

MYSQL™ PROVIDES SCALABILITY, RELIABILITY, & ENTERPRISE SUPPORT AT A LOWER COST

Open Databases
White Paper
April 2009

Abstract

Enterprises are looking for the best, most cost-effective way to manage the ever-increasing amount of data, an ever-growing pool of users, and unrelenting demands for high availability. Database administrators must ensure their database system can handle this level of growth with limited downtime—and without a high cost. As proprietary database systems have grown increasingly complex with features that many companies don't need, they have also become harder to manage and more expensive in terms of upgrades and staffing, training, and support. The Sun MySQL™ database delivers the scalability and features that companies need but with less complexity and lower costs than proprietary systems. The world's widely used open-source database, MySQL has more than 12 million installations.

Table of Contents

Executive Summary	1
Open Databases	1
Staying Ahead of Enterprise Data Growth	1
Why Open Source for Database Systems?	2
MySQL: Sun’s Industry-Leading Open Database	3
Scalability and Availability	4
What to Consider When Considering an Open Database	5
Open Database Case Studies	6
MySQL Helps Zappos Cost-Effectively Scale to \$1 Billion in Sales	6
Alcatel-Lucent Taps MySQL for Mission-Critical System	6
Swedish National Police	7

Executive Summary

More data, more user connections, more critical applications: It all adds up to greater demands on database systems than ever. Developers of proprietary solutions have tried to keep up by piling on more features, but have succeeded mainly in ratcheting up prices. Increasingly, organizations are looking to open source for effective yet cost-efficient database alternatives. Many have chosen MySQL™ as a result. As of April of 2009, MySQL has an install base of over 11 million and is downloaded more than 60,000 times per day.

MySQL is the most widely used open-source database in the world. This white paper looks at how open source in general leads to more streamlined database solutions that are easier and less costly to deploy and support, and also at how open source has made MySQL in particular a higher-performing, more reliable, and more affordable alternative to proprietary solutions. It also shows how MySQL is benefiting a number of companies, including:

- Saving a major online retailer more than \$1 million a year
- Delivering a solution for a telecommunications technology company at a cost ten times less than others the company evaluated
- Reducing costs by more than 50% as part of an overall open source strategy for a large national law enforcement organization in Western Europe

Open Databases

Staying Ahead of Enterprise Data Growth

Explosive data growth has left many organizations struggling to manage a staggering amount of information, cope with an unprecedented number of database users, and meet a steady, uncompromising demand for availability. The Data Warehousing Institute puts the average data growth rate for a typical business today at 46%, which may be conservative in many cases; within Sun Microsystems, data grew 200% in 2008. And internetnews.com reports that IDC research has per-person data requirements growing at an annual rate of 60%.

For the database administrator, the challenge associated with data growth today is threefold:

1. How much data and how many user connections can my database system handle?
2. How much downtime can the system sustain before the financial toll on the business becomes substantial?
3. How much is it going to cost to successfully manage challenges 1 and 2?

Meeting today's database challenge requires technology that can scale to support vast stores of data that are continuing to grow all the time; deliver on the promise of constant availability; and—increasingly important—keep database system costs in check even as demands placed on it continue to develop.

Increasingly, many are finding that the way to achieve scalability and availability while keeping costs in check is an open-source database. An open-source solution can not only scale to handle massive amounts of data and growing numbers of user connections, but can also help lower the total cost of ownership (TCO) of database systems.

Why Open Source for Database Systems?

With many proprietary database systems, one of the greatest problems is posed by their complexity. As with many other proprietary technologies, these systems have grown overly complex as more and more features have been added to them and as they have evolved to meet the demands placed on them by data growth. These systems may require near-constant, highly skilled IT attention to ensure their proper operations, which makes them time-consuming and difficult to work with. Often, they have more database features and capabilities than are actually needed or used by the people who invest in them. And little has been done historically to make all those features and capabilities easy to configure and use.

An open-source database, on the other hand, is inherently easy to work with. It is likely to focus, as open-source projects tend to do, on a manageable number of features that will attract the interest of a broad community of developers and users. To appeal to that broad community, an open-source database will necessarily be simple and straightforward. From deployment, configuration, and setup to ongoing operations and technology management, an open source database can reasonably be expected to perform reliably with little need for tweaking and tuning by skilled personnel.

Another side effect of excessive complexity in many proprietary database systems is high cost. Each year, the companies that sell proprietary databases introduce new “upgraded” systems—or systems with more features and more complexity—and charge ever-increasing prices for them. And the cost of ongoing operations—for things like staffing and training and support is steep, because of the growing complexity of the systems. There's also the cost of downtime, which is a real risk with systems that require so much care and support to run properly on a day-to-day basis.

With an open-source database, streamlined product design and ease of deployment can be coupled with the backing of strong support organizations to help keep acquisition and operating costs down.

MySQL: Sun's Industry-Leading Open Database

MySQL is Sun's open-source alternative to proprietary solutions for online and transactional database needs. MySQL has long been proven in the marketplace. It's the most widely used open source database in the world, with 12 million active installations and 70,000 downloads a day. Its popularity is a reflection of the growing need for a database system that delivers the scalability and features that customers need, but with less complexity and lower costs than proprietary systems. Available in an enterprise version that includes production support and monitoring tools with the database server, MySQL is used by companies like Google, Yahoo!, Alcatel-Lucent, and Nokia Siemens Networks for mission-critical applications.

An often-expressed concern about open source has to do with quality assurance (QA). Since, by definition, open source involves code contributions from an open community of users and developers, how is the quality of the resulting product maintained? With its large installed base and so many major companies relying on MySQL, quality assurance is critical. To help achieve the highest level of quality, every community contribution goes through Sun's stringent quality control process before being incorporated into the product. Beyond that, the answer to the quality assurance question lies to a great extent in the installed base itself. With millions of people using the software every day, there is a tremendous, transparent effort going on all the time to identify and resolve issues. Compare this with a company offering a major proprietary database project; there may be thousands of people in QA at that company working on the same thing. But MySQL has 70,000 people downloading the product every day, using it, and reporting on their experiences to improve the quality continually. The process is entirely open to them; if a user wants to know how many bugs are in any particular MySQL solution at any given time, the information is all available on line—what the bugs are, who's working on them, and how fast they're being fixed. It's a level of engagement and transparency that's unheard of with proprietary solutions, and it ultimately contributes to the quality of the product.

MySQL draws on a number of unique characteristics of open source to enable several important product benefits that have become increasingly less common in proprietary databases: reliability, performance, and ease of use. Without the inherent complexity of proprietary systems, MySQL performs reliably and at a higher level than those systems without requiring extensive troubleshooting due to bottlenecks or other slowdowns. It also incorporates a number of performance-enhancing mechanisms such as high-speed load utilities, distinctive memory caches and full text indexes.

Because it requires less attention to perform at a high level, MySQL is easier to use than proprietary solutions. And its open design allows it to incorporate a variety of

types of database storage engines—high-speed read-and-load engines, transaction engines, clustered engines, and memory engines—so that it can act as a single solution for an organization that has a variety of requirements. No other database offers this ability to choose from a broad range of native and third-party storage engines to meet specific needs. Through this capability, MySQL can be configured to meet the most demanding performance expectations for applications such as high-speed transactional processing systems, high-volume web sites and large data warehouses.

MySQL includes among its native storage-engine options a high-performance query engine (MyISAM) that is combined with features such as a fast data insert capability and strong support for fast full-text searches and other specialized Web functions. This has made it the de facto standard for high-traffic Web sites. Additionally, the third-party choices for storage engines for MySQL include InnoDB, one of the most powerful transactional database engines available.

The ability to offer a wealth of options for configuring MySQL to meet specific application needs is a direct result of MySQL coming out of the open-source community. In open source, people are working all the time to extend the core functionality—whether by building new engines, developing plug-ins that add functions, or making code contributions that will fundamentally improve the product.

Scalability and Availability

MySQL is, then, an open-source database that can perform as well as or better than proprietary alternatives. But what about the key concerns for scalability and availability when the number of users and amount of data involved are constantly increasing? How does this open source solution handle these requirements?

MySQL is proven in the field as a database server that can run anything from deeply embedded applications with a footprint of just +3-4MB to massive data warehouses holding terabytes of information. It also can be configured with a variety of high-availability capabilities such as high-speed master/slave replication and specialized cluster servers with instant failover. This level of flexibility to attain optimum scalability and availability for a particular set of needs is a hallmark of the open approach to database technology that underlies MySQL.

Most critically, MySQL enables extremely high levels of scalability and performance at costs far lower than those of proprietary database solutions—in terms of both the upfront license costs and the ongoing operational costs that make up the total cost of ownership. Savings will vary between companies and projects, but on average, companies that deploy MySQL have experienced per-project savings of between \$250,000 and \$500,000. Larger projects or enterprise deployments may see savings in excess of \$10 million.

These savings are attributable to the following factors:

- **Large reductions in database licensing costs.** MySQL Enterprise, Sun's subscription-based, all-inclusive, open-source database offering, is less expensive than proprietary alternatives because it is sold on a per-server basis, not by the number of CPUs, chips, or cores. MySQL also offers an option for organizations to deploy an unlimited number of servers for a fixed annual fee. Other databases include extra priced options, that drive up the cost—and the vendors also typically charge an additional annual maintenance fee.
- **Significantly less system downtime.** As businesses have moved from batch processing to real-time enterprise systems, the impact of any outage can take a tremendous toll in lost productivity, lost revenue, increased call center loads, and higher IT costs. MySQL reliability circumvents these risks. Using it can reduce downtime by 60% or more, leading to potential savings of more than \$100,000 for most companies. For larger organizations, improvements may be measured in millions of dollars of costs avoided.
- **Lower hardware costs.** Organizations can reduce their hardware expenditure significantly by deploying applications on open-source technology that runs on low-cost commodity hardware. MySQL is specifically recognized for its superior performance on low-cost, Intel-based hardware.
- **Reduced administration, engineering, and support costs.** A combination of low complexity and high reliability lowers the cost of developing, maintaining, and supporting database applications using MySQL.

What to Consider When Considering an Open Database

If you're thinking of moving to an open-source database system, there are several things to consider to ensure that you make the most effective choice. One concern that some have expressed about using an open source database is that the feature set will be inadequate. Look for a system with a rich—or rich enough—feature set, i.e., one that offers the capabilities to meet your specific needs, but without the superfluous features that drive up the cost of proprietary solutions. And if you need a database to support large amounts of data for mission-critical applications, choose a solution that has proven scalability to support large data stores, as well as one with demonstrated system reliability and availability.

Support is also an important factor in selecting an open database. Look for a product that offers a full complement of technical support services, including 24/7 support, just as a proprietary solution would. And look for a database that is backed by both an internal staff overseeing support and ongoing development and a large community of users who can contribute to the quality of the product on a continuing basis. For more information or to download MySQL go to mysql.com.

Open Database Case Studies

MySQL Helps Zappos Cost-Effectively Scale to \$1 Billion in Sales

Zappos is one of the world's largest online retailers with over \$1 billion in annual sales. The company sells shoes, handbags, and eyewear as well as other apparel. However, its primary focus is delivering superior customer service. Zappos believes delivering the best customer service is key to a successful online shopping experience.

Since Zappos was founded in 1999, the company has used MySQL as its primary database to power its Web site, internal tools, and reporting tasks. As Zappos has grown, MySQL has been able to scale with the business, making it a perfect solution even at the current sales volume. It's been an important piece of infrastructure that Zappos has scaled as the company has grown.

MySQL plays a critical role in delivering that customer service by providing Zappos with:

- High performance and scalability enabling millions of customers to shop on Zappos.com every day
- 99.99% database availability so that Zappos' customers don't experience service interruptions that impact revenue
- A cost-effective solution that saves Zappos over \$1 million per year, allowing them to spend more money on their customer service and less on their technical infrastructure

Compared to proprietary enterprise systems, Zappos estimates it is saving about \$1 million per year in licensing fees and salaries of dedicated DBAs that can only manage individual systems.

Alcatel-Lucent Taps MySQL for Mission-Critical System

Alcatel-Lucent is a worldwide leader in providing end-to-end communications technology solutions that enable service providers to deliver communication services to end users. The company was looking for a database solution to use with its mission-critical XDMS building-block system for creating new applications. Alcatel-Lucent selected MySQL Cluster Carrier Grade Edition, a telecommunications-specific version of MySQL, based on the product's:

- High performance and fast replication capabilities, that enable real-time access to subscriber data
- High availability and demonstrated scalability with large volumes of data
- Ability to scale out incrementally, adding capacity on an as-needed basis to accommodate millions of users
- Low cost — at least ten times less than shared-disk clustered databases that the company evaluated

This is not the first time Alcatel-Lucent has turned to MySQL Cluster Carrier Grade as the database solution for a critical system. Before selecting it for the XDMS system, the company was already using it to power subscriber databases and mission-critical HLR and HSS solutions. The XDMS system will be used for new and innovative telephony and TV-oriented applications.

Swedish National Police

With an existing infrastructure consisting of a heterogeneous mix of ancient architectures (including mainframes), the Swedish National Police was facing serious cost increases for its legacy proprietary applications. In order to save taxpayers' money and minimize risk, the organization needed to become more committed to one clear, scalable strategy and phase out its outdated systems.

The Swedish National Police is now deploying open-source software and MySQL. The organization has committed to building all future IT systems on an open-source software infrastructure stack, including MySQL and JBoss. Compared to proprietary solutions, the group conservatively estimates that it will achieve at least 50% savings in the first five years.



Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN (9786) Web sun.com

© 2009, Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, and MySQL are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the United States and other countries. Information subject to change without notice. Printed in USA 04/09 XX0000-0/#K