


Nexus 7 from Google.
The perfect tablet for the holidays.



nexus

Buy Now

[Slashdot.org](#)

General Topics

- [TV](#)
- [Resources](#)
- [Jobs](#)
- [Log In](#)
- [Join](#)

Search this website ...

Search

- [Business Intelligence](#)
- [Cloud](#)
- [DataCenter](#)

10

Choosing a Database That's Right for Your Business

by [Robin Miller](#) | April 23, 2012

4

 [Repost This Article](#)

Database selection is crucial for any business that anticipates handling large amounts of data. But how do you make the right choice?

The connection was reset

The connection to the server was reset while the page was loading.

- The site could be temporarily unavailable or too busy. Try again in a few moments.
- If you are unable to load any pages, check your computer's network connection.
- If your computer or network is protected by a firewall or proxy, make sure that Firefox is permitted to access the Web.

Try Again

Database selection can be a major head-scratcher. You talk to one expert and hear: “[MySQL](#) will do everything you need, and then some.” Another says, “[Microsoft SQL Server](#) is easier to administer than MySQL and at least as powerful.” A third might mention the lesser-known MongoDB, pointing out its use by Disney, craigslist, Foursquare, and Intuit.

Josh Berkus (who stars in a little video embedded in this article) prefers [PostgreSQL](#), with which he’s been closely associated for many years. At the same time, however, he acknowledges that just about any popular modern database will handle almost all commercial database tasks.

Natalie Kaminski, founder and CEO of New York-based [Fincode Solutions](#), a self-described “boutique software and web application development firm,” prefers Microsoft SQL Server: “We prefer Microsoft SQL Server since it’s pretty simple to set up and maintain. We are not limited by what it is compatible with, as you may be with ‘packaged’ bundles.”

Microsoft SQL Server is popular for SMB applications, especially in companies where most of the IT staff has strong Microsoft experience and is not fully conversant with open source.

The “packaged bundles” Kaminski mentioned are all-in-one solutions from vendors such as SAP, Oracle and IBM. Those bundles can sometimes prove more expensive than taking an “a la carte” approach to purchasing IT solutions. Nonetheless, other IT pros (including Berkus) point out that packaged solutions are easy for a CEO or CTO to select, as they require only one set of negotiations with one vendor—and one bill to pay each month or year.

A growing number of vendors are positioning the cloud as a solution for businesses’ relational database need. [Amazon’s RDS](#) (Relational Database Service) offers MySQL and Oracle; Oracle now offers [Oracle DB as a public cloud service](#), and other companies including [Rackspace](#) and [JoyentCloud](#) offer various cloud-based database services to corporate customers, from one-person businesses to Fortune 100 multinationals.

To Cloud, or Not to Cloud

But Kon Leong, President and CEO of [ZL Technologies](#), warns that the “take it to the cloud” trend is still more hype than reality. His company specializes in a very particular database niche—archiving data—and markets itself primarily toward Fortune 500-sized firms. Nonetheless, Leong’s advice could prove useful for companies of any size:

1. You can no longer throw anything away. In our litigious society, he said, you need to keep records well beyond the old seven-year limit. How long? No one knows. Some lawsuits deal with issues several decades in the past—which means a whole lot of archived data.

2. Not only will you keep accumulating data, you may need to retrieve it quickly. Besides Leong’s company, others specialize in archiving voluminous amounts of data. Whatever option you choose, they can range from scalable, cloud-based databases (which you might already use) to a hybrid cloud installation.

The idea of a hybrid cloud is simple: you have data you’re using now, which is critical to your day-to-day operations. Beyond a certain point in time (and this point may be different for different parts of your business, as well as different data components within each part) you can relegate data to a virtual storage facility—the archive—where you can still call up what you need. At the same time, the archive doesn’t have (and likely doesn’t need) sophisticated GUI front ends or load-balanced servers to handle sudden spikes in user requests.

In fact, you might even be able to store your archived data on hard drives you keep in a safe deposit box, in addition to storing it either on your premises or remotely in “the cloud.” There’s simply no “one size fits everyone” answer when it comes to databases. That’s why a select portion of consultants and integrators specialize in data storage.

Each Expert has Different Database Preferences

One database expert espouses Microsoft SQL Server, while another is a major player in the PostgreSQL community. Albert Factora, a Senior Partner at [Groove Analytics](#) has yet another take—and a recommendation or two. He wrote:

“There are three critical attributes that must always be addressed: security, scalability and cost, and there are two main types to choose from: relational or non-relational. A relational database is ideal for already-homogenized transaction-based information such as banking or accounting records.”

Relational databases, he continued, are appropriate for many data-warehousing implementations because “most of the data looks the same,” with a finite number of attributes for each object-type and “easy ways to relate objects on one collection (i.e. table to objects in another collection.”

So-called [NoSQL](#) systems, on the other hand, are appropriate for data with a variable number of attributes or columns for each main item in a particular data set—especially “for data where there is simply no way of knowing in advance what attributes will exist for a particular object within a set.”

“Fortunately, there are open source solutions to support both types of data,” Factora suggested. “The more mature platforms of both types can scale to very high record counts, and include substantial tools for implementing security at all levels.”

For data that’s loosely structured, with objects that feature a variable number of attributes (with some attributes composed of large binary data—think sound or video files), then a NoSQL database might prove best. “A top contender in this arena is [MongoDB](#), an open-source platform with commercial licensing and support available,” he added. “10Gen is the company which handles this. MongoDB offers high performance, high availability and scalable solution licensing.”

For applications that need a relational database solution, MySQL is often indicated. “This mature open-source relational platform offers little or no cost for a development environment, and native versions are available that can run on all major desktop and server platforms,” he wrote. “MySQL instances are also available through virtual data centers supported by companies like Amazon, Dell or Rackspace.” Another benefit: MySQL can scale in response to business strategy, and supports industry-standard security policies. Oracle (which took over MySQL when it acquired Sun) offers commercial support and licensing for production implementations.

MySQL is also everywhere. It is the database behind an overwhelming majority of the world’s blogs and CMS-based websites.

Finding an Appropriate Expert

If you know databases and their proposed purpose within your organization, selecting a vendor and installing one is a fairly easy process. If not, then some additional questions come into play: Are you looking primarily to run your business and stay on top of sales, ordering, and inventory? Are you going to do lots of marketing research based on customer information? Or both?

That “both” option is entirely feasible with modern relationship databases. Take this hypothetical corporate database with three purposes:

1. *Everyday operations*
2. *Marketing research*
3. *Archival “CYA” data storage*

It’s entirely possible that a CIO or IT administrator would rely on more than one database (as long as they all

handle data in the same formats) for these three types of data use. But which platforms to choose?

This is where memberships in local business organizations, such as a Chamber of Commerce, come in handy for business owners. You can ask fellow members for referrals to competent IT (and specifically database) consultants and packagers. An IT person or department is another obvious and appropriate resource when shopping for IT products and services.

Kaminski also suggests that moving from one database to another or moving data from spreadsheets to a database "can be painful when the data was not strictly monitored and maintained over its lifetime." (Quick example: naming the same company in different ways on multiple spreadsheets.)

In other words, by making good choice, you save yourself from having to move from one database to another for a long, long time.

10 comments 6 people listening

Sign in

livefyre 

+ Follow conversation

Post to  

Post comment as



anonymous 74 pts

This line is funny: "the database behind an overwhelming majority of the world's blogs and CMS-based websites."

So you should trust your mission critical BI data to... the database behind an overwhelming majority of the world's blogs and CMS websites. Riiiiiiiiight.

Let's contemplate the pro-MySQL arguments in this article. You could make those exact same arguments if you were trying to convince somebody to buy a fleet of Toyota Yaris cars when they need to haul concrete bricks around. Toyota Yaris cars are everywhere, they're ubiquitous and supported by Toyota... and they are used by an overwhelming non-minority of college students to get to class. Therefore, they are ideal for hauling around thousands of pounds of concrete bricks!



anonymous 74 pts

I like MongoDB... because mongo is web scale:

<http://www.youtube.com/watch?v=b2F-DitXtZs>



fistynuts 7 pts

This is utter drivel.

