very low cost. Our Firebird environments have been very reliable, and we have not felt the need for any comprehensive

technical support. However, occasionally we do get development issues resolved by the Firebird community, which is very active. We are now planning to expand another five more applications to Firebird during the next six months.” (eCommerce company)

- **Oracle.** Overall, Oracle’s Berkeley DB offers good DBMS features and functionality, with broad support for embedded databases, in both market share and technology innovation. Although Berkeley DB does not offer SQL, stored procedures, triggers, views, or functions, it has broad coverage for transactions, database availability, in-memory database, security, and zero-administration to support most embedded database requirements. Berkeley DB also offers support for Java and XML data that makes it an attractive platform to deploy various kinds of applications. Oracle’s acquisition of Sleepycat should further help in extending the Berkeley DB technology and increasing its adoption. Oracle is also likely to integrate Berkeley DB with other Oracle products in the coming years, offering a broader spectrum of data management solutions.<sup>6</sup>

“We have been using Berkeley DB for the past two years for our enterprise software. We use Berkeley DB to share information among various applications related to resources. Berkeley DB is a small, easy, and fast database that is critical for our product. We did not need complex queries or complex databases; we just needed a simple and easy-to-use database with no administration requirements, and Berkeley DB delivered that.” (Technology company)

## SUPPLEMENTAL MATERIAL

### Online Resource

The online version of Figure 2 is an Excel-based vendor comparison tool that provides detailed product evaluations and customizable rankings.

### Data Sources Used In This Forrester Wave

Forrester used a combination of three data sources to assess the strengths and weaknesses of each solution:

- **Vendor surveys.** Forrester surveyed vendors on their capabilities as they relate to the evaluation criteria. Once we analyzed the completed vendor surveys, we conducted vendor calls where necessary to gather details of vendor qualifications.
- **Product demos.** We asked vendors to conduct demonstrations of their product's functionality. We used findings from these product demos to validate details of each vendor's product capabilities.
- **Customer reference calls.** To validate product and vendor qualifications, Forrester also conducted reference calls with two of each vendor's current customers.

### The Forrester Wave Methodology

We conduct primary research to develop a list of vendors that meet our criteria to be evaluated in this market. From that initial pool of vendors, we then narrow our final list. We choose these vendors based on: 1) product fit; 2) customer success; and 3) Forrester client demand. We eliminate vendors that have limited customer references and products that don't fit the scope of our evaluation.

After examining past research, user need assessments, and vendor and expert interviews, we develop the initial evaluation criteria. To evaluate the vendors and their products against our set of criteria, we gather details of product qualifications through a combination of lab evaluations, questionnaires, demos, and/or discussions with client references. We send evaluations to the vendors for their review, and we adjust the evaluations to provide the most accurate view of vendor offerings and strategies.

We set default weightings to reflect our analysis of the needs of large user companies — and/or other scenarios as outlined in the Forrester Wave document — and then score the vendors based on a clearly defined scale. These default weightings are intended only as a starting point, and readers are encouraged to adapt the weightings to fit their individual needs through the Excel-based tool. The final scores generate the graphical depiction of the market based on current offering, strategy, and market presence. Forrester intends to update vendor evaluations regularly as product capabilities and vendor strategies evolve.

## ENDNOTES

- <sup>1</sup> View the vendor summary for more detailed analysis on how Ingres fared in this evaluation. See the June 14, 2006, Tech Choices “[Ingres Leads In Open Source Databases.](#)”
- <sup>2</sup> View the vendor summary for more detailed analysis on how MySQL fared in this evaluation. See the June 14, 2006, Tech Choices “[MySQL Leads In Open Source Databases.](#)”
- <sup>3</sup> View the vendor summary for more detailed analysis on how PostgreSQL fared in this evaluation. See the June 14, 2006, Tech Choices “[PostgreSQL Leads In Open Source Databases.](#)”
- <sup>4</sup> View the vendor summary for more detailed analysis on how Derby fared in this evaluation. See the June 14, 2006, Tech Choices “[Derby Is A Strong Performer In Open Source Databases.](#)”
- <sup>5</sup> View the vendor summary for more detailed analysis on how Firebird fared in this evaluation. See the June 14, 2006, Tech Choices “[Firebird Is A Strong Performer In Open Source Databases.](#)”
- <sup>6</sup> View the vendor summary for more detailed analysis on how Oracle fared in this evaluation. See the June 14, 2006, Tech Choices “[Oracle Is A Strong Performer In Open Source Databases.](#)”

# FORRESTER®

Helping Business Thrive On Technology Change

## Headquarters

Forrester Research, Inc.  
400 Technology Square  
Cambridge, MA 02139 USA  
Tel: +1 617/613-6000  
Fax: +1 617/613-5000  
Email: [forrester@forrester.com](mailto:forrester@forrester.com)  
Nasdaq symbol: FORR  
[www.forrester.com](http://www.forrester.com)

## Research and Sales Offices

Australia	Israel
Brazil	Japan
Canada	Korea
Denmark	The Netherlands
France	Switzerland
Germany	United Kingdom
Hong Kong	United States
India	

*For a complete list of worldwide locations,  
visit [www.forrester.com/about](http://www.forrester.com/about).*

For information on hard-copy or electronic reprints, please contact the Client Resource Center at +1 866/367-7378, +1 617/617-5730, or [resourcecenter@forrester.com](mailto:resourcecenter@forrester.com). We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research (Nasdaq: FORR) is an independent technology and market research company that provides pragmatic and forward-thinking advice about technology's impact on business and consumers. For 22 years, Forrester has been a thought leader and trusted advisor, helping global clients lead in their markets through its research, consulting, events, and peer-to-peer executive programs. For more information, visit [www.forrester.com](http://www.forrester.com).