
Virtualization for MySQL on VMware®: Best Practices and Performance Guide

**Mike Frank
Senior Product Manager
Sun Microsystems MySQL**

**Jennifer Glore, MySQL Technology Lead
Adina Kalin, Staff Engineer
Sun Microsystems ISVe**

A MySQL® Business White Paper

Table of Contents

1	INTRODUCTION	3
2	WHY RUN MYSQL DATABASE ON VMWARE® INFRASTRUCTURE.....	3
3	SOFTWARE COMPONENTS.....	4
4	HARDWARE INFRASTRUCTURE	5
5	TEST OVERVIEW	6
6	TEST RESULTS.....	8
7	BEST PRACTICES	9
8	CONCLUSIONS	10
9	RESOURCES.....	10

1 Introduction

Recently Sun/MySQL and VMware jointly ran tests for MySQL on VMware® ESX in order to:

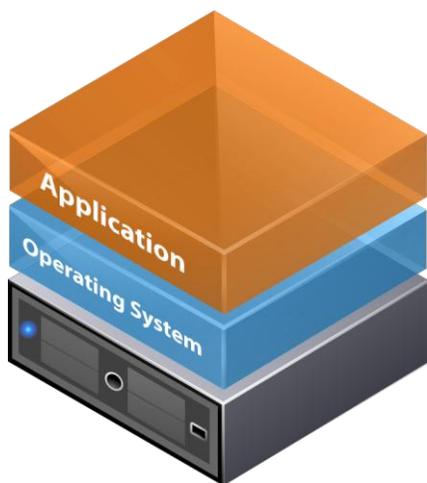
- Provide a comparison of virtualized instances of MySQL Server instances running on VMware ESX with non-virtualized MySQL Server instances running on an equivalent hardware server.
- Provide “Best Practice” configuration guidance for running MySQL Server in a virtualized VMware environment

In a nutshell, performance variance is minimal on a properly configured system, as shown in tests on a typical small to medium business level server with MySQL running in a VMware virtualized environment. This document explains how we reached this conclusion by describing the setup, configuration, and results from our testing. This information should be invaluable to anyone running or considering running MySQL Server on VMware Infrastructure.

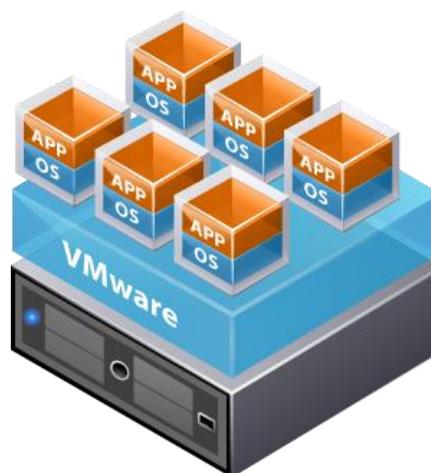
The evaluators also performed additional tests regarding the performance of MySQL with VMware® High Availability (HA) and VMware® Distributed Resource Scheduler (DRS) for load balancing, migration, and recovery. Details of that testing series will be covered in a second document.

2 Why run MySQL Database on VMware® Infrastructure

Virtualization dramatically improves the efficiency and availability of resources and applications in IT organizations. Internal resources are underutilized under the old “one server, one application” model and IT administrators spend too much time managing servers rather than innovating. An automated datacenter, built on a VMware virtualization platform, allows them to respond to market dynamics faster and more efficiently than ever before. [VMware Infrastructure](#) delivers resources, applications—even servers—when and where they’re needed. VMware customers typically [save 50-70% on overall IT costs](#) by consolidating their resource pools and delivering highly available machines with VMware Infrastructure. The VMware Infrastructure product suite delivers cost savings by maximizing infrastructure efficiency, reducing operational costs and providing cost effective business continuity.



Traditional Architecture



Virtual Architecture

VMware solutions can provide many benefits to MySQL administrators. VMware virtualization provides a layer of abstraction between the resources required by MySQL Server and the operating system and the underlying hardware providing those resources.

Benefits of virtualizing MySQL include:

- **Consolidation:** Consolidate multiple physical servers onto one, with little to no impact on user response time and throughput.
- **Ease of Provisioning:** Encapsulate standard MySQL Server/OS configuration to an image that can then be cloned or moved across hardware resources, thereby reducing the cost of application provisioning and deployment.
- **Manageability:** Use VMware® VMotion to move virtual machines from server to server with no loss of user connectivity, data, or transactions, eliminating planned downtime for common operations such as hardware maintenance.
- **Availability:** Reduce unplanned downtime and maintain higher service levels. VMware HA ensures that in the case of an unplanned hardware failure, affected virtual machines are rapidly restarted on another host in a VMware cluster.

3 Software Components

The software stack used to perform the tests comprised the following products:

Database Server: MySQL Enterprise Server 5.1

- MySQL Enterprise Server software is the most reliable, secure and up-to-date version of MySQL for cost-effectively delivering E-commerce, Online Transaction Processing (OLTP), and multi-terabyte Data Warehousing applications. It is a fully integrated transaction-safe, ACID compliant database with full commit, rollback, crash recovery and row level locking capabilities. MySQL delivers the ease of use, scalability, and performance that has made MySQL the world's most popular open source database. <http://www.mysql.com/products/enterprise/server.html>

Operating System: Linux - RedHat 5.2 ES 64 bit (<http://www.redhat.com>)

Testing Software: SysBench 0.4.10 (<http://sysbench.sourceforge.net/>)

Virtualization software: VMware Infrastructure components including:

- VMware® ESX 3.5 (<http://www.vmware.com/products/vi/esx/>): VMware ESX is a “bare-metal” hypervisor that partitions physical servers into multiple virtual machines. Each virtual machine represents a complete system, with processors, memory, networking, storage and BIOS. VMware ESX enables multiple virtual machines to:
 - Share physical resources
 - Run unmodified operating systems and applications
 - Run the most resource-intensive applications side by side on the same server
- VMware® vCenter Server: VMware vCenter Server, formerly VMware VirtualCenter, centrally manages [VMware Infrastructure](#) environments simplifying day to day IT operations management and making virtual environments easy to control and administer.
 - [Complete virtualization management](#) for the VMware Infrastructure.

- Decrease planned and unplanned downtime for improved [business continuity](#).
- Run fewer servers and dynamically power down unused servers to [reduce energy costs](#).
- Virtual Machine File System (VMFS): A high-performance clustered file system for ESX

4 Hardware Infrastructure

The hardware used to perform the tests comprised the following products:

Server: Sun Fire x4200 M2, 2*dual 2.6GHz AMD Opteron, 16GB RAM,
<http://www.sun.com/servers/entry/x4200/>



The Sun Fire was equipped with two internal 10000 rpm 100 gigabyte disks so we could have one disk for the Linux operating system and one for ESX.

SAN: StorageTek 6140 configured RAID 0+1, directly attached to the server.

http://www.sun.com/storage/disk_systems/midrange/6140/

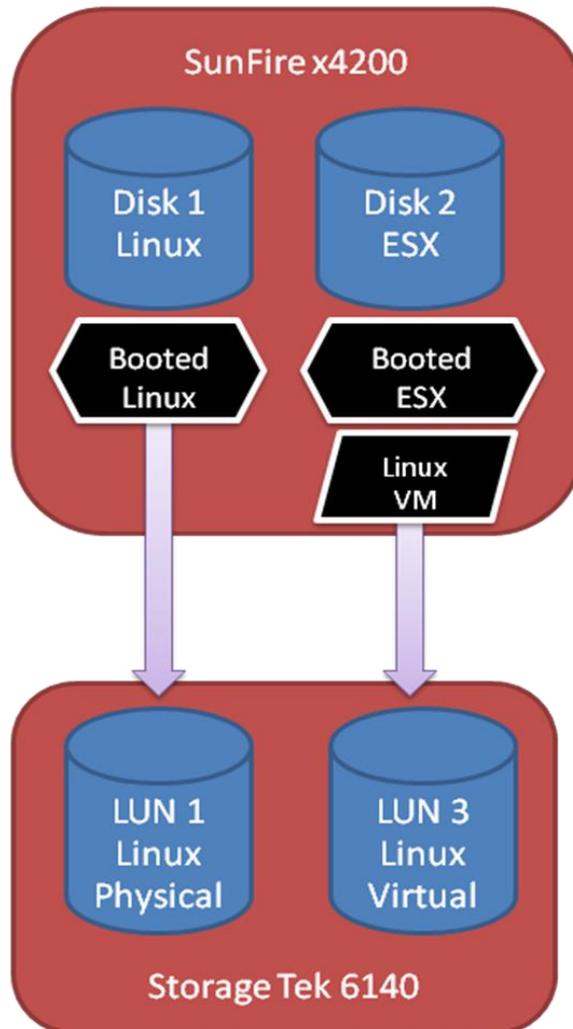


HBA Card: QLogic Corp QLA2432 (rev 02)

5 Test Overview

Our goal was to make this testing a fair comparison of MySQL behavior in virtual environments versus non-virtual environments. All tests were performed on exactly the same hardware system; using matching internal boot disks and SAN LUNs. These two internal disks were set up so we could boot in two different modes:

1. Linux on the physical hardware
2. VMware ESX hypervisor on the physical hardware



The environment was set up to achieve “apples to apples” comparisons for physical and virtual environments:

Physical Environment:

When the Linux operating system was booted directly on the hardware we disabled two CPUs, thus matching the number of CPUs for both the physical and virtual environments.

To disable the CPUs on Red Hat Linux use the following command:

```
echo 0 > /sys/devices/system/cpu/cpu2/online and echo 0 >
/sys/devices/system/cpu/cpu3/online
```

Virtual Environment:

Within VMware ESX, we created two virtual machines with resources and operating system versions equivalent to their physical counterparts: two cpus and 4GB RAM.

Storage Environment:

On SAN storage, we created two LUNs of 100GB each; in other words, one LUN for each test scenario: Linux physical and Linux Virtual. The SAN LUNs were configured exactly the same and varied in name only.

We used the ext3 filesystem:

```
/etc/fstab ( Linux ):
/dev/sde1 /var/lib/mysql ext3 defaults 2 2
```

MySQL Environment:

MySQL Enterprise Server 5.1.30 was installed with each operating system, both physical and virtual, and the my.cnf was set up so that the database files were on a LUN on the StorageTek 6140 unit (since each DB instance had its own LUN as explained earlier). Again the MySQL configurations matched for each environment. We used the default my.cnf.

Test Execution

For each of the two target platforms (Linux physical and Linux virtual), we completed runs for 4, 8, 16 and 32 concurrent threads within SysBench. To push read and write IO we ran SysBench in OLTP mode and used the complex (“advanced transactional”) mode. Using the [complex](#) (“advanced transactional”) mode forces each thread to perform transactions using BEGIN/COMMIT statements and a combination of select, update, delete and insert within these transactions.

Steps for Running SysBench

Before running the SysBench scripts:

The SysBench database, sbtest, was created and all user permissions granted as follows:

```
mysql -uroot -p
create database sbtest;
grant all on *.* to 'root'@'%';
```

(note: that this only needs to happen once per platform)

Before each test script:

The sbtest table was dropped and recreated with sysbench/prepare.

```
mysql -uroot sbtest -e "drop table sbtest"
```

Then the sysbench was prepared

```
./sysbench
--test=oltp
--mysql-table-engine=innodb
--oltp-table-size=10000000
--mysqlsocket=/var/lib/mysql/mysql.sock
--mysql-user=root prepare
```

And finally the test was run varying --num-threads with values of 4,8,16,and 32

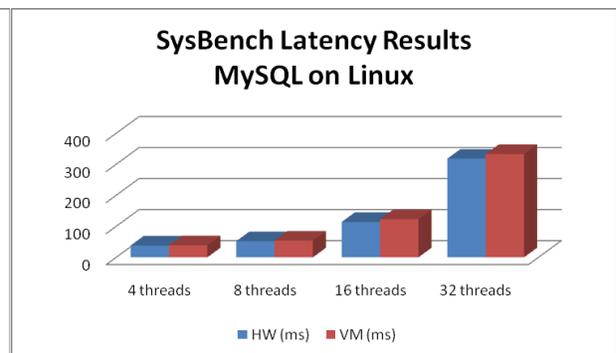
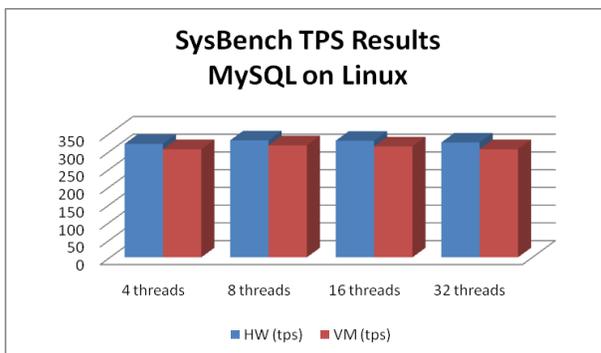
Note: that when you run sysbench the -oltp-table-size should be larger than in the prepare phase. We used -oltp-table-size=80000000

```
./sysbench
  --num-threads=4
  --max-time=900
  --max-requests=500000
  --test=oltp
  --mysql-user=root
  --mysql-host=localhost
  --mysql-port=3306
  --mysql-table-engine=innodb
  --oltp-test-mode=complex -oltp-table-size=80000000 run >
MYRESULTS.txt
```

6 Test Results

The SysBench test scenario was a stress test with no delay that lasted 15 minutes. This heavy stress level scenario was chosen to cover the majority of possible user situations. However, this peak transaction load is atypical of most production systems and represents a peak load that few customers would see for any sustained length of time within their environments. As shown in the results included below, the difference in transactions per second is just under 4.9% on average between a physical and virtual environment. This difference is likely unnoticeable to end users, especially in a more typical situation with a lower transaction volume. Furthermore, in our tests, the number of threads (connections) didn't have any significant effect on performance for either HW or VM results and the differences remained in the same range.

SysBench Threads	4	8	16	32
HW (tps)	320	330	329	324
VM (tps)	305	316	313	305
TPS delta (%)	4.69%	4.24%	4.86%	5.86%
HW Latency (ms)	38	52	114	318
VM Latency (ms)	39	54	123	333
Latency Delta (%)	-2.63%	-3.85%	-7.89%	-4.72%



7 Best Practices

Within our testing scenario, we implemented several best practices that if ignored may compromise performance.

Storage configuration

- RAID
 - Use a RAID configuration that works well for databases. We used 0+1
 - Use a hardware RAID and not a software-based RAID (e.g. a volume manager)
- HBA Card
 - To optimize the throughput of your HBA card(s), you may want to change their maximum queue depth setting within your ESX configuration.
 - In our tests we adjusted the Queue Depth for our QLogic HBA
/etc/vmware/esx.conf
Set ql2xmaxqdepth=128
- General
 - For optimum performance separate the OS and database disks.
 - Spread your virtual disks across your physical disks
 - Be careful not to dynamically reconfigure virtual disks with the MySQL database running. This applies to regular hardware as well, but it's a mistake often made on VMs.

Proper use of the Sysbench tool

As specified above, before any run of sysbench, the sbtest table should be reset (dropped and rebuilt). As well, table size in sysbench/prepare should be smaller than the size you give into sysbench/run. Failure to do this may result in inconsistent results.

VMware Best Practices

General best practices for VMware performance testing include the following:

- Ensure that database queries use indices effectively, thus reducing any unnecessary I/O operations. Overall, additional I/O operations may cause an increase in CPU use compared with physical systems.
- For Red Hat based versions of linux, be aware of the effects of the frequent timer interrupts, which may need to be managed via the divider command. More information on this issue is available in the VI performance paper listed in the resources section below.
- Reduce the number of vCPUs in your virtual machines to the minimum, as larger numbers of vCPUs will incur more scheduling overhead. For example, if running on a four vCPU virtual machine with only 30% CPU utilization, reduce the number of vCPUs to two and expect to run at around 55 to 60% CPU utilization.
- Disable the floppy drive, CD ROM and USB adaptors in your virtual machines. This reduces overhead slightly, and can also reduce the chances of IRQ conflicts on the system. IRQ conflicts are also documented in VMware performance literature listed below.
- Use physical systems that have multiple physical NICs available, and separate network traffic across the two NICs to reduce network contention.

Additional performance best practices can be found in the documents listed in the resources section.

8 Conclusions

These results show that MySQL performed well on VMware Infrastructure, with minimal performance impact. Customers should consider deploying MySQL on this architecture, as it has the potential to reduce capex and opex costs, increase service levels and simplify the manageability of their applications.

9 Resources

Customers can find more information about VMware and MySQL products via the links listed below:

MySQL: <http://www.mysql.com>

VMware: <http://www.vmware.com/>

MySQL Enterprise Server: <http://www.mysql.com/products/enterprise/server.html>

MySQL 5.1 Reference Manual Section 15.1 Common Issues with Virtualization
<http://dev.mysql.com/doc/refman/5.1/en/ha-vm-commonissues.html>

MySQL Virtualization Forum
<http://forums.mysql.com/list.php?149>

MySQL 2006 Quick Poll on Virtualization
http://dev.mysql.com/tech-resources/quickpolls/virtualization_software.html

VMware Infrastructure 3 product web site:
http://www.vmware.com/products/data_center.html

VMware Infrastructure 3 download website:
<http://www.vmware.com/download/vi/eval.html>

VMware support website:
<http://www.vmware.com/vmtn/>

VMware Performance Tuning Paper:
http://www.vmware.com/pdf/vi_performance_tuning.pdf

System Compatibility Guide for a complete list of compatible hardware:
http://www.vmware.com/pdf/vi35_systems_guide.pdf

Storage/SAN Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/pdf/vi35_san_guide.pdf.

I/O Compatibility Guide for a complete list of compatible networking devices:
http://www.vmware.com/pdf/vi35_io_guide.pdf.

Todor Ivanov's Weblog
<http://ivanov.wordpress.com/2008/04/01/database-virtualization/>

Database Systems on Virtual Machines: How much do You Lose?
<http://www.cs.uwaterloo.ca/~ashraf/pubs/smdb08overhead.pdf>

Database Virtualization: A New Frontier for Database Tuning and Physical Design
<http://whitepapers.techrepublic.com.com/abstract.aspx?docid=289253>