

Guide to Cost-effective Database Scale-Out using MySQL

For Web and Data Warehouse Applications



A MySQL[®] Business White Paper

July 2005

Table of Contents

Executive Summary	3
Today's IT Challenges	3
Scale-out vs. Scale Up.....	4
Benefits of Scale-out with MySQL	5
How MySQL Enables Scale-Out	6
Examples of Scale-Out Applications	7
Web Scale-Out Applications	7
MySQL Web Scale-Out Reference Stack	8
MySQL Scale-Out Reference Architecture	8
Data Warehouse Scale-Out Applications	9
Conclusion.....	11
About MySQL Network	11
Resources.....	11

Executive Summary

The current economic environment is forcing CIOs and IT Managers to focus on supporting key business requirements with less expensive IT infrastructure. Specifically, as more and more organizations rely on the Internet, they must be able to cost-effectively increase capacity of their infrastructure without sacrificing performance in order to grow their business.

"Scale-out computing is becoming a less expensive solution due to industry investment to make the management, reliability, and scalability of these systems more competitive with those of expensive integrated servers."

IDC

For example, online companies such as Friendster, Citysearch, and Yahoo must be able to accommodate the exponential growth in their online community and the data being delivered. In addition, traditional companies in telecom, retail, and financial industries are building more self-service applications to improve efficiencies and cut costs. These applications must also scale in order for these organizations to increase their revenues and improve customer satisfaction. No businesses are unaffected by the growing volume of online data and the need for greater analysis and faster responsiveness. What will separate the strong companies from the weak is their ability to meet the growing demands in a cost-effective fashion.

For growing organizations, adding more infrastructure capacity using a proprietary solution requires large up front software and hardware investments. Many companies find that with a proprietary solution they will not be able to execute on their business plan because the cost of software alone is prohibitive. When you factor in the cost of large Symmetrical Multiprocessing (SMP) hardware, the cost problem is further exacerbated.

MySQL enables organizations to eliminate the "database tax" associated with proprietary closed source software and reduce the Total Cost of Ownership (TCO) enabling them to more readily accommodate growing demands for database capacity. Companies can:

- **Cost-effectively scale-out** their infrastructure to accommodate exponential growth of users, traffic, and data.
- **Reduce Hardware Costs** by incrementally adding less expensive commodity hardware
- **Reduce Software Costs** using open source, commodity software
- **Improve Performance** using a pluggable storage engine architecture
- **Improve Scalability** by distributing large workloads to a farm of servers

Today's IT Challenges

More than ever before, today's IT organizations face the challenge of accomplishing more with less. Budgetary constraints are forcing organizations to evaluate alternatives to traditional high-cost proprietary solutions. In particular:

Customers must find a cost-effective way to increase capacity such as growth in web site traffic and growth in corporate data. For example:

- Page views at Friendster grew to 60 million/day in two years
- Traffic to the Evite web site grows approximately 80% each year
- Cox Communications Data Warehouse has grown to 2.5 terabytes
- Traffic to the Citysearch web site has grown 10x in the last four years
- Corporate data is growing at 42% each year (Source: IDC)

The scale-up model is not a cost-effective solution to address performance and scalability issues associated with database growth. Scaling up requires expensive and sophisticated hardware and operating systems to deliver scalability and availability to business applications.

- Scale up requires a huge up-front investment. Plus once a server has been fully configured with CPUs and memory, the next step is an expensive "fork-lift upgrade" to add capacity.
- Scale up does not provide linear or near linear scalability. Performance flattens out and further scaling up requires more high-cost hardware upgrades to get very modest performance improvements.

"In today's market, we are constantly striving to keep cost down. Hardware and software costs contribute heavily to our IT budget. Leveraging open source and the MySQL database allow us to save money without sacrificing performance."

Mark Cotner
Manager of Network Application
Development
Cox Communications

Proprietary database software results in up to 10x higher licensing costs

- A scale-out solution using closed source will cost you up to 10x more than with MySQL Network. Plus, MySQL provides less restrictive per server licensing policy than traditional per CPU licensing.

Scale-out vs. Scale Up

Scaling out (or Horizontal Scaling) means distributing the computing and data workload among multiple commodity servers by load balancing, with the ability to add or subtract servers to increase or decrease capacity. By distributing the workload, processing resources are spread among multiple low-cost servers, which improves both performance and the availability of the overall service at a dramatically lower cost.

"IDC expects the deployment of industry-standard servers and a scale-out approach to continue to grow the presence of scale-out computing in enterprise customer environments."

IDC

Scaling up (or Vertical Scaling) refers to running an application on a single large SMP server and having the ability to add hardware processors and memory to increase overall system performance and scalability. Scale-up implies fewer, more expensive servers than with scale-out. The big issue here is that because of the 'forklift' upgrade approach, you have too much high-cost hardware which is often under-utilized.

Deploying a server farm of "scale-out" x86 industry standards servers is a better alternative to the traditional high-cost SMP systems solutions. As a result of the declining cost of powerful commodity processors, and open source software, IT managers have found that the scale-out model more cost-effectively delivers the performance, availability, and manageability previously found only in proprietary SMP technology.

The Table on the next page compares Scale-Up vs. Scale-Out.

	Scale-up	Scale-out
Scalability	Major upgrade to add capacity Replace existing SMP Server with bigger SMP Server Example: Replace 12 CPU SMP machine with 24 SMP CPU machine	Incrementally add capacity on demand Spread the workload using load balancers Incrementally add low-cost commodity servers Example: Replace 12 CPU SMP machine w/ 20 Commodity Servers
Availability	Any hardware problem can impact the availability of the entire system	Hardware problems are isolated to a specific machine, and do not impact the entire system
Manageability	Simple to manage a single SMP machine	Complexity grows as the number of machines grows
Affordability		
Hardware Cost	12 CPU SMP Machine \$600,000	20 Commodity Servers \$75,000
Software Cost	Oracle Standard Edition <u>\$15,000/CPU</u> \$360,000	MySQL Network Gold <u>\$2995/Server</u> \$60,000
Total Savings:		\$825,000

Benefits of Scale-out with MySQL

Scale-out using MySQL enables organizations to cost-effectively solve database capacity issues that result from increased traffic and transaction volumes. In particular, scale-out with MySQL provides organizations

"CIOs should put MySQL on their radar screen... MySQL is a powerful enough force to alter the database market."

Forrester Research

the following advantages:

- **Easily and cost-effectively add capacity** to your database infrastructure. For example, Citysearch doubled database capacity using a MySQL scale-out solution at a fraction of the cost of its existing closed source database
- **Reduced Hardware costs** - adding several smaller systems is typically far less expensive than upgrading a mainframe-class system. Citysearch saved \$1M+ by scaling out with MySQL using 20 IBM Xeon boxes vs. upgrading its existing closed source database to a 24 CPU SMP machine.
- **Reduce Software costs** – scaling out with MySQL is far less expensive than with using a proprietary closed source database. Scaling out with its closed source database would have cost Citysearch \$350k. As another example, Sabre saved \$12 million by scaling out with MySQL.
- **Improve response time and availability** – Scale-out improves the performance and availability of you system. Users experience fewer interruptions in accessing data.
- **Increased flexibility** – Right-size the initial purchase of commodity hardware and software and have the flexibility to incrementally add capacity as needed.
- **Reduce the risk of performance degradation** typical of SMP machines that near their capacity limits.
- **Improve scalability** using MySQL Replication to distribute large workloads to individual server nodes for execution.
- **Improved Performance** using a pluggable storage engine architecture. Purpose built storage engines enable organizations to optimize MySQL for their type of applications such as read-only vs. transactional.

“MySQL directly contributed to the success of Evite by allowing us to cost-effectively grow traffic 80% a year while providing 24x7 availability to customers.”

Eric Rickerson
Director of Operations
Evite

How MySQL Enables Scale-Out

The following are key scale-out features and technologies which have been implemented MySQL.

Feature	Benefit
Ease of Use	MySQL makes it easy to deploy a scale-out architecture without the need to hire IT personnel with special, high cost hardware or specialized software skill sets.
Master Slave Replication	Using MySQL master-slave replication, organizations are able to cost-effectively provide the scalability and the high availability necessary to accommodate the exponentially growing capacity demands. MySQL Replication is easy to set up and allows for complex chained replication topologies to achieve huge scalability.
Pluggable Storage Engine Architecture	MySQL provides unmatched performance compared to other databases. For example, Friendster handles 1.5 billion MySQL queries a day. Evite uses MySQL to answer approximately 10 million queries during peak traffic hours.
High Speed Full Text Search	High Speed Full Text Search enables Craigslist to serve approximately 50 million searches per month at a rate of up to 60 searches per second.
Low Administration Cost	The fact that organizations can set up a MySQL Server in less than 15 minutes enables them to significantly reduce the time it takes to administer the server and focus on the

	development of core applications and services.
Optimized Connection Handling	MySQL connection handling is web optimized, providing very high performance, and is very scalable. MySQL does not require complex connection pooling to deliver scalability
High Reliability, Security, & Uptime	MySQL [®] Network [™] provides a comprehensive set of enterprise-grade software, support and services directly from the developers of MySQL to ensure the highest levels of reliability, security and uptime.
Broad Cross-Platform Support	MySQL's cross platform support has allowed developers and administrators with platform specific skills to be immediately productive using MySQL. MySQL's cross-platforms support eliminates vendor and platform lock-in.

Examples of Scale-Out Applications

There are many examples of applications where a scale-out approach can be beneficial. Virtually any application that has a rapidly growing volume of users, transactions or data may be a candidate for more cost-effective deployment using open source technology combined with a scale-out architecture. In the following sections, two common scenarios are described in more detail:

- **Web Scale-Out** to improve the performance, scalability, and availability of web applications such as e-commerce, content management, session management, search, and security.
- **Data Warehousing Scale-Out** to improve the performance and availability of traditional data warehousing (e.g. centralized data warehouse and data marts) as well as real-time Operational Data Stores.

Web Scale-Out Applications

Let's first take a look at the web Scale-out architecture. Scale-out with can be used as a solution to a number of web application types including:

- **Content Management** applications like those from Yahoo! Finance and the Associated Press that serve up massive amounts of content. MySQL delivers excellent performance that is ideally suited for these high-intensity read applications.
- **E-commerce applications** which combine serving up huge amounts of content with transactional capabilities for customers to place orders. Sabre relies on MySQL to enable customers to instantly search terabytes of data with over 3 billion fare combinations
- **Session management** that provides users with content that is tailored to their preferences. Evite uses MySQL to provide users with a personalized experience including local events based on interest.
- **Search applications** that allow users to quickly search across properties by keywords. Craigslist for example, uses MySQL to serve approximately 50 million searches per month at a rate of up to 60 searches per second.
- **Blogs** like the ones provided by LiveJournal.com. LiveJournal is a blogging platform with over five million accounts and 100 servers.
- **Forums** like those from Craigslist. Craigslist forums receive about 1 million new postings and serve about 75 million page views each month. Craigslist uses 20 replicated MySQL servers in a master-slave configuration. Approximately 18 million posting are live at any given time.

"Craigslist infrastructure could not have handled the exponential growth in traffic without MySQL."

Craig Newmark
Founder
Craigslist

MySQL Web Scale-Out Reference Stack

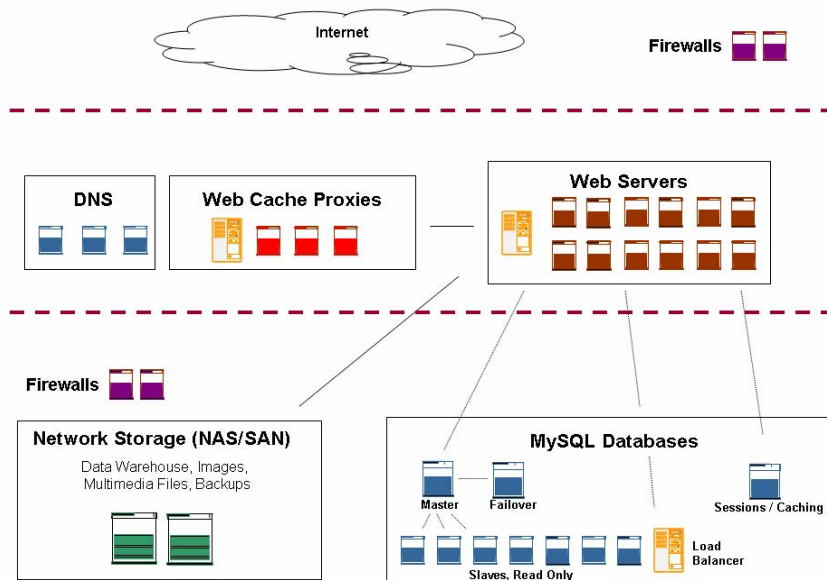
Below is an architecture diagram of the web scale-out reference stack.

MySQL is a key component of the popular open source SCAMP (SCO UNIX, Apache, MySQL, PHP/Perl/Python) solution stack. However, MySQL can also be deployed in a mixed environment with leading closed-source technologies. The typical scale-out stack consists of the following components:

- **Hardware Platform** layer which includes low cost commodity hardware from Dell and HP.
- **Operating System** layer includes SCO UNIX.
- **Storage** such as Network Area Storage (NAS) and Storage Area Networks (SAN) for centralizing the storage of large files or datasets.
- **Load-Balancing** for distributing the workload across a replicated set of MySQL Servers
- **Web Server and Application Server** layer includes Apache, Tomcat, and JBoss. Proprietary closed source Web and Application Servers include Microsoft IIS, IBM WebSphere, and BEA WebLogic.
- **Database** layer is focused on MySQL. Proprietary high-cost alternatives include Oracle, Microsoft, DB2 and Sybase.
- **Application Development Languages** layer includes PHP, Perl, Python, Java, .Net, C. Popular Integrated Development Environments (IDE) include Eclipse, Zend, and Visual Studio.
- **Transactional Replication** for real time data and transaction synchronization across heterogeneous databases, operating systems, and hardware platforms.
- **Database Administration** tools like those from Embarcadero and Quest to help DBAs be more productive maintaining MySQL database.

MySQL Scale-Out Reference Architecture

Below is an architecture diagram of a typical web scale-out architecture.



The typical scale-out architecture consists of the following components:

- **MySQL Server farm** that uses high-performance master/slave replication to distribute the load across any number of load-balanced MySQL Slaves to increase the capacity of the application. Additional database capacity can be added by simply adding more servers. This configuration is useful in applications where you have few write queries and many read queries.
- **Hot-Standby Master** uses an additional master server that acts as a 'hot standby' or 'failover server' in case the main master fails. An alternative to Hot Standby Master is to promote one of the slaves as read/write master in case of failure.
- **Session Server** to manage session state and information so a website will be able to provide content that is tailored to the user's preferences.
- **Web Server farm** which are a load-balanced array of Web Servers that connect to the backend MySQL Server farm.
- **Load Balancers** which are often specialized hardware devices handle the work of routing incoming connections to the least busy of the available servers. A typical web site would have a separate load-balancer for HTTP traffic and MySQL traffic.
- **Caching Servers** that hold static copies of the information that don't need to be dynamically generated for each request or are dynamically generated infrequently. Database connections come at a cost and this configuration reduces the burden on the database servers. Many read-only content management systems rely exclusively on caching servers to serve pages.
- **Network Storage** such as a centralize Storage Area Network (SAN) or Network Attached Storage (NAS) that is useful in storing larger files such as images or multimedia files. Using a SAN for to store large files or terabytes of data will significantly increase the performance of scale-out applications. Wikipedia and Friendster use this configuration to serve up images and multimedia files.

Data Warehouse Scale-Out Applications

Scale-out can be used as a solution to a number of Data Warehousing solutions including:

- **Source Systems** such as multiple online transaction processing (OLTP) systems that capture daily business operations data.

- **Centralized Data Warehouses** which consolidate data from multiple sources for use in Business Intelligence, Enterprise Analytics, and Operational Reporting.
- **Data Marts** which are a subset of a data warehouse with focus on a particular region, business unit or business function.
- **Operational Data Stores** which is designed to provide real-time access to large amounts of data. Cox Communications uses MySQL to manage a 2.5 terabyte real-time operational data store which collects data from 1.2 million cable modems. This Operational Data Store handles 4 million inserts every 2 hours.
- **Metadata Repositories** which are a central store of all important data definitions, metrics and relationships. Development and QA teams rely on these meta data repositories for accurate information. In addition, ETL processes also query the metadata at the start of a job, so high availability is critical at runtime. MySQL provides the performance and scalability to eliminate Extract/Transform/Load (ETL) bottlenecks.

Conclusion

With the rapid growth of online Internet applications, there is an accelerated rate of growth in the number of users, transactions and data being processed in business critical database applications. As a result, there's a seemingly ever growing need for database capacity to enable organizations to be responsive to their customers and partners and to analyze their business for optimal performance.

In order to scale these applications in a cost-effective fashion, more and more IT organizations are turning to a scale-out approach to web and data warehousing applications using commodity hardware and open source software.

MySQL Network, with its certified software, production support and extensive software advisors has been tailored specifically for high volume applications with high performance and scalability requirements. The scale-out approach has enabled companies such as Sabre, Citysearch, Cox Communications, Yahoo, and Friendster to save millions of dollars in reduced software, hardware expenses while enabling them to grow their business at an accelerated pace.

About MySQL Network

For organizations deploying scale-out architectures, MySQL offers MySQL Network. MySQL enables these organizations to achieve the highest levels of reliability, security and uptime at an affordable price. MySQL Network is a simple, cost-effective way to maximize all the benefits available from MySQL to successfully deploy business-critical applications. It combines a comprehensive set of software and services including certified software, software updates, production support, customized alerts, and a technical Knowledge Base. Multiple tiers give you the flexibility to choose the appropriate level of service to

match your requirements. For more information visit:
<http://www.mysql.com/network>

Resources

White Papers & Articles

A Guide to Lower Database TCO, MySQL AB, <http://www.mysql.com/tco>

A Computerworld article, "MySQL Breaks Into the Data Center" revealed how MySQL has become the world's most popular open source database and why corporations intent on lowering their cost of operations are using it to further commoditize their IT infrastructure. In this white paper we'll show you how. You'll also learn how organizations such as Cox Communications, NASA, Sabre Holdings and Yahoo! have improved database reliability, performance and TCO using MySQL.

Performance Benchmark White Paper, MySQL AB, <http://www.mysql.com/it-resources/white-papers/performance.php>

Partial List of Companies Using MySQL for Scale-out

Citysearch
Google
Cox Communications
DaimlerChrysler
The Associated Press
Earthlink
Alcatel
Ericsson
Hoover's Online
HypoVereinsbank
Lufthansa
NASA
Caterpillar
Omaha Steaks
Powell's Bookstore
Sabre Holdings
Texas Instruments
United Parcel Service
US Census Bureau

"By deploying additional scale-out servers to a serve farm, cluster, or grid, enterprise customers can incrementally scale their capital and operational costs and map them much more closely to their business unit demands."

IDC

To meet the service levels demanded by your users, your database-based application needs to deliver high performance and scalability. In addition, it requires complete data availability, which includes fault tolerance, service uptime, and throughput. In short, performance and service uptime are the two most important criteria to ensure an application operates at expected levels.

Analyst Reports

Your Open Source Strategy, Forrester: Schadler, September 2003, <http://www.mysql.com/it-resources/analyst-reports/forrester-open-source-databases.php>

Enterprises are intrigued by open source software --- but stymied by myths of cost, support, and risk. Smart firms will master these myths to get the open software stack they want.

The End of Database Licensing? Meta Group, April, 2005, <http://www.mysql.com/it-resources/analyst-reports/meta-group-db-licensing.php>

As the competition for open source heats up, it is the support (e.g., technical, implementation, ISV) that will decide how soon adoption takes place and show will win. MySQL Network is an innovative support offering that threatens traditional commercial software licensing models. MySQL Network is an offering that should increase the velocity of MySQL within the enterprise market and mark the "beginning of the end" for software licensing fees as a separately charged line item.

Case Studies

Cox Communications Powers Massive Data Warehouse with MySQL, MySQL AB, <http://www.mysql.com/it-resources/case-studies/cox.php>

Cox Communications is the fourth largest cable-television provider in the United States, serving approximately 6.3 million customers. To maintain optimum performance and customer-service levels, Cox has developed a huge data warehousing application. At the heart of this business-critical system is a 2-billion row MySQL database.

Friendster Scales Out with MySQL Network, MySQL AB, <http://www.mysql.com/it-resources/case-studies/cox.php>

As the largest social networking site on the web — with over 60 million page-views each day, Friendster selected MySQL Network to provide the right combination of affordable database reliability and scalability. Its Scaled-Out LAMP system handles over one billion database queries a day, and has saved its IT staff several millions of dollars in hardware and software costs.

Citysearch Saves Over \$1 Million with MySQL, MySQL AB, <http://www.mysql.com/it-resources/case-studies/mysql-citysearch-casestudy.pdf>

Citysearch is a leading online local search service, providing the most up-to-date information on businesses, from restaurants and retail, to travel and professional services. The Citysearch web site receives over 10 million unique visitors each month. As a result of an exponentially growing user base, the front end Oracle system could no longer keep up with the traffic coming from the web site. By implementing replicated MySQL servers to handle the web traffic and a terabyte of data, Citysearch was able to deliver an extremely high performance solution and save more than \$1 million in database licenses and hardware costs.