

Microsoft® SQL Server® in a VMware® Environment on Dell™ PowerEdge™ R810 Servers and Dell EqualLogic™ Storage

A Dell Technical White Paper

Dell | Database Engineering Solutions

Anthony Fernandez

April 2010



THIS WHITE PAPER 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.

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the *DELL* logo, and the *DELL* badge, *PowerEdge*, *PowerConnect*, and *Dell EqualLogic* are trademarks of Dell Inc. *Microsoft*, *Windows*, and *SQL Server* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries. *ESX*, *VMware*, *vSphere*, and *VMotion* are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Table of Contents

- Executive Overview 2
- Introduction 3
- Audience and Scope 3
- Reference Architecture Overview 3
- Deployment Considerations..... 4
 - Server Considerations..... 4
 - Network Considerations 5
 - CPU Considerations..... 5
 - Memory Considerations 5
 - External Storage Considerations 6
- SQL Server Performance Evaluation..... 6
 - Physical versus Virtual OLTP Performance 8
 - Scale Out Virtual Machines on a Single Server 8
- Conclusion 11
- References..... 11

Executive Overview

This white paper describes the benefits and performance of the latest Dell™ PowerEdge™ R810 running Microsoft® SQL Server® in VMware® vSphere™ 4.0 Update 1 (U1) and EqualLogic™ PS6000XV iSCSI storage arrays. The Dell PowerEdge R810 offers an ideal platform to consolidate demanding virtual database workloads in space conscious datacenters. The R810 also offers energy efficiency features in a dense 2U, form factor rack server with up to a four-socket capacity and up to 512GB of RAM.

A series of evaluation and performance tests were executed to determine the performance the PowerEdge R810 can deliver in a virtual database environment. The tests compared the overall performance improvement against the previous server generation, as well as the added capacity.

Introduction

The PowerEdge R810 with the latest Intel® Xeon® 6500 and 7500 Series family of processors, with up to 512GB of RAM, is an ideal candidate for IT-Green initiatives such as datacenter consolidation, and power and cooling reduction while maintaining high levels of productivity and performance.

Virtualization allows operating systems and applications to be decoupled from physical hardware. With the computing power and high-memory capacity, the PowerEdge R810 provides a platform where multiple virtual machines can operate in parallel. Dell EqualLogic iSCSI SAN provides another layer of virtualization at the storage layer that allows ease of administration, performance and scalability.

This white paper describes the reference architecture that designs and implements a virtual infrastructure running database workloads; the tests simulated users and applications that depend on database engines.

Audience and Scope

This white paper is intended for customers, partners, solution architects, storage administrators, and database administrators who are evaluating, planning, or deploying virtual machines. It describes an overview of deployment considerations and key points for sizing and deploying SQL Server in a virtualized environment.

Reference Architecture Overview

The Dell Reference Architecture describes the building blocks needed to deploy a virtual infrastructure utilizing a multi-tier architecture that contains redundant components that provide high availability at every level. Figure 1 below shows a high-level diagram of the components used in the architecture.

Figure 1: Reference Architecture



Table 1 lists the reference architecture components.

Table 1 Hardware Components

	Description
Server	PowerEdge R810
Processors	Four Intel® Xeon® X7550, 2.92GHz, 8MB Cache, 6.40GT/s QPI, Turbo, HT
Memory	64GB 8x8GB RDIMM
Internal Drives	Two 2.5" 15K 73GB RAID 1 for the operating system Four 2.5" 15K 143GB RAID10 for ESX backup data store (optional)
Hard Drive Controller	PERC H700
BIOS	Virtualization Technology (VT) Enabled
Remote Management Controller	iDRAC6 Enterprise
Operating System	VMware vSphere 4.0 U1
Additional Network Cards	Two Intel Pro 1000VT Quad-Port 1GbE
Network Switches	Two PowerConnect™ 6248P (Stacked)
Storage	Dell EqualLogic PS6000XV iSCSI SAN (3 Members)

Deployment Considerations

This section describes the high-level considerations for designing a virtual infrastructure that utilizes the latest PowerEdge R810. It describes key points and recommendations based on lab testing.

Server Considerations

The PowerEdge R810 server 2U form factor has up to six internal 6GBps SAS drives, and it allows organizations to maintain a small hardware footprint. The R810 offers a balance between performance and I/O capacity; with six PCIe slots the server has high network and I/O data path availability.

The server was configured with all six internal 2.5" drives. Drive 0 and Drive 1 were configured to use the ESX™ OS in a RAID 1 virtual disk, and Disk 2- 6 was configured in a RAID 10 virtual disk that was utilized for the virtual machine back-up space.

Figure 2: R810 Front View



Table 2 Internal Hard Drive Configuration


Virtual Disk Configuration	Description
Virtual Disk 1: Operating System	Two 73GB 15K SAS HDD (RAID 1)
Virtual Disk 2: Backup Data store	Four 173GB 15K SAS HDD (RAID 10)

Network Considerations

For consolidated virtual environments running database workloads through the iSCSI layer, sizing the bandwidth and throughput requirements based on the database workload I/O characteristics is essential. The number of add-on network NIC’s required for drive network traffic depends on the number virtual machines running on the system, in addition to the workload characteristics of the virtual database servers.

For example, Figure 3 shows the architecture network cable configuration. The embedded network adapters are connected to the public and VMotion™ networks. The VM iSCSI traffic is load balanced across two Intel Quad-Port NICs on slot 1 and slot 5 located on PCIe Riser 1 (Left) and 2 (Right). Each physical NIC is mapped to a virtual NIC, and is configured within a virtual switch.

Figure 3: Network Diagram



CPU Considerations

The Dell PowerEdge R810 can be configured with one, two or four Intel Xeon 4-core and 8-core processors for a total of 24 or 32 Cores. For high-performance environment that host a large number of virtual machines with high CPU utilization, the X7560 2.26GHz 8-core processor is recommended. For non-CPU intensive workloads, especially in a consolidation environment where the goal is to reduce legacy hardware, the E7540 or X7542 are recommended.

Memory Considerations

Memory sizing is primarily determined by the amount of RAM needed for each virtual machine hosted on the system. Even though memory sharing is supported by VMware, it is recommended to have the memory required for the virtual machine database workload without over committing; VM memory sizing for running a database workload should be carefully analyzed to mitigate any storage bottlenecks.

External Storage Considerations

The external storage ensures that the virtual environment performs to the required level. EqualLogic storage offers many virtual environment benefits, such as ease of use and configuration. It is important to group the VMs according to workload and have that workload assigned to a specific pool of storage resources to avoid conflicts. For example, if there is a host performing large sequential operations, such as reporting and data mining, on the same storage array, the VMs that have a random access pattern impact the entire pool of resources negatively, reduce the performance of the entire pool, and lengthen the operation time.

SQL Server Performance Evaluation

This paper evaluates the overall performance capabilities of the recommended configuration. The goal is to simulate a real-world environment that consolidates the SQL Server instances running on physical hardware into a virtual environment. The tests analyzed the overall impact of running multiple virtual machines with SQL Server on the PowerEdge R810.

Figure 4 shows the logical representation of virtual machine operating system files and database files. Each virtual machine is configured with two virtual hard drives (VMDK), one for the operating system and a separate one for the database files.

Figure 4: Storage Disk Configuration



The first step is to test that the I/O path from the virtual machine to the storage is clear of bottlenecks introduced by human error. To do this, the team characterized the overall storage performance using the Microsoft SQLIO tool. SQLIO simulates I/O transactions, and can find the maximum threshold of the storage subsystem independent of the CPU or server memory capacity. The team ran SQLIO with a small random I/O and a large sequential I/O on a separate physical server running Windows 2008 R2 to

characterize the maximum IOPS and MBPS storage capabilities. The team then configured ESX and the vKernel data path. A volume hard disk was created and dedicated to the database on the shared storage. A test VM running Windows 2008 R2 was also deployed on this system and the additional hard disk was formatted with NTFS. SQLIO was run to attain the same I/O characteristics as the physical server. The number of NICs for the iSCSI traffic was increased so that load balancing and multi-pathing could be implemented at the ESX level. The team tested again within the VM to verify that the random I/O characteristics remained unchanged, and that the large sequential throughput increased as new data paths were added.

The next step was to deploy a single VM with Windows 2008 R2 and install SQL Server 2008 SP1; a database was created on the shared storage data volume. To simulate a user workload, a TPC-E schema was deployed on the database VM to simulate an online trading broker.

Two categories of tests were created. The first category evaluated the overall impact of running SQL Server in a virtual environment compared to a legacy physical system. The second category evaluated the performance of consolidating multiple workloads into a single PowerEdge R810; the number of VMs was scaled up to assess the impact of running parallel virtual machines.

The table 3 lists the parameters and assumptions used to generate a consistent load. The assumptions are based on a consolidation effort from a legacy environment to a virtual environment for non-mission critical level databases. These tests assume a low to medium level of access by applications and users. Table 4 lists the legacy hardware configuration, and Table 5 lists the virtual machine configuration.

Table 3 Configuration Assumptions

Configuration	Description
Database size	50GB
User load	20 application connections
Average System Utilization	5% - 10%

Table 4 Legacy Hardware Configuration

Configuration	Physical System Description
Legacy Server	PowerEdge 1950
CPU	Two Intel Xeon 5160 Dual Core 3.00GHz
Memory	4GB of RAM
Storage	Internal Storage (2 HDD's)
Operating System	Windows 2003 Server

Table 5 Virtual Machine Configuration


Configuration	Virtual Machine Description
CPU	2 vCPU
Memory	4GB of RAM
Storage	Shared Storage RAID 10 Pool (3 Members)
Operating System	Windows 2008 Server R2

Physical versus Virtual OTLP Performance

The first test evaluates the legacy system performance and compares it to a virtual instance of SQL Server. Figure 5 below shows the side-by-side comparison of a 9th generation server versus a virtual machine running on the PowerEdge R810. The 9th generation server has two Intel Xeon dual-cores, X5160 at 3.00GHz, and 4GB of RAM, and two 3.5” internal hard drives for the OS and database files. The VM was configured with 4 vCPUs, and 4GB of RAM. The database files reside on the external iSCSI SAN.

A set user load of 20 users was utilized on both the physical legacy server and the virtual machine. The CPU and transactions per second (TPS) were recorded for each test. The test yielded a 23% improvement of TPS on the legacy server over the virtual machine system with a 25% CPU increase. The ESX host had a total of 5% CPU utilization. In the next test the team scaled out the virtual machines to replicate the workload, so as to measure the benefits of consolidating in a virtual environment.

Figure 5 Legacy Physical versus Virtual Machine




Scale-Out Virtual Machines on a Single Server

The scale out test evaluates running multiple virtual machines in a consolidated configuration on the PowerEdge R810. The PowerEdge R810 supports running with one, two and four sockets. To take advantage of the higher memory configurations that can drive the workloads of multiple virtual machines, it is recommended the server be deployed with two- or four-socket configurations.

The first test compared the physical versus virtual to assess the overhead for this type of workload. This test evaluates the overall performance of the system as the user load is increased on the physical legacy server versus running individual virtual machines, with each having a copy of the database and a constant number of users. The bottom line shows the performance of the PE1950 relative to the R810 as the user load is increased. The results for the PE1950 showed decreasing performance as the users were added due to CPU constraints of reaching over 90% utilization at an 80 user load. In contrast, the PowerEdge R810 continued to scale-out virtual machines to 13 VMs for this workload, and achieved three times more transactions than the legacy server.

Figure 6 Legacy Servers versus Scale-Out VMs



In addition to the increased performance capacity of the R810, the power consumption was compared to the legacy servers and the findings are presented in Figure 7 below.

Figure 7 Performance per Watt Comparison





Figure 7 shows that as the number of VMs hosted on the R810 increases, the performance to watt benefit continues to increase; the power envelope does slightly increase as new VMs are deployed. By

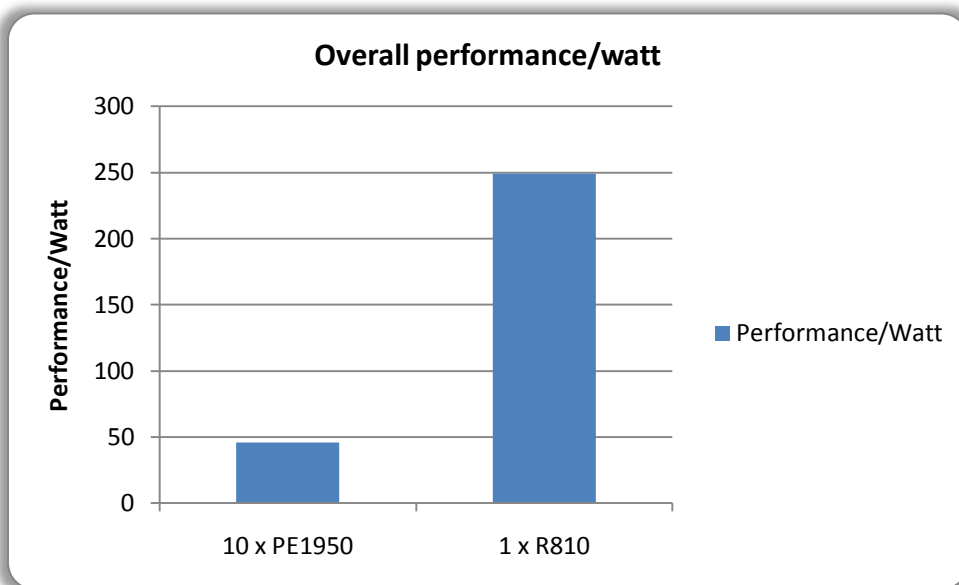
contrast, the legacy servers' power envelope is multiplied by the number of servers. Figure 8 shows the combined power consumption and transactions per second of ten PE1950's. The combined performance those servers is 34% more transactions per second than a single R810 hosting 10 VM's, but the real benefit is the 88% reduction in power of the R810 compared to the PE1950's.

Figure 8 Performance and Power Consumption Comparison



Plotting this result yields a performance/watt graph as shown in Figure 9. This graph shows the benefits of a single R810 in a consolidated virtual environment as compared to multiple legacy physical servers.

Figure 9 Overall Performance/Watt Legacy versus Virtual



Conclusion

This white paper describes the benefits of the latest PowerEdge R810, deployed in a virtual environment that is running database workloads as compared to legacy physical hardware. The tests performed at the Dell Engineering labs show that the performance improves and power consumption goes down when several database servers running on legacy hardware into are consolidated into a single R810 server running VMware virtualization technology. IT organizations can benefit from consolidation efforts by having fewer physical servers to help meet energy efficiency goals, while at the same time maintaining separation of applications. In addition to power and cooling benefits, fewer R810 servers can be configured in a highly-available solution that provides fault tolerance and dynamically distributes loads across fewer servers. With VMware features such as VMotion, VMware High Availability (HA) and Fault Tolerant (FT), organizations can ensure application availability while reducing complexity.

The latest generation of Dell servers, such as the R810, provides the CPU and memory capacity to drive demanding database workloads while helping organizations meet their legacy hardware consolidation goals. In addition, this configuration provides scaled-out performance to meet future growth.

Dell helps to simplify growth with the latest Dell EqualLogic iSCSI storage SAN. Data storage is an important aspect of consolidating and virtualizing database workloads. In the past, a poorly configured Fibre Channel (FC) SAN could negatively impact the organization's database, application performance, and ultimately the bottom line. The Dell EqualLogic iSCSI virtualized SAN simplifies management, and scales-out as the organization grows without the need to rearchitect the SAN. Dell EqualLogic in conjunction with a virtual environment provides a complete solution to help organizations scale, and grow in capacity seamlessly while providing high performance and simplified management.

References

Dell Solutions Engineering

- <http://www.dell.com/sql>
- <http://www.dell.com/vmware>
- <http://www.dell.com/virtualization>

Dell Services

- <http://www.dell.com/services>

Dell Business Ready Configurations

- <http://content.dell.com/us/en/business/virtualization-business-ready-configurations.aspx>