


2015 DATA CONNECTIVITY OUTLOOK



Today's Database Marketplace

It's no secret that we are living in a world of Big Data. The sheer scale of the Internet, the proliferation of mobile devices, the explosion of social media, and the emerging world of the Internet of Things promise to make Big Data even bigger over time.

What was once a relatively quiet world of running your business on a traditional relational database has become noisy – full of new choices and new technologies.

Last year, in our inaugural database survey, we highlighted the strength of the leading relational vendors, and the variety of architectures and new entrants into the market. This recalled the “database wars” of the '90s, and the resulting consolidation, making us wonder how the next inevitable shakeup might play out in the coming years.

In this year's survey, the choices of database features and vendors continue to expand. Customers must research and decide how to choose between traditional relational databases, NoSQL solutions and Big Data platforms, and whether to host the data locally or in the cloud. While traditional relational databases still predominate, many emerging products with specific advanced capabilities targeting analytics are creating a niche and gaining traction in the marketplace.

Our 2015 survey, based on feedback from our diverse customer population, is designed to help you sort through the hype and see which vendors are creating buzz and which are more likely to consolidate in the coming years.

About This Survey

In order to better understand in the changing nature of the database world, and to help enterprises navigate these changes, Progress DataDirect recently conducted its second annual Data Connectivity Survey for 2015. The survey respondents encompassed 580 customers, comprised of CxOs, Managers, Data Scientists, Business Analysts and Developers. The survey questions focused on currently installed technology as well as planned direction for the next two years.


The objective of this survey is to provide a view of the current state of the database marketplace as well as anticipated trends.

Relational Databases:

Top Players and Emerging Contenders

We start with relational databases, the bread and butter of online transaction processing platform (OLTP) systems today. The traditional SQL Server and Oracle databases continue to have the dominant share. Among our respondents, 59% and 41% respectively have those databases installed. MySQL (if you combine the Community and Enterprise versions) comes in on par with Oracle. IBM DB2 comes in next with only one-third of the presence of MySQL or Oracle. Open source PostgreSQL rounds out the top players.

Relational



Emerging Products Show Notable Growth

Beyond the top players is an interesting collection of emerging relational database products, each with single-digit presence. But what separates these from the front runners is the growth rate. While our respondents are not planning to grow their footprints for the front runners, most of the new entrants are showing growth – and a few are showing significant growth. Amazon Redshift, only installed with 3% of those in our survey, is showing growth to 5% over the next two years. While these are admittedly small numbers, the fact that they are showing any growth is notable, particularly when compared to DB2, which is projected to fall by half, from 14% of our respondents to only 7% over the next two years.

It is interesting to examine the emerging products to understand why they are growing, and why the traditional leaders seem to be waning. SAP HANA, for example, provides a column-oriented, in-memory, relational database. Sybase IQ, also a product from SAP, is a column-oriented, petabyte-scale, relational database. Amazon Redshift provides a cloud-based, column-oriented, subscription-based database with massive, parallel processing.

Each product has a unique specialization that is making it attractive as the world moves toward Big Data and web scale.

One feature that most of the emerging products has is that they are column-oriented. By storing data in columns, they tend to shine as analytical platforms. Online transaction processing platforms (OLTP) just aren't as efficient as the platforms designed and optimized for analysis.

As Big Data continues to find its way into corporations, the analysis of that data is becoming more important and driving the growth of online analytical platforms [OLAP].


And it is important to note that the survey is by no means predicting the demise of Oracle or SQL Server. The world still runs on those products; their market share is substantial. However, 2015 is shaping up to be the year when the shift toward specialized databases becomes significant.

NoSQL Shows Growth as a Big Data Store

In addition to the growth in adoption of specialized relational platforms, we are also seeing growth in non-relational platforms, often referred to as NoSQL. These systems aren't bound by the rigid schemas and tabular structure of a generalized relational database. They are highly scalable and can reasonably be referred to as Big Data stores for their ability to provide web-scale data storage. And in contrast to the traditional relational databases discussed above, each of the NoSQL databases is promising growth in deployment over the next two years.

Open source MongoDB is the leader of the pack, having been adopted by 26% of our customer base. SimpleDB, a cloud-based offering from Amazon, comes in next at 12%, and Cassandra, an Apache open source architecture, comes in third at 9%. Hbase, Couchbase, DynamoDB and Riak round out the results.

NoSQL



So why does over half of our population have NoSQL in their shops? The respondents indicated that the ease of development and maintenance of applications was the primary motivation. They found that it is much easier to throw data into a NoSQL database without having to worry about rigid schemas. Additionally, respondents liked the scalability of these platforms.

Commercially supported distributions of free open-source solutions showed growth as well.

One theme that appeared was the growth in adoption of open source solutions. The answer is clear: if you can have a database that is easy to develop against, easy to maintain, and low cost or free, you have a winner.

Hadoop Leads the Way in Adoption

The third category of data store examined was Hadoop. The Hadoop distribution from Apache was the leader in adoption with 28% usage among our respondents. The other choices in our survey represent commercial distributions of Hadoop. Companies bundle their own management, software, and support to add to value to the core Apache Hadoop software.

Oracle Big Data Appliance came in second behind Apache Hadoop, with 18% of our respondents. BDA is a full-rack system with Sun servers, networking, and software included. Cloudera and IBM BigInsights tie for third place in our Hadoop category at 10% each.


What is particularly noticeable about this category is the growth rate. Strong growth among our customer base is projected. Amazon Elastic Map Reduce wins the growth contest with usage set to grow 117% among our respondents over the next two years. Cloudera Impala, the SQL query engine that runs with Hadoop, is the runner up with plans for 105% growth reported.

It is also worth noting that partnerships in the Hadoop space abound. Oracle teams with Cloudera for management software for its BDA offering. Pivotal HD can have Hortonworks under the covers. Offerings from Microsoft and Teradata use Hortonworks as well.

Hadoop is a relatively new data source – its birth at Yahoo was barely 10 years ago – and key components such as the HDFS file system are only a couple of years old. A variety of commercial Hadoop distributions are available when companies are ready to deploy business-critical applications.

The availability of free open source software enables anyone to try out Big Data.

Hadoop



Big Data Keeps Getting Bigger

It has been 45 years since the relational database was introduced and 35 years since the debut of Oracle's first commercial product. The world has changed significantly since then. As technology has advanced, the reduced cost of storage, bandwidth, and mobile device and infrastructure has led to major shifts toward web scale, Big Data, cloud computing, and analytics. Creative database projects initiated at Google, Yahoo, Facebook and others have led to several new open source Big Data architectures.

The velocity, volume, and variety of data is increasing. Velocity from millions of mobile devices and millions of sensors in the Internet of Things can strain real-time systems. The volume of data can quickly fill terabytes. And the data are no longer structured transactions; they can be posts, photos, transactions, tweets, videos—you name it.

The flood of data is accelerating. How can an enterprise effectively prepare for the future?

Be Prepared for Future Opportunity

In order to make sense of the plethora of options, IT service managers and business line managers should start by getting together and asking questions, such as:

- ▶ How might the increased opportunity for real-time, Big Data, and NoSQL analytics help our business?
- ▶ How has the velocity, volume, and variety of our data change over the last few years and how do we expect it to change going forward?
- ▶ How fast is data growing, and how will we store it in the next 3 years? Should we start looking at Hadoop?
- ▶ How will we expand our data and analytics capacity to meet the coming demand, and will we do it in our own data center, in the cloud or a hybrid?
- ▶ Where will we find the expertise, or how will we train our staff in these new technologies?
- ▶ Will we use free open source software, or will we rely on turnkey commercial distributions?

The last several years has seen an explosion in technology and created many opportunities to enhance the way businesses can be run, and there is no indication that it is slowing down. The winners will be those managers who embrace the changes and capitalize on the new opportunities that the new world of data will afford them.

Learn more

<https://www.progress.com/solutions/data-connectivity>

About Our Survey

To build the 2015 Progress Data Connectivity report, Progress surveyed 580 customers in small, medium, and large businesses. The survey was global with regional distribution of:

- ▶ North America—44%
- ▶ Europe—23%
- ▶ Asia/Pacific—21%
- ▶ Latin America—10%

To ensure a comprehensive survey across organizations, survey participants spanned roles throughout the organization. The distribution of roles consists of:

- ▶ CXO—9%
- ▶ IT & Business Mgmt—24%
- ▶ Product or Program Mgmt—22%
- ▶ Development—35%
- ▶ Other—9%

PROGRESS

Progress [NASDAQ: PRGS] is a global software company that simplifies the development, deployment and management of business applications on-premise or in the cloud, on any platform or device, to any data source, with enhanced performance, minimal IT complexity and low total cost of ownership.

WORLDWIDE HEADQUARTERS

Progress Software Corporation, 14 Oak Park, Bedford, MA 01730 USA Tel: +1 781 280-4000 Fax: +1 781 280-4095 On the Web at: www.progress.com

Find us on [f](https://www.facebook.com/progresssw) facebook.com/progresssw [t](https://twitter.com/progresssw) twitter.com/progresssw [y](https://www.youtube.com/progresssw) youtube.com/progresssw

For regional international office locations and contact information, please go to www.progress.com/worldwide

Progress and EasyI are trademarks or registered trademarks of Progress Software Corporation or one of its affiliates or subsidiaries in the U.S. and other countries. Any other marks contained herein may be trademarks of their respective owners. Specifications subject to change without notice.

© 2015 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

Rev 04/01 | 150326-0063