

PERFORMANCE STUDY

Scaling IBM DB2 9 in a VMware Infrastructure 3 Environment



IBM **DB2** Information Management
Software

Contents

- Executive Summary..... 1**
- Introduction..... 1**
- Test Configuration and Methodology2**
 - Hardware configuration and setup2
 - Storage Configuration3
 - Software Configuration.....3
 - Test Methodology4
- Test Results5**
 - Results with One-VCPU Virtual Machines5
 - Results with Two-VCPU Virtual Machines6
 - Comparing One-VCPU and Two-VCPU Virtual Machines8
 - Comparison to Native Environment..... 10
- Conclusion 11**

Executive Summary

AMD, IBM, and VMware recently conducted a performance study on an IBM System x3755 using Second-Generation AMD Opteron™ processors and VMware® ESX Server 3.0.1 software. The performance study simulated a Web-based transactional workload using IBM DB2 9. It is clear from the empirical results in this study that running VMware virtualized environments on an IBM System x server based on the AMD Opteron processor and using the IBM DB2 9 database, IT organizations can:

- Leverage the benefits of server virtualization to achieve near-native database performance with minimal overhead.
- Gain throughput scalability by running multiple uniprocessor and SMP virtual machines on a single physical server.
- Exploit the NUMA memory architecture contained in the AMD Opteron processor to improve DB2 performance in a virtualized environment. .

Introduction

AMD, IBM, and VMware continue to collaborate on leveraging IBM System x servers, AMD Opteron processors and VMware Infrastructure 3 to offer a solid, robust virtualization platform to small and medium-sized enterprises. Running VMware virtual machines on IBM System x servers based on AMD Opteron processors and using IBM DB2 9 to host Web-based database application workloads can provide dramatic price/performance ratios thanks to the flexibility of virtual platforms, which are designed to enhance server consolidation and disaster recovery.

More than three years ago, DB2 was a featured application for the launch of the IBM eServer 325 — the first AMD Opteron processor-based server from a tier one vendor — and was one of the first database engines to support the AMD64 platform. This relationship continues to grow with new AMD processors that provide plenty of muscle while saving on power consumption and cooling costs. The Second-Generation AMD Opteron processors with Direct Connect Architecture introduce several new features, including quad-core upgradeability, AMD Virtualization™ (AMD-V™), and energy-efficient DDR2 memory. In addition, AMD Opteron processors are designed to advance AMD's performance-per-watt capabilities and leverage the proven technologies introduced in 2003 with first-generation AMD Opteron processors.

IBM has optimized DB2 9 for the AMD architecture, especially the dual-core AMD Opteron processor. DB2 recognizes the configuration of an AMD Opteron processor and transparently adjusts its behavior to take advantage of the NUMA architecture features, such as location of the memory controller. In DB2 9, internal data structures are organized so processes and memory are affinity-tuned. In other words, with memory affinity, DB2 9 places memory close to the process that is going to be using it. Furthermore, with the new Self Tuning Memory Manager feature in DB2 9, DB2 can dynamically react to changes in such underlying system resources as CPU and RAM to deliver optimal performance.

VMware Infrastructure is the most widely deployed, production-ready software suite available on commodity hardware for optimizing and managing IT environments through virtualization. Using virtualization-based distributed infrastructure services such as VMotion, VMware High Availability, and VMware Distributed Resource Scheduler, VMware Infrastructure 3 can provide transformative cost savings to the dynamic data center while increasing serviceability, efficiency, and reliability. VMware ESX Server provides a built-in dynamic, self-optimizing NUMA load balancer that fully exploits the advanced scalability features of servers using processors that provide NUMA support.

VMware Infrastructure 3 also supports 64-bit virtual machines as well as SMP virtual machines to deploy mission-critical enterprise applications.

We conducted a performance study using an AMD Opteron processor-based IBM System x3755 system running VMware ESX Server 3.0.1 software to simulate a Web-based transactional workload using IBM DB2 9. On this system, multiple virtual machines were created and configured to support a Web-based transaction-processing environment. The number of virtual machines running on the VMware ESX Server host was gradually increased, maintaining constant workloads on each virtual machine, with the goal of achieving the highest possible throughput score. By running a number of virtual machines simultaneously, we demonstrated how DB2 9 and VMware ESX Server running on a single AMD Opteron processor-based IBM System x3755 host could support increasing amounts of work while still effectively utilizing the underlying physical server and storage resources. This white paper describes the results obtained from these performance tests, which can help systems engineers and IT architects determine effective workload throughput and virtual machine utilization for their datacenters.

Test Configuration and Methodology

Hardware configuration and setup

The IBM System x3755 host used for testing is a 4-socket dual-core system with 2.8 GHz, AMD Opteron processors for Socket F (1207) and 16GB of RAM. VMware recommends using a dedicated Ethernet adapter for the ESX Server service console and a separate Ethernet adapter for virtual machines. For these tests, the workload was designed such that the network traffic had minimal impact on the scalability of the DB2 software. Therefore, we configured only one of the two onboard Ethernet adapters for use with both the ESX Server service console and the virtual machines. We used one PCI slot for the QLogic Fibre Channel host bus adapter (HBA) to provide connectivity to the Network Appliance storage area network (SAN). Table 1 summarizes the configuration of the System x3755 host tested.

Component	Details
CPU	4 AMD Opteron 8000 Series processors
Memory	16GB
Internal Disks	2x 300GB SCSI
Ethernet adapters	2x 10/100/1,000 Mbps internal adapters
Fibre Channel HBA	QLogic QLA2340 HBA

Table 1 — Hardware configuration used in test environment

Storage Configuration

We connected a single Network Appliance FAS960c storage server to the IBM System x3755 host using Fibre Channel Protocol. We attached 28 133GB 10K rpm disks to the storage server. We defined a single aggregate on which we created a RAID-4 FlexVol for each virtual machine's database tables and logs. We used only 10 percent of the SAN storage during this testing. Table 2 summarizes the SAN configuration we used in the test environment.

Component	Details
SAN	Network Appliance FAS960c SAN
Disks	28x 133GB 10,000 rpm
LUNs	RAID-4 LUNs

Table 2 — Storage configuration used in test environment

Software Configuration

We installed ESX Server 3.0.1 on one of the internal hard disks of the IBM System x3755 host used in testing. We used the out-of-the-box ESX Server configuration with no special updates or tuning.

After installing ESX Server, we created one Red Hat® Enterprise Linux 4 Update 3 virtual machine using VMware VirtualCenter 1.2. We configured the virtual machine with one virtual CPU, 1.6GB of RAM, one virtual disk, and one virtual Ethernet adapter. The virtual disk used for the operating system and the application environment was located on the internal SCSI hard disks of the IBM System x3755 host. We added a second virtual disk to the virtual machine by mapping a Network Appliance SAN LUN to that virtual machine. This disk housed both DB2 data and transaction log files for the virtual machine.

We used Zend Core for IBM 1.5.0 to implement the PHP-based database application. Zend Core for IBM delivers a rapid development and production PHP foundation for applications using PHP with IBM data servers. The product includes tight integration with the entire IBM data server family including DB2 9. The package includes Apache 2, DB2 Express, DB2 Runtime Client, Cloudscape, PHP 5, and PHP extensions such as `ibm_db2`.

We used default values for all configuration parameters in this study. Table 3 lists the software versions used in the test environment.

Software	Version
VMware ESX Server	3.0.1
VMware VirtualCenter	1.2
RedHat Enterprise Linux	RHEL4 Update 3 ¹
IBM DB2	9.1
Zend Core for IBM	1.5.0

Figure 3 — Software configuration used in test environment

¹ DB2 9.1 and VMware ESX Server 3.0.1 are also validated and supported on RHEL 4 Update 4 and SLES10

We installed VMware Tools in the guest operating system and set up eight identical uniprocessor virtual machines. We later changed these virtual machines were to use two or four virtual CPUs in SMP virtual machines as needed. We powered up and ran only the virtual machines used during a particular test; we made sure all unused virtual machines on the ESX Server host were shut down.

Test Methodology

The workload we used for demonstration in this paper simulates an online transaction processing environment using the DVDStore application. A driver simulates users logging in and browsing a catalog of available products via simple search queries. Returning customers see their online purchase transactions history, while new users may register to create customer accounts. Users may select items they would like to purchase and proceed to check out or continue to view available products. Restocking is also done once a minute to replenish the product quantities. Figure 1 shows the average percentage of execution for each transaction. To download the version of the DVDStore application including the database scripts used for this study, go to <http://devzone.zend.com/content/ds2.tar.gz>.

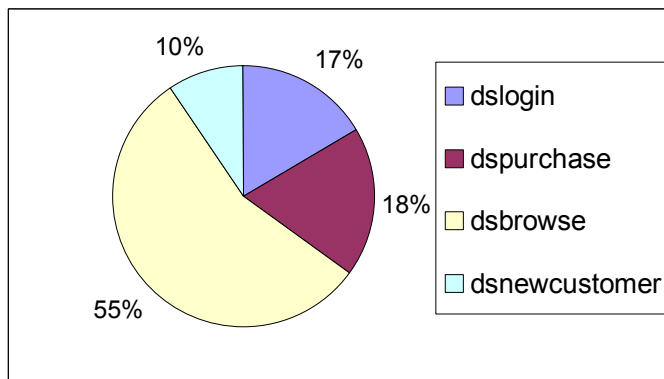


Figure 1 — Workload transaction distribution

We created a 20GB database representing 20 million customers, one million orders, and 100,000 products. The database utilizes many DB2 features such as automatic storage, stored procedures, triggers, and referential integrity. The tuning of the database memory heaps is not predefined but is autonomically determined using the dynamic Self Tuning Memory Manager available in DB2 9.

The entire system configuration is a three-tier model. Tier 1 executes the driver that emulates the activities of Web users. Tier 2 comprises Web application servers that intercept the requests and send database transactions to a DB2 9 data server. We configured all three tiers to run within the same virtual machine instance.

The simulated workload ran 10 users for each virtual CPU used in the test. We powered on additional virtual machines to scale the workload. In other words, we increased the number of simulated users in increments of 10 users, starting with 10 users for one virtual CPU and increasing to 80 users for eight virtual CPUs. To simulate a real-world scenario, we set the think time for a default of 0.25 seconds and configured the percentage of new users at 10 percent. Each test ran for 10 minutes.

The parameter we used to measure the performance and scalability of ESX Server virtual machines was operations per minute (OPMs) reported by the application. OPMs can also be thought of as transactions per minute. We measured CPU utilization and other system metrics on the ESX Server host using esxtop.

Test Results

Results with One-VCPU Virtual Machines

The benefits of running IBM DB2 on VMware ESX Server become readily apparent when we ran the test with multiple virtual machines. Figure 2 shows the throughput scalability we achieved by successively increasing the number of virtual machines from one to eight. The throughput demonstrates near-linear scalability as virtual machines are added.

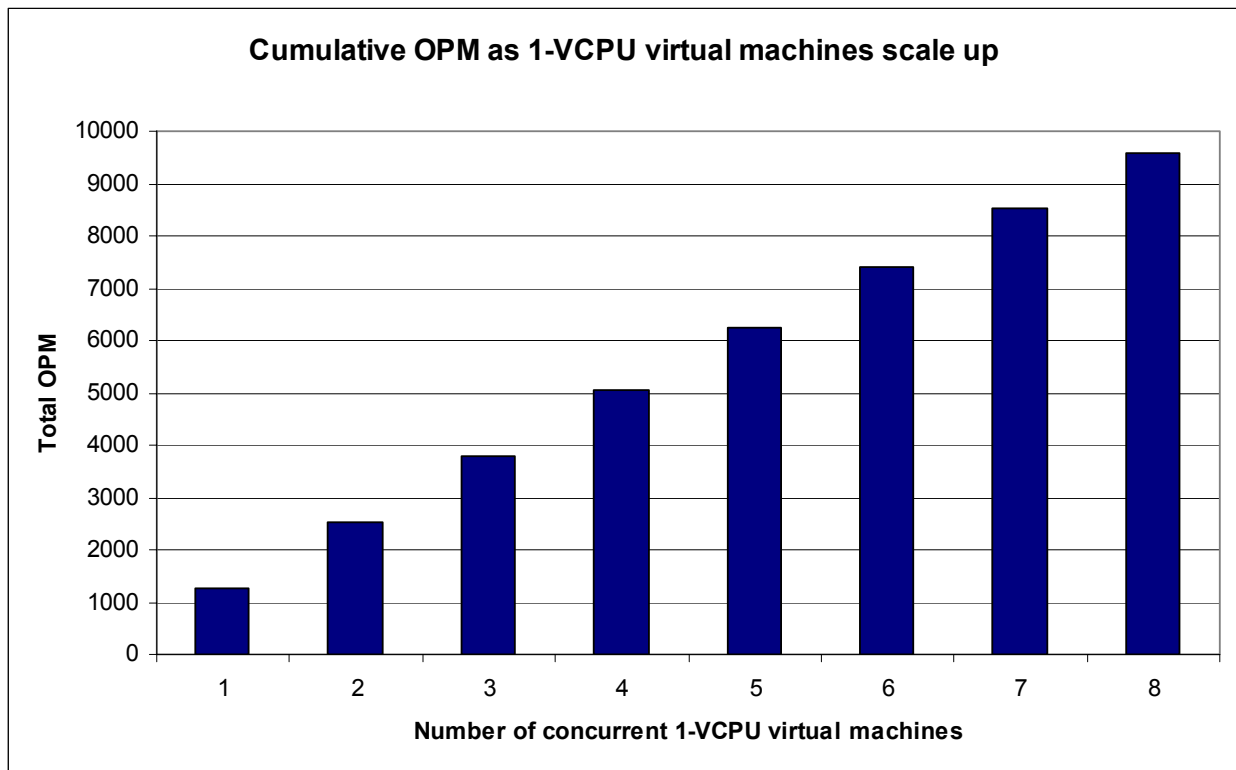


Figure 2— Scalability of one-VCPU virtual machines

Another way to establish VMware ESX Server scalability is to look at the throughput that each virtual machine sustained as the number of virtual machines scaled up. Figure 3 shows the average throughput per virtual machine as the number of virtual machines and number of users increased. The average throughput per virtual machine decreased by about 6 percent when the number of virtual machines increased from one to eight. This decrease in average throughput per individual virtual machine as the number of virtual machines increased is expected because the ESX Server host had to sustain more work with a finite pool of system resources.

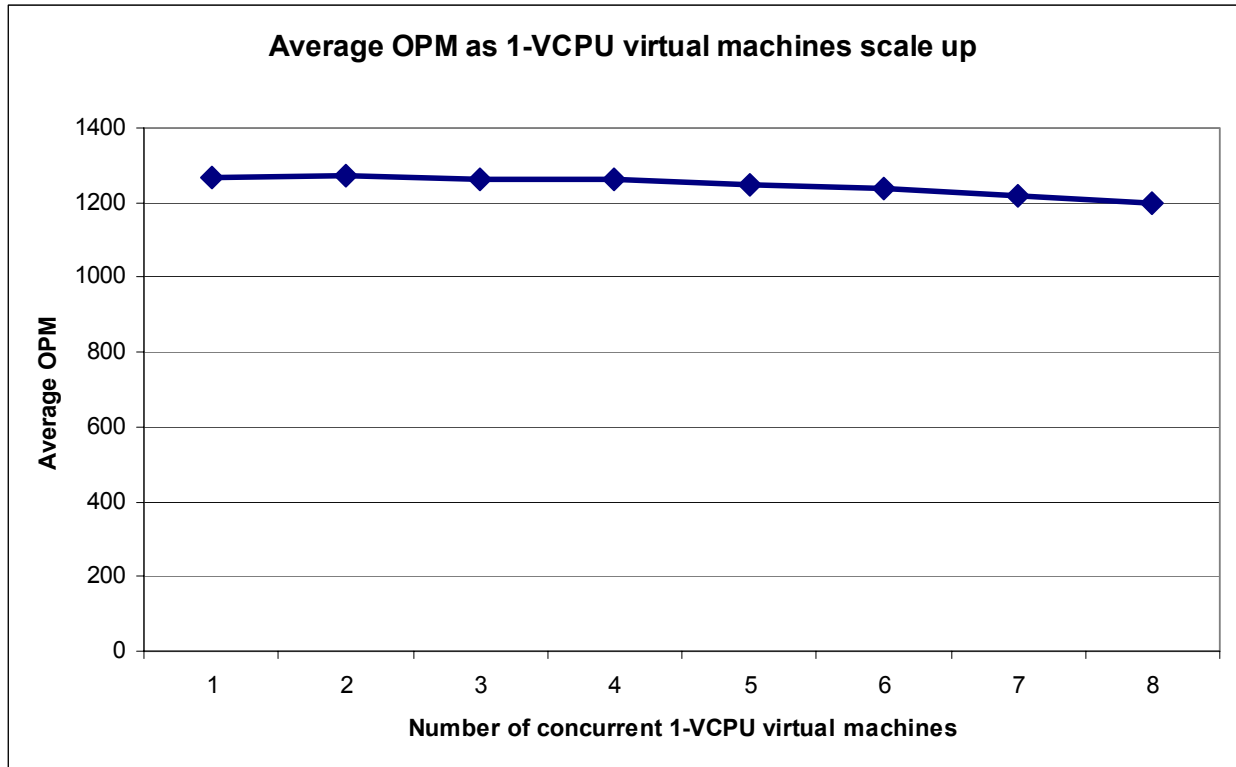


Figure 3 — Average throughput per one-VCPU virtual machine

These two methods — performance per ESX Server host and average performance per virtual machine — used a throughput metric to help determine efficient ESX Server system utilization and to establish the quantifiable metrics used to determine the number of supported users in production environment. The data presented in this figure represents average values. With any workload, the load level will vary during the normal course of a business day. So system administrators should allow for spikes in resource utilization when determining how many virtual machines to run on an ESX Server host.

Results with Two-VCPU Virtual Machines

This section shows the results obtained when running two-VCPU SMP virtual machines.

Figure 4 shows scalability achieved by successively increasing the number of two-VCPU SMP virtual machines from one to seven. The throughput demonstrated near-linear scalability up to four virtual machines. Beyond four virtual machines, the ESX Server host ran in an overcommitted CPU scenario on the four-socket dual-core System x3755 server used in this testing. Consequently, it is no surprise that the total throughput gradually decreased as the number of virtual machines increased beyond four.




Figure 4 — Scalability of two-VCPU SMP virtual machines

Figure 5 shows the decrease in the average throughput per virtual machine as the number of two-VCPU SMP virtual machines increased. The average throughput per virtual machine decreased by only about 6 percent until the number of virtual machines increased from one to four.




Figure 5 — Average throughput per two-VCPU SMP virtual machine

Comparing One-VCPU and Two-VCPU Virtual Machines

During these tests, we simulated 10 users for each virtual CPU used in the test run. In other words, for the same number of concurrent virtual machines used, we doubled the number of simulated users for two-VCPU SMP virtual machines as compared to one-VCPU uniprocessor virtual machines. Figure 6 illustrates that the throughput was almost double in two-VCPU SMP virtual machines until the System x3755 host became overcommitted. These results validate that virtual environments can achieve SMP scalability similar to that seen in native environments.




Figure 6 — Scalability comparison with constant number of virtual machines

Figure 7 compares the scalability of one-VCPU uniprocessor virtual machines to two-VCPU SMP virtual machines under constant workload (that is, with the same number of simulated users). The throughput rate for two-VCPU SMP virtual machines was slightly less compared to that for one-VCPU uniprocessor virtual machines. The results for the SMP case could have been improved if the virtual CPUs were pinned to specific physical CPUs.




Figure 7 — Scalability comparison with constant number of users

Comparison to Native Environment

Figure 8 compares the throughput performance of a four-core bare-metal configuration to various virtual machine configurations totaling four virtual CPUs. In the virtual environment, this test was run with four one-VCPU uniprocessor virtual machines, two two-VCPU SMP virtual machines and one four-VCPU virtual machine. The results show that you can achieve comparable performance in virtualized environments by scaling the number of virtual machines. For SMP-enabled virtual machines, the virtual processors from the same virtual machine were co-scheduled. Co-scheduling led to the problem of processor fragmentation, which was increasingly noticeable on SMP virtual machines with higher VCPU counts. This resulted in reduced throughput rates and thus higher virtualization overhead for SMP virtual machines as the number of virtual CPUs increased.




Figure 8 — Physical to virtual scalability comparison

Conclusion

The series of tests and corresponding analysis we describe in this paper supports the conclusion that running VMware virtual machines on IBM System x servers based on AMD Opteron processors and using IBM DB2 9 can provide an effective production-ready platform for hosting multiple virtualized transaction processing workloads.

The performance metrics (throughput per virtual machine and aggregate ESX Server host throughput) we present in this white paper can help IT organizations size and design virtualized environments to meet their SLA targets. The number of actual users and transactions that can be supported in production environments will, of course, depend on the specific applications used and the intensity of user activity. The results clearly demonstrate that in a heavily virtualized environment, good throughput was retained even as the number of total supported users and the number of virtual machines increased.

For more information on AMD products, go to <http://www.amd.com/opteron>.

For more information on AMD-V™ virtualization technology, go to <http://www.amd.com/virtualization>

For extensive DB2 performance optimization on AMD Opteron processors, go to <http://www.devx.com/amd/Article/32087>

For more information on IBM System x3755, go to <http://www-03.ibm.com/systems/x/rack/x3755/>.

For more information on DB2 and to download a trial version of DB2 9, go to <http://www-306.ibm.com/software/data/db2/9>.

For more information on virtualization with VMware Infrastructure 3, go to <http://www.vmware.com>.

Revision: 20070329 Item: PS-008-ISV-01-03



VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com
© 2007 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806, 6,944,699, 7,069,413; 7,082,598 and 7,089,377; patents pending.
VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.



IBM **DB2** Information Management
Software