

Dell PowerEdge M820 Virtualization Consolidation Study

Consolidating 10th and 11th generation PowerEdge 2U rack servers into 12th generation PowerEdge M820 blade servers

Louis Barton
Solutions Performance Analysis (SPA)



This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2012 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, EqualLogic, PowerConnect, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, SQL Server, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. VMware, ESX, ESXi, vCenter, and vSphere are either trademarks or registered trademarks of VMware Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

November 2012 | Rev 1.0

Contents

Executive summary	4
Introduction	5
Test description	6
Sample environment characterization	6
Virtualized SQL system configurations	8
Test methodology	9
Overview	9
Data collection overview	9
Test results	10
Virtual machine scaling to determine optimal load for each system	10
Performance: Aggregate throughput (OPM)	11
Conclusion	14
Appendix A: References	15
Appendix B: Additional configuration details	16

Tables

Table 1. DVD Store database size description	6
Table 2. System test configurations	8
Table 3. Supporting infrastructure configurations	8
Table 4. Detailed system test configurations	16
Table 5. BIOS settings	16
Table 6. Software configuration	16
Table 7. Shared storage configuration	17

Figures

Figure 1. DVD Store test environment	7
Figure 2. Aggregate throughput for PowerEdge M820, R710, and 2950 III	11
Figure 3. OPM scaling by number of VMs per SUT	12
Figure 4. Performance ratio in server consolidation format	13
Figure 5. Virtual Machine Capacity (10U rack space)	14

Executive summary

Dell's Solutions Performance Analysis (SPA) team conducted a study to measure the virtualized environment to compare the legacy Dell™ PowerEdge™ 2950 III and R710 servers to the new Dell PowerEdge M820 blade server. The configurations chosen were aimed at maximizing the performance of each system by equipping them with the top-speed processors available, a 10Gb network connection to a backend iSCSI SAN, and sufficient memory to run the workload at 100% processor utilization for all processor architectures tested.

The study indicates that the Intel® Xeon® E5-4600 processors offer significant performance gains. The Intel Xeon E5-4600 processors are also able to scale out further, supporting a higher number of virtual machines (VMs) at optimal performance.

Key findings of the study include:

- The Dell PowerEdge M820 with Intel Xeon E5-4650 processors provided 266% higher virtualized online transaction processing (OLTP) performance than the PowerEdge 2950 III with Intel Xeon X5460 processors, and 90% higher virtualized OLTP performance than the PowerEdge R710 with Intel Xeon X5690 processors. The M820 configuration with Intel Xeon E5-4650 processors obtained 664,640 orders per minute (OPM) in a virtualized environment, as compared to the R710 configuration, which scored 350,351 OPM; and the 2950 III configuration, which scored 181,477 OPM.
- Optimal VM performance is achieved on a single PowerEdge M820 running on 30 VMs. PowerEdge 2950 III optimal VM performance is achieved running on 4 VMs and PowerEdge R710 running on 10 VMs. That is a 650% increase in optimal performing M820 VMs over 2950 III VMs and a 200% increase in optimal performing VMs over R710. A full chassis of M820 blade servers can support the equivalent of 60 x 2U 2950 III rack servers and 24 x 2U R710 rack servers running at peak performance levels.
- In the same 10U rack space, a full chassis of PowerEdge M820 blade servers can sustain 240 optimally performing VMs, PowerEdge R710 can sustain 50 VMs, and PowerEdge 2950 III can sustain 20 VMs. That is a 380% increase in optimal VM count over the R710 and an 1100% increase over the 2950 III.

Introduction

In order to provide a better understanding for virtualization performance, this study describes specific application performance in a virtualized environment. The servers used in this study are the Dell PowerEdge M820 with Intel Xeon E5-4600 series processors, the previous generation Dell PowerEdge R710 with Intel Xeon X5600 series processors, and the Dell PowerEdge 2950 III 10th generation rack servers with Intel Xeon 5400 series processors. The goal of this study is to show the benefits of consolidating a number of rack servers down to a few blade servers using the Dell PowerEdge M820 blade server. Because of the four-socket configuration and increased memory capacity, the performance of many previous and current generation servers can be consolidated into a single Dell PowerEdge M820 blade server.

Considerations include:

- Comparing Microsoft® SQL Server™ virtualized performance between the Dell PowerEdge M820 with Intel Xeon E5-4600 processors against the previous-generation Dell PowerEdge R710 equipped with Intel Xeon X5600 processors, and the Dell PowerEdge 2950 III 10th generation server with Intel Xeon 5400 series processors.
- The optimal number of VMs for each server is determined by running the DVD Store Version 2 (DS2) workload at increasing VM counts until the system reaches a saturation point. It is at this point the aggregate OPM of all VMs peaks. Increasing the VM count after this point will show a decline in throughput as response times increase and as the system exceeds its saturation point.
- System configurations between the three platforms in this study differ in both the hypervisor version and database management systems. The PowerEdge M820 blade server is running VMware® ESXi™ 5.0 and Microsoft SQL Server 2012. The PowerEdge R710 and 2950 III are running VMware ESX™ 4.1 update 2 and Microsoft SQL Server 2008 R2. The SPA team chose these system configurations to show a solution that offers a representative snapshot for each platform.

This study is not intended to be a comprehensive analysis of virtualization performance between all Dell PowerEdge platforms. Only the M820 blade server and the R710 and 2950 III rack mount servers were chosen to characterize the relative gains in performance in the current generation of Dell servers versus previous generations. In addition, because the database generated was using DS2 custom scripts, the numbers published are not directly comparable to other published DVD Store results by either Dell or its competitors.

Test description

DS2 is a complete online e-commerce test application with a backend database component, a Web application layer, and driver programs. The goal in designing the database component as well as the mid-tier application was to use many advanced database features (for example, transactions, stored procedures, triggers, and referential integrity) while keeping the database easy to install and understand. The DS2 workload may be used to test databases or as a stress tool for any purpose.

The distribution includes code for Microsoft SQL Server, Oracle®, and MySQL™ databases. Included in the release are data generation programs; shell scripts to build data for 10MB, 1GB, and 100GB versions of DS2; database build scripts and stored procedures; PHP Web pages; and a C# driver program.¹

Database workload sizes can be small, medium, or large depending on environment setup. See Table 1 for details.

Table 1. DVD Store database size description

Database	Orders	Size	Products	Customers
Small	10MB	20,000	1,000/month	10,000
Medium	1GB	2,000,000	100,000/month	100,000
Large	100GB	200,000,000	10,000,000/month	1,000,000

Sample environment characterization

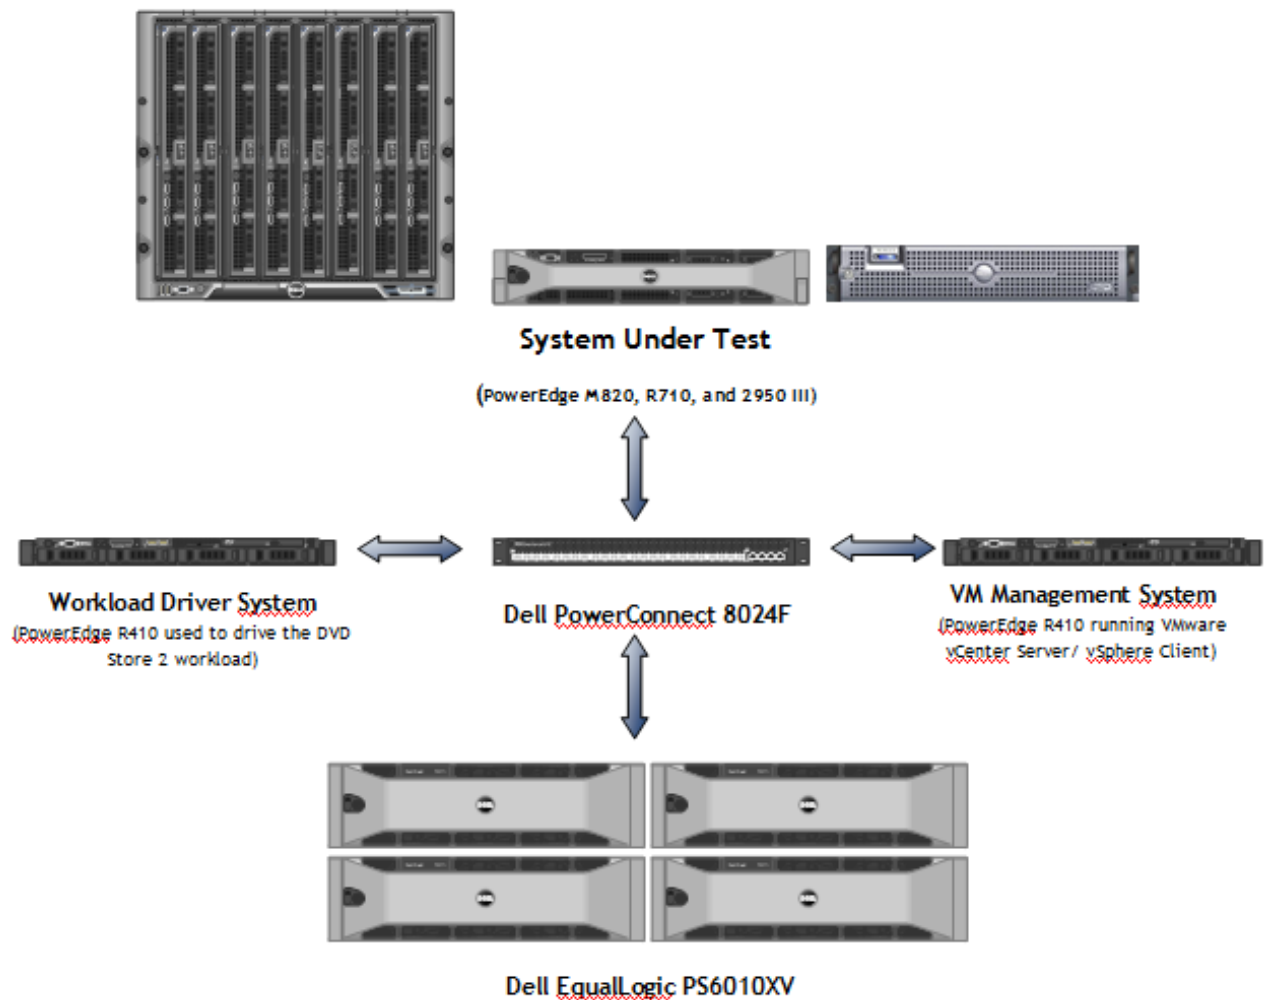
To measure consolidated VM infrastructure host performance in a simulated e-commerce scenario, a full end-to-end solution was implemented. This environment consists of a physical server running VMware ESXi 5.0 or VMware ESX 4.1 update 2. This is referred to as the system under test (SUT) in Figure 1. The PowerEdge M820 blade server is running ESXi 5.0; the PowerEdge R710 and PowerEdge 2950 III are running VMware ESX 4.1 update 2.

In addition to the SUT, the environment contains a VM management system and a workload driver system. The environment controller is running Microsoft Windows Server® 2008 SP1. The intent of this system is to manage the VMware vCenter™ and vSphere™ client. The workload driver system is running Windows Server 2008 R2 Enterprise. This system is used to create simultaneous SQL sessions that will simulate a customer interacting with the DS2 Web application, which in turn exercises the backend SQL database on the VM.

¹ Test description information from Dell Enterprise Technology Center website:
<http://en.community.dell.com/techcenter/extras/w/wiki/dvd-store-readme.aspx>

In addition to these three physical systems, the environment contains a 10Gb Ethernet Dell PowerConnect™ 8024F switch, and four Dell EqualLogic™ PS6010XV SANs. The PowerConnect switch enables communications between the VMs and controller systems (workload driver and environment controller), and between the SUT and the iSCSI-based EqualLogic storage backend.

Figure 1. DVD Store test environment



Virtualized SQL system configurations

Configurations used in the test environment are shown in Table 2 and Table 3.

Table 2. System test configurations

System	PowerEdge 2950 III	PowerEdge R710	PowerEdge M820
Hypervisor	VMware ESX 4.1 update 2		VMware ESXi 5.0
Processor Model	2 x Intel Xeon X5460	2 x Intel Xeon X5690	4 x Intel Xeon E5-4650
Processor Frequency / Cache Size	3.16GHz, 12M L3	3.46GHz, 12M L3	2.7GHz, 20M L3
Processor TDP	120W	130W	130W
Physical Cores / Processor	4	6	8
Logical Processors	16	24	64
Memory Details	8 x 4GB DDR2 FBDIMMs	12 x 8GB PC3L 2Rx4 RDIMMs	32 x 8GB PC3L 2Rx4 RDIMMs
Memory Frequency	667MT/s	1333MT/s	1333MT/s

Table 3. Supporting infrastructure configurations

Environment controller system	DVD Store 2 workload driver system
PowerEdge R410 Windows Server 2008 R2 Enterprise 2 x Intel Xeon Processor E5620, 2.4GHz Memory: 8GB RAM Disk configuration: 2 x 146GB SAS (RAID 1)	PowerEdge R410 Windows Server 2008 R2 Enterprise 2 x Intel Xeon Processor E5620, 2.4GHz Memory: 16GB RAM Disk configuration: 2 x 146GB SAS (RAID 1)
Networking	Virtualization Shared Storage
PowerConnect 8024F running at 10GbE 1 network port used per EqualLogic SAN (4 total) 1 network port used for SUT 1 port used as management port (via vSphere) Jumbo frames enabled on switches and SANs	Dell EqualLogic PS6010XV 4 x Enclosures 7 x Volumes 16 x 600GB SAS 15k HDD per enclosure RAID 10

Test methodology

Overview

The following methodology was used in the testing procedure:

- Determined the optimal number of virtual machines that each SUT can support. Each SUT was tested with 2-40 VMs and the total OPM recorded. The number of VMs that provided the highest total OPM was the number chosen to test that particular server for the remainder of the study.
- Exercised all enabled VMs for a 20-minute run time.
- All database instances restored from the same restore file and handled identically throughout the benchmark run to promote an optimal and reliable test environment.
- Recorded performance metrics and SUT performance counters.

Note: The performance metric recorded is OPM (operations per minute), which is a measure of the number of orders processed during the benchmark runtime.

Data collection overview

The following steps were used in collecting data:

1. Built a 1GB SQL database on a VM using DS2 custom database scripts.
 - a. Generated the necessary database tables and a database build script in a Linux VM using the DS2 provided custom database PERL scripts.
 - b. Moved the files generated by the custom scripts from the Red Hat® Enterprise® Linux 6.1 VM to the Microsoft Windows Server 2008 R2 VMs.
 - c. Executed the database build script which populates the database with customer, order, and inventory data.
2. Backed up the generated SQL database to a restore file, allowing the database to be restored from a pristine state. This database backup was then distributed to all VMs in the test bed.
3. Deleted VM databases, and then restored from the backup restore file and rebooted.
4. Performed the official run starting all VMs simultaneously.

Note: The VMware ESXtop tool gathered system-related performance counters for profiling system utilization data during the recorded run (such as CPU, I/O, and memory).

5. Determined the number of VMs needed to reach a system saturation point.
 - a. This is determined by scaling the number of VMs from a minimum to a maximum VM count and recording the total OPM at each increment.
 - b. The score will peak as you increase the VM count.
 - c. Once it reaches the saturation point, the throughput scores will decline as the system exceeds its saturation limits.

6. Checked the official run data for stability; collecting data if stable, and restarting the process if high variation was present.

Data collected during the official run included:

- Average of the total ESXtop CPU counters during the official run (CPU percentage)
- VM throughput in OPM; the metric used in this study is an aggregate OPM measurement of all enabled VMs and is represented as performance in the test results

Test results

Virtual machine scaling to determine optimal load for each system

In comparing the 12th generation Dell PowerEdge M820 blade server with the 11th generation Dell PowerEdge R710 server and the 10th generation Dell PowerEdge 2950 III server, it was necessary to determine the optimal loading of VMs on each system due to the architectural differences in the processors. The M820 features the new Intel Xeon processor E5-4600 product family, the R710 comes equipped with the Intel Xeon 5600 series processors, and the PowerEdge 2950 III comes equipped with the Intel Xeon 5400 series processors. The E5-4650 processor has 8 physical cores with 20MB L3 cache. The X5690 processor has 6 physical processing cores with 12MB L3 cache and the X5460 has 4 cores with a 12MB L3 cache. The first test in this study was to determine how many VMs to assign each server for the remainder of the testing.

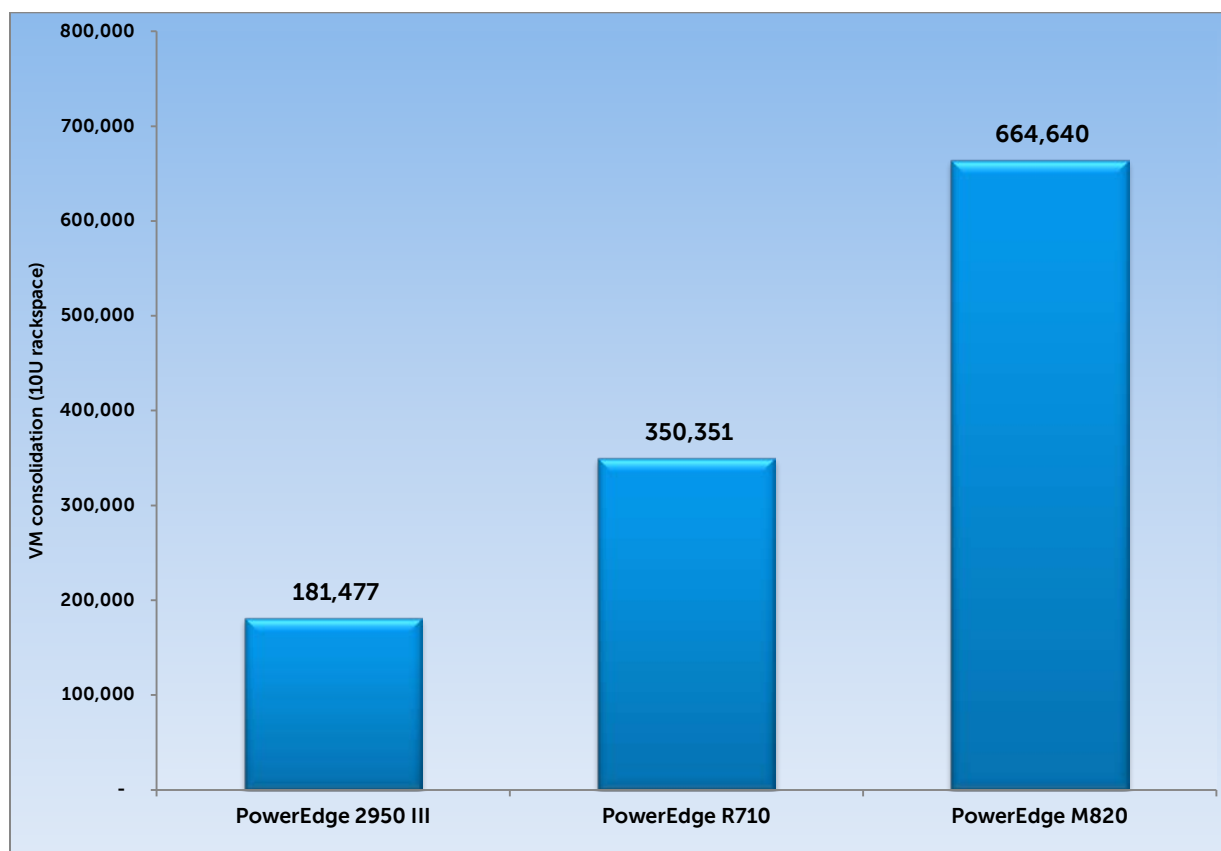
To allow each system to perform at its highest level for the scaling study, each system was set to Maximum Performance mode in the BIOS and configured with 2 DIMMs per memory channel. The Intel Xeon X5460 processor has 2 memory channels per processor, the Intel Xeon X5690 processor has 3 memory channels per processor, and the Intel Xeon E5-4650 has 4 memory channels per processor. For a quad-processor configuration, the M820 is configured with 32 DIMMs. The dual-socket R710 is configured with 12 DIMMs and the 2950 III is configured with 8 DIMMs.

Performance: Aggregate throughput (OPM)

The Dell PowerEdge M820 features the Intel Xeon processor E5-4600 product family, which has a greater core count than the previous generation PowerEdge R710 with the Intel Xeon 5600 series processor and double the amount of cores as R710's predecessor: the PowerEdge 2950 III with Xeon 5400 series processors. With additional processor capabilities, the M820 supports three times the number of VMs as the R710, and eight times the VM count of the 2950 III.

Figure 2 clearly shows this advantage, comparing the total OPM achieved by the M820, R710, and 2950 III. Running at the peak number of VMs for each system's processor architecture, the M820 was able to process 664,640 OPM compared to the R710, which was able to process 350,351 OPM; and the 2950 III, which generated a maximum throughput of 181,477 OPM. The PowerEdge M820 was able to handle 90% more operations than the PowerEdge R710, a substantial improvement over just one generation of servers. That delta is further expanded, which showed uplift of 266% over the PowerEdge 2950 III server.

Figure 2. Aggregate throughput for PowerEdge M820, R710, and 2950 III



Using the DS2 throughput metric (OPM) of all VMs on a given platform is one method of quantifying performance. This aggregate measurement can then be used to compare multiple platform performance. Additionally, the SPA team derived two other metrics from this single data point: server consolidation and VM capacity.

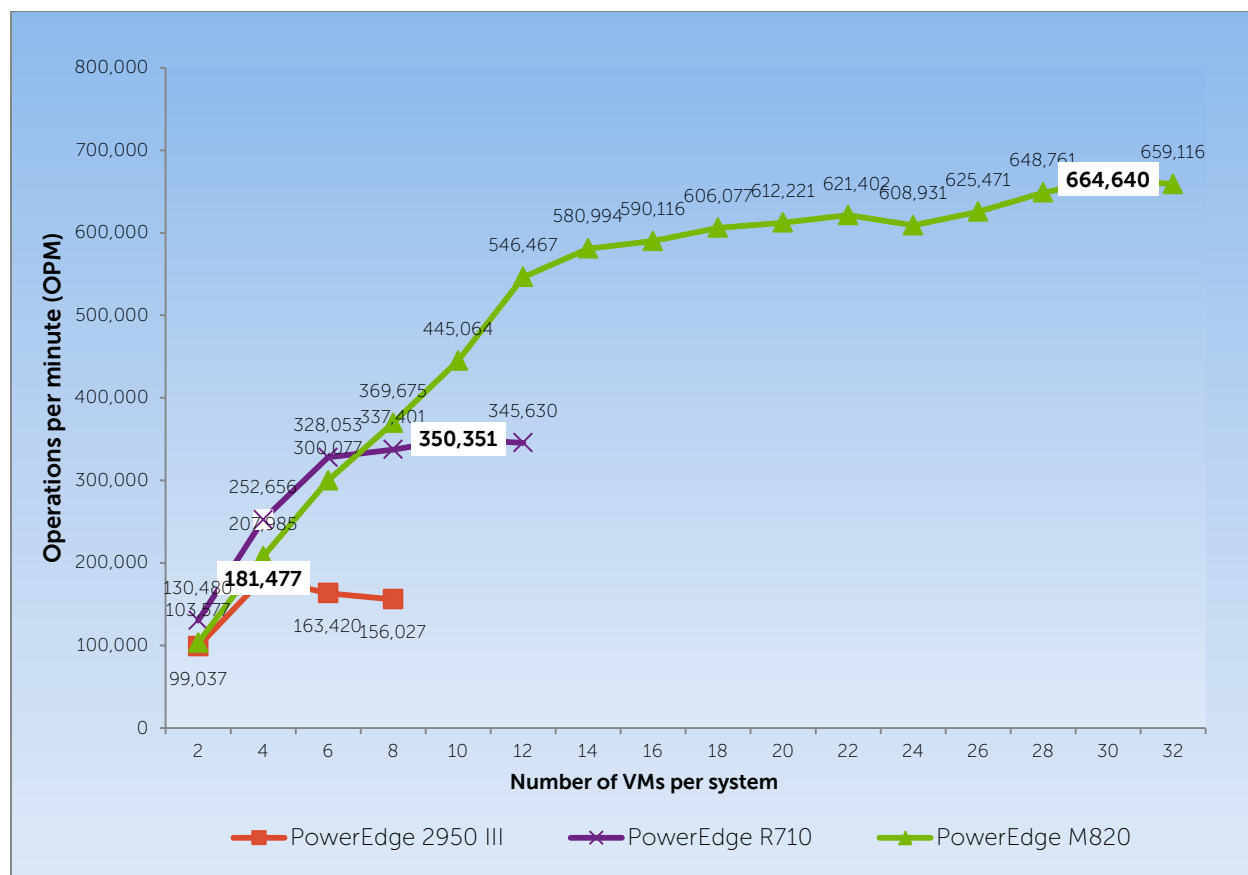
Server consolidation

The first metric is server consolidation. This measurement is the number of servers that can be consolidated from previous generation platforms to the current generation platform. This comparison demonstrates how many PowerEdge 2950 III and PowerEdge R710 servers can be consolidated into a fully populated blade chassis of PowerEdge M820 servers.

Looking closer at the data, the peak performance data points per platform tell a story beyond raw performance. At the VM count of those peak performance points, there is a statement on raw performance (OPM) and consolidation (VM count). The PowerEdge M820 blade server throughput peaks at 30 VMs running the DS2 workload. PowerEdge R710 peaks at 10 VMs and PowerEdge 2950 III peaks at 4 VMs. As shown in Figure 3, those peak performance levels can be translated into server consolidation metrics. This investigation determines how many previous generation servers can be consolidated in a fully populated chassis of PowerEdge M820 blade servers.


This first comparison in Figure 4 highlights a very important feature of the new Dell PowerEdge M820. With more high-performing processing cores, the PowerEdge M820 is more scalable and can support a greater number of heavily loaded VMs without sacrificing individual VM performance.

Figure 3. OPM scaling by number of VMs per SUT



The measurements show that a fully populated chassis of 8 PowerEdge M820 blade servers are equivalent to 24 PowerEdge R710 servers or 60 PowerEdge 2950 III servers. Figure 4 represents a consolidation ratio of nearly three fully loaded racks of PowerEdge 2950 III servers down to one blade chassis of PowerEdge M820 servers. Figure 4 also illustrates server consolidation transitioning from the previous generation 2U dual-socket rack servers to the current generation quad-socket blade servers.

Figure 4. Performance ratio in server consolidation format




VM Capacity

The optimal throughput performance data points can also be used to determine virtual machine capacity. This is a measurement of the maximum number of peak performing VMs per platform that extrapolates that data into a like-to-like rack space comparison. This study determines how many fully loaded VMs the PowerEdge R710 and PowerEdge 2950 III are supported in the same 10U rack space as a fully populated blade chassis of the PowerEdge M820.

Using the performance data previously stated in the preceding performance and consolidation sections of this paper, we can derive a VM capacity metric. As shown in Figure 5, the PowerEdge M820 server supports 240 peak performing VMs in a 10U rack space, the PowerEdge R710 server supports 50 VMs and the PowerEdge 2950 III server supports 20 VMs in the same 10U rack space.

Figure 5. Virtual Machine Capacity (10U rack space)



Conclusion

The goal of this study is to characterize the performance improvements of the new 12th generation Dell PowerEdge M820 blade server over the previous generation Dell PowerEdge R710 server with Intel Xeon 5600 series processors and two-generation-old Dell PowerEdge 2950 III with Intel Xeon 5400 series processors using a virtualized database workload. Actual customer environments are unlikely to be running continuously at 100% processor utilization levels, but the relative scaling achieved by the PowerEdge M820 in this study shows the merits of the Intel Xeon processor E5-4600 product family and the overall design of the Dell PowerEdge M820 blade server.

Based on the data contained in this study, the PowerEdge M820 with E5-4650 processors provides considerable performance, consolidation, and capacity advantages while reducing the overall data center footprint over the previous generations of rack mount servers. The PowerEdge M820 blade server with Intel Xeon E5-4650 processors is the quad-socket platform of choice for maximum performance and scalability with additional processing cores and higher capacity memory support.

Appendix A: References

The DVD Store Virtualized SQL test methodology and overall structure of this white paper are due in large part to previous studies, as well as technical manuals and documents. The following list contains the key sources of information this study drew upon.

- “Dell PowerEdge R910 SQL OLTP Virtualization Study: Measuring Performance and Power Improvements of New Intel Xeon E7 Processors and Low-Voltage Memory” by Waseem Raja and John Beckett, attachments.wetpaintserv.us/jD7st0SRn16bZhoLImoZoQ505396
- “Configuring VMware vSphere Software iSCSI With Dell EqualLogic PS Series Storage,” en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19861454.aspx
- “How do I configure the optimal switch settings for IP based SAN?” by Naveen Iyengar, en.community.dell.com/techcenter/enterprise-solutions/w/oracle_solutions/1422.aspx
- “Dell EqualLogic Configuration Guide: A guide to building an iSCSI based SAN solution with Dell EqualLogic PS Series Arrays,” en.community.dell.com/techcenter/storage/w/wiki/equallogic-configuration-guide.aspx
- “Dell PowerEdge R720 SQL OLTP Virtualization Study: Measuring Performance and Power Improvements of the New Intel Xeon E5-2600 Series Processors” by Don Hoffman, <http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/Dell-PowerEdge-R720-SQL-OLTP-Virtualization-Study.pdf>
- “Dell PowerEdge R820 Virtualization Consolidation Study” by Don Hoffman, Louis Barton, and John Barnhart, http://en.community.dell.com/techcenter/extras/m/white_papers/20106856.aspx

Appendix B: Additional configuration details

Table 4. Detailed system test configurations

System	PowerEdge 2950 III	PowerEdge R710	PowerEdge M820
Processor	2 x Intel Xeon X5460 (3.16 GHz, 4c) 12M L3, 120W	2 x Intel Xeon X5690 (3.46 GHz, 6c) 12M L3, 130W	4 x Intel Xeon E5-4650 (2.7 GHz, 8c) 20M L3, 130W
Memory	8 x 4GB 667MT/s DDR2 FBDIMMs - 32GB	12 x 8GB PC3L-1333 MT/s 2Rx4 RDIMMs (TJ1DY) - 96GB	32 x 8GB PC3L-1333MT/s 2Rx4 RDIMMs (TJ1DY) - 256GB
RAID Controller	PERC 6/i	PERC H700	PERC H710
RAID FW	6.2.0-0012	12.3.0-0032	21.0.1-0132
HDD	2 x 10k SAS 36GB RAID 1	2 x 10k SAS 300GB RAID 1	2 x 10k SAS 300GB RAID 1
PSU	870W (1 of 2 slots populated)	870W (1 of 2 slots populated)	1100W (1 of 2 slots populated)
Optical drive	IDE CD-ROM	SATA CD/DVD	SATA CD/DVD
BIOS	2.4.3	6.1.0	1.0.0
iDRAC FW	BMC v 2.28	1.80.17	1.05.05
Network Adapters	Broadcom® NetXtreme II™ (5708)	Broadcom NetXtreme II (BCM5709C)	Broadcom NetXtreme II (BCM5720)
Network FW	7.4.8	7.4.8	7.4.8
CPLD	A00 (1.0.0)	A00 (1.0.0)	A00 (1.0.0)

Table 5. BIOS settings

System	PowerEdge 2950 III	PowerEdge R710	PowerEdge M820
Virtualization	Enabled	Enabled	Enabled
Turbo mode	N/A	Enabled	Enabled
Memory frequency	667MT/s	1333MT/s	1333MT/s
Power management	Maximum performance	Maximum performance	Maximum performance
C1E	N/A	Enabled	Enabled
Hardware prefetchers	Enabled	Enabled	Enabled

Table 6. Software configuration

System	PowerEdge 2950 III	PowerEdge R710	PowerEdge M820
Hypervisor version	ESX 4.1 Update 2	ESX 4.1 Update 2	ESXi 5.0
DBMS version	SQL Server 2008 R2	SQL Server 2008 R2	SQL Server 2012
Virtual guest OS (DS2 VMs)	Windows Server 2008 R2	Windows Server 2008 R2	Windows Server 2008 R2
Virtual guest OS (DB table creation)	Red Hat Enterprise Linux 6.1 Server	Red Hat Enterprise Linux 6.1 Server	Red Hat Enterprise Linux 6.1 Server

Table 7. Shared storage configuration

System	PowerEdge 2950 III	PowerEdge R710	PowerEdge M820
SAN type	Dell EqualLogic PS6010XV	Dell EqualLogic PS6010XV	Dell EqualLogic PS6010XV
Number of enclosures	4	4	4
Number of LUNs	7	7	7
Number of disks/enclosure	16	16	16
Hard drive description	600GB / 3.5" SAS 15k	600GB / 3.5" SAS 15k	600GB / 3.5" SAS 15k
RAID	10	10	10