

A Proposed Assessment Model of Database Performance on Virtual Machine vs. Traditional Server

Nevine Makram Labib and Amr Ismail Hassan
Computer and Information System Department, Faculty of Management Sciences
Sadat Academy for Management Sciences, Cairo, Egypt

Abstract

Virtual machine technology allow managing a many servers on the same physical hardware it means best utilization ,these technology allow server consolidation, easier deployment, and more flexible provisioning.so in the last years many organizations using virtualization technology with database are increasing and that effect on the database performance.

In this research, an experimental study of running a database workload on a virtual machine and traditional server, then an attempt has been made to evaluate database performance and compares the performance between them.

From the experimentation results, database in traditional server has lower efficiency compared to the database in virtualization server.

Keywords: Virtualization Technologies, Virtual Machine, Oracle Architecture and Oracle Performance.

Introduction

The term of virtualization is add a flexible and programmable layer of Software between the physical hardware and the software running above it and separates the user-perceived notion of resources from their actual physical implementation. The virtual machine monitor (VMM) is between software systems and the computing resources (such as CPU and disk) that they use. VMM allows the resources of a physical machine to be divided into multiple partitions. Each partition, called a virtual machine (VM), can have an independent operating system running different software applications isolated from other VMs running on the same physical machine. These technologies are being increasingly used to improve software systems, including database systems, and lower their total cost of ownership [1].

System performance has become increasingly important as computer systems get larger and more complex. The Problems of Performance is result of contention for, or exhaustion of, some system resource. When a system resource is exhausted, the system is unable to scale to higher levels of Performance the Oracle Database, is design to prevent system resources from becoming exhausted, that causing downtime[2].

Computer system is generally categorized into four Areas: CPU, memory, I/O subsystem, and network. In an Oracle database, the workload is measurement by “rates” of work being performed. Service time is not limited to just CPU time, disk I/O is also a form of service time. Any resource can provide service time and since resources have limited capacity they can also be Sources of queue time [3].

I/O (input/output) any operation program that transfers data to or from a computer. There are some devices are basically input-only devices (keyboards and mouse) others are output-only devices (printers) and others provide both input and output of data (hard disks, diskettes, writable CD-ROMs) [4].

The I/O devices are a very important part of the operating system .Using standard interfaces for a wide range of devices, making it easier to add newly developed devices to existing systems. Device drivers are developed modules that can be installed into an OS to handle a particular device or category of similar devices [5].

Redo logs files are a logs of history of all changes made to the database. Each redo log file consists of redo records. A redo record describes or represents a change made to a single block in the database [6].

It is crucial structure for recovery operations is the redo log. it consists of two or more files that store all changes made to the database as they occur. Every instance has an associated redo log to protect the database in case of an instance failure [7].

Literature Review

1- Automatic Virtual Machine Configuration for Database Workloads

Automatic Virtual Machine Configuration for Database Workloads Virtual machine monitors are tools for deployment database management systems and enterprise software applications. In which several database management system instances running in a virtual machine, sharing a common pool of physical computing resources. To optimizing the performance of these database management systems by controlling the configurations of the virtual machines in which they run. These virtual machine configurations determine how the shared physical resources will be allocated to the different database instances.

We introduce a virtualization design advisor that uses information about the anticipated workloads of each of the database systems to recommend workload-specific configurations offline.

We have implemented our approach using both PostgreSQL and DB2, and we have experimentally evaluated its effectiveness using DSS and OLTP workloads.

Problem Detention

N virtual machines, each running an independent DBMS, are competing for a pool of physical resources. Each DBMS, we are given a workload description consisting of a set of SQL statements. The problem is to allocate a share, or fraction, of each physical resource to each of the virtual machines. Each workload has an associated cost, which depends on the resources allocated to the virtual machine in which the workload runs.

Conclusions

Configuring multiple virtual machines that are all running database systems and sharing a pool of physical resources. Our approach is solved by implemented a virtualization design advisor that takes information about the different database workloads and uses this information to determine how to split the available physical computing.

The advisor depends on the cost models of the database system query optimizers to enable it to estimate workload performance under different resource allocations. We also presented performance measurements to refine the cost models Introducing a correcting cost model for evaluating the virtualization design advisor, demonstrating its accuracy and effectiveness.

2- Performance Evaluation of Oracle VM Server Virtualization Software

Software Server resources often are underutilized. It is easy to accommodate the need for additional server resources by adding more hardware. But in today's cost conscious landscape, with an emphasis on energy efficiency and managing power consumption, server virtualization emerges as a viable solution to maximize server utilization, achieve efficiency, provide scalability, while also being easy to deploy.

Tolly Group engineers examined the performance of virtual servers running on Oracle VM compared to typical bare-metal servers that consume more real estate and cost more to deploy. Also repeat the same tests on two virtual machines to prove the performance scalability of the Oracle VM virtualization technology.

Conclusions

Single virtual machine was able to produce 92.5% of the transaction throughput achieved by a bare-metal server running Oracle Database 11g in a 32-bit Linux environment with 30 active users. Running two instances of Oracle VM on a single server, it show combined throughput of 164.9% compared to the bare-metal server throughput. With 50 users, the single Oracle VM server achieved 93.6% of the transactions delivered by the bare-metal server. And with two instances of Oracle VM the throughput was 160.9% compared to the bare-metal server throughput. Oracle VM software can deliver comparable performance to a bare-metal server without the expense of deploying and supporting extra hardware. The high performance offered by Oracle VM is due to running Oracle Enterprise Linux in a virtualized mode. Since Linux is aware

that it is virtualized, it interacts with the hypervisor, this is the officially recommended and fully supported way of Oracle VM deployment.

3- A Performance Evaluation of Database Systems on Virtual Machines

Virtual machine technologies offer simple and practical to address many manageability problems in database systems. These technologies allow for server consolidation, easier deployment, and more flexible provisioning. Database systems are increasingly being run on virtual machines. However, it is also important to understand the cost of virtualization. Virtual machine technologies add a layer between applications and the hardware that they use (e.g. CPU, memory, and disk). This added complexity results in a performance overhead for software systems running in a virtual machine. We present an experimental study of the overhead of running a database workload in a virtual machine. Using a TPC-H workload running on PostgreSQL in a Xen virtual machine environment, these overheads do not translate to a high overhead in query execution time. Conclude that the advantages of running a database system in a virtual machine do not come at a high cost in performance.

Problem Detention

We use PostgreSQL and Xen virtualization. To answer the questions:

1. How much performance degradation will a database system experience when Moving from setting (Non-virtualized host) to setting (Xen-virtualized host)?
2. What are the possible causes of this overhead?

To measure the performance gap between a database system running on a non-virtualized physical machine and on a virtual machine given that the two systems are identically configured, both in software and hardware. The two systems should run the same operating system, with equal amounts of CPU power, memory, etc., available.

Experimental Results

1- Warm Experiments

The file systems and database cache. To eliminate the impact of the complex and rather expensive disk I/O factor after the first run the data required by all queries is present in the main memory, obviating the need for disk I/O in subsequent runs.

2- Cold Experiments

Clearing the Postgres and Linux caches in this manner ensures that all data required by a queries read from disk, a database system usually needs to retrieve large amounts of data (read data from the disk) or perform updates (write data back to disk).

Conclusions

The benefits of virtualization are coming with a very low cost. By understanding different virtualization technologies and how they interact with the hardware, operating system, applications, and the inter-virtual machine interactions, Moving in this direction increasingly to build better systems, this study explores various aspects of interaction between the Xen virtual machine monitor and the Postgres database system using the TPC-H database benchmark.

Problem Definition

The I/O (redo log file and disk I/O) is the major factor of the performance for oracle database engine with fixing the other factor (CPU, memory and network).on this research the I/O and redo log file will be measured on traditional server and virtual machine

Measuring the Database Performance

- 1- **Database Performance on Virtual Machine Disk I/O as following:**
 - a) **Load Profile:**

Table 1, Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	66.4	0.1	0.02	0.10
DB CPU(s):	3.1	0.0	0.00	0.00
Redo size (bytes):	548,150.8	2,239.2		

b) Instance Activity Stats:

Table 2, Instance Activity Stats

Statistic	Total	per Second	per Trans
physical read total IO requests	202,518	142.89	0.29
physical read total bytes	3,313,290,240	2,337,754.60	4,784.85
physical write total IO requests	274,702	193.82	0.40
physical write total bytes	4,215,892,480	2,974,602.68	6,088.34

2- Database Performance on Traditional server Disk I/O as following:

a) Load Profile:

Table 3, Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	2.0	0.0	0.00	0.01
DB CPU(s):	0.5	0.0	0.00	0.00
Redo size (bytes):	1,079,554.7	2,209.6		

b) Instance Activity Stats:

Table 4, Instance Activity Stats

Statistic	Total	per Second	per Trans
physical read total IO requests	4,471	3.24	0.01
physical read total bytes	65,732,096	47,601.01	194.45
physical write total IO requests	305,546	221.27	0.90
physical write total bytes	2,189,511,680	1,585,572.04	6,477.00

Comparison between performance of the database on virtual machine vs. traditional server

1. Disk I/O

- **Instance Activity**

Table 5, comparison between physical total request and physical total bytes

Function Name	Virtual machine			traditional server		
	Total	Per sec	Per trans	Total	Per sec	Per trans
physical read total IO requests	202,518	142.89	0.29	4,471	3.24	0.01
physical read total bytes	3,313,290,240	2,337,754.60	4,784.85	65,732,096	47,601.01	194.45
physical write total IO requests	274,702	193.82	0.40	305,546	221.27	0.90
physical write total bytes	4,215,892,480	2,974,602.68	6,088.34	2,189,511,680	1,585,572.04	6,477.00

- For Virtual Machine

IOP/S= physical read total IO requests + physical write total IO request
 $142.89+193.82=336.71$ requests

MP/S=physical read total bytes + physical write total bytes

$2,337,754.60+2,974,602.68=5312357.28$ byte ~ 5.066M

- For Traditional Server

IOP/S= physical read total IO requests + physical write total IO requests

$3.24+221.27=224.51$ requests

MP/S=physical read total bytes + physical write total bytes

$47,601.01+1,585,572.04=1633173.05$ byte ~1.55M

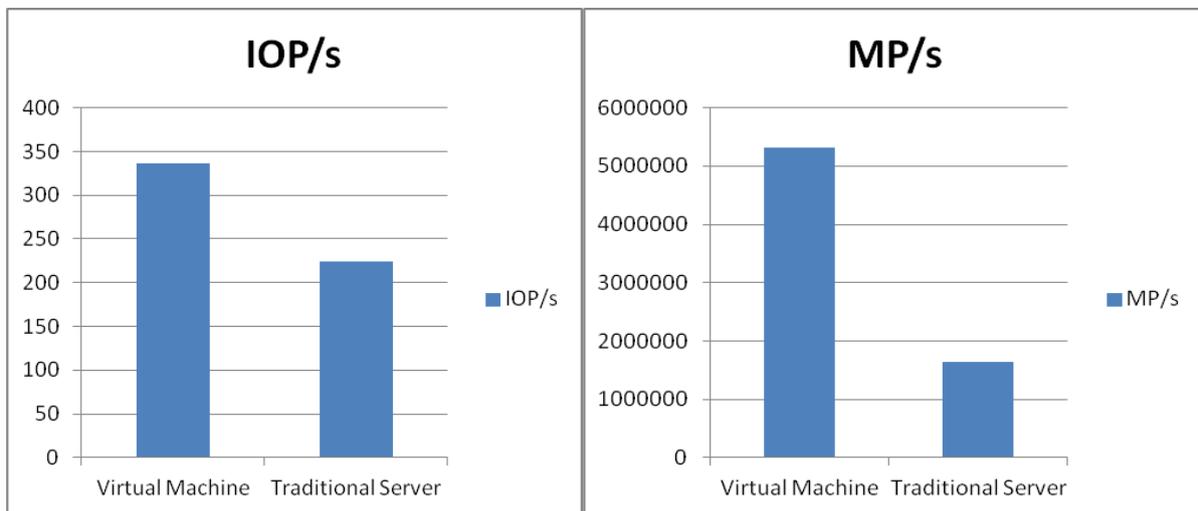


Fig. 1 Summary of the comparison between physical total request and physical total bytes

2. Redo log file

The LGWR is active background process doing physical I/O ,according to the table of (Load Profile and IO Stat by Function / File type summary) and duo to the high rate of redo traffic generated per second in the virtual machine was 548,150.8 byte (.6M Data per sec) and in traditional server was 1,079,554.7byte (1.106M Data per sec).From this result we found that the VM is less than TS in physical I/O.

Conclusion

In this work we measure the differences between oracle database performance on virtual machine and traditional server in I/O and redo log file with fixing the other factor CPU, memory, network, database workload time and number of concert user (1000).

And we found that:

Number of IOP/s in virtual machine is 336.71 requests (~ 5.066M per sec) and in traditional server are 224.51 requests (~1.55M per sec).and Redo log file in the virtual machine was 548,150.8 byte (.6M Data per sec) and in traditional server was 1,079,554.7 byte (1.106M Data per sec).From this results the virtual machine performance is higher than the traditional server with 69% in I/O. according to the high in I/O it generate redo log file in the virtual machine generate redo log file with 45% than the traditional server. Based on these results the virtual machine technology in all system related to the database is better performance in oracle database rather than traditional system.

References

- [1] David Marshall, Stephen S. Beaver, Jason W. Mc Carty, VMware ESX Essentials in the Virtual Data Center, August 25,2008, pp: 3, 4,5,6,7
- [2] Oracle Database Performance Tuning Guide 10g Release 1 (10.1), December 2003 pp: 46
- [3] "Virtual-machine," VMware Corp, [Online]. Available: <http://www.vmware.com/virtualization/virtual-machine.html>. [Accessed 15-9-2012].
- [4] Margaret Rouse definition input/output (I/O), [Online]. Available: <http://whatis.techtarget.com/definition/input-output-I-O>
- [5] Abraham Silberschatz, Greg Gagne, and Peter Baer Galvin, "Operating System Concepts, Eighth Edition ", Chapter 13 [Online]. Available: http://www.cs.uic.edu/~jbell/CourseNotes/OperatingSystems/13_IOSystems.html
- [6] Redo log [Online]. Available: http://en.wikipedia.org/wiki/Redo_log
- [7] Steve Fogel Oracle® Database Administrator's Guide 10g Release 2 (10.2), May 2006 pp: 876