

Reduce Costs & Increase Microsoft® SQL Server® OLTP Workload Service Levels:

PowerEdge™ 2850 Consolidation to PowerEdge™ 11th Generation

A Dell Technical White Paper

Dell | Database Solutions Engineering

Raghunatha M

Leena Basanthi K

April 2010



Executive Summary

Businesses of all sizes often face challenges with legacy hardware and software that are outdated. Deploying and maintaining multiple individual applications in a dispersed legacy environment increases IT staffing overhead as the number of servers for both collection and data storage grows. To overcome these issues, server manufacturers have created larger and larger servers capable of handling the workload of multiple smaller servers. Consolidation is the grouping of multiple applications or databases onto fewer servers. This practice allows replacement of larger servers for multiple smaller servers, which helps data centers reduce the number of servers they support. In addition, consolidation provides for better utilization due to advances in hardware and software technologies. The key enablers that favor server consolidation are the hardware processor, memory and disk space.

This white paper showcases how organizations can benefit significantly by moving from legacy 2-socket servers (Dell™ PowerEdge™ 2850 2-socket server) to new Dell PowerEdge 11th generation 4/2-socket servers (Dell PowerEdge R810 4/2-socket server) and the performance gains achieved. This document provides a consolidation factor, as measured by the CPU processing power of the new Dell 11G server and its ability to handle multiple copies of basic building blocks running concurrently while delivering at least the same performance as the legacy server for each workload.

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, *PowerEdge*, and *EqualLogic* are trademarks of Dell Inc. *Microsoft*, *Windows Server*, and *SQL Server* are registered trademarks of Microsoft Corporation in the United States and/or other countries. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the United States and/or other countries. *Broadcom* is a registered trademark of Broadcom Corporation or its affiliates in the US and 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.

April 2010

Contents

- Executive Summary 1
- Introduction 3
- Legacy Server to New Generation Server 3
- Need for Consolidation 4
 - Multiple Instances Per Physical Server 4
- Building Block Approach for Workload Consolidation 4
 - Base Building Block Architecture..... 5
 - Scaled-Up Building Block Architecture 5
- Test Configuration..... 7
- Test Scenarios 7
 - Legacy (PowerEdge 2850) Test Scenarios..... 8
 - 11G (PowerEdge R810) Test Scenarios 10
- Consolidation Factor 14
- Conclusion 15
- References 15

Introduction

The recent trend toward data center and server consolidation makes it important for a database server to be scalable and to perform well for a wide variety of applications. One of the important components in the IT environment is the database. Databases tend to be very widespread across the enterprise because they are very efficient for handling relational storage of data and are designed to be platforms for a wide variety of applications. Because databases form the foundation of so many business systems, organizations can easily lose control of the number of databases that need to be maintained since many databases are created to solve a specific problem they may be experiencing. Thus, databases are one of the prime candidates for consolidation. Organizations should take the initiative to optimize the IT infrastructure by opting for a comprehensive solution such as server consolidation. Modern server resources will process the workload faster by reducing administration, maintenance and hardware costs. The Dell PowerEdge portfolio offers a new generation of servers with flexible services such as easy configuration, powerful open-standard systems management and scalable storage. This server product line uses the new range of Intel® Xeon® Nehalem-based chips and storage.

The first step discussed in this paper was to determine the amount of database workload a legacy server, in this case the PowerEdge 2850, can sustain. Using this information, a basic building block was derived which is focused mainly on database size, SQL Server target memory and user load with adequate other resources like storage disks, network bandwidth and server physical memory. Finally, the number of basic building blocks that a legacy server can sustain at the CPU saturation level was determined.

The second step discussed in this paper was to determine the CPU processing power of the new Dell 11G server and its ability to handle multiple copies of basic building blocks running concurrently while delivering at least the same performance as the legacy server for each workload. The resulting number is called the *Consolidation Factor*.

This paper will focus on consolidation strategies for online transaction processing (OLTP) applications storing data in the SQL Server Database Engine. OLTP applications typically focus on fast response times with smaller but more frequent queries and updates causing relatively small I/Os. But as OLTP database transactions are more I/O bound, with sufficient disks in backend storage subsystems, database host CPUs can be utilized efficiently and can handle larger database applications. In the subsequent sections, lab test results were discussed to showcase the approach to saturate the CPU and scale up by adding more CPUs to increase the server processing power, provided that adequate disks are added to handle more I/O requests thus improving the overall processing power to handle multiple SQL database workloads. Along the same lines, this paper showcases the server consolidation factor by how many legacy platforms, the PowerEdge 2850 running OLTP workloads, can be replaced with the PowerEdge 11G server, the R810. The OLTP database transactions are more I/O bound.

Finally, based on the results, one may determine how many standalone legacy servers running OLTP SQL database workloads can be consolidated on a single PowerEdge 11G R810 server.

Legacy Server to New Generation Server

Most of the IT organizations face challenges in maintaining the legacy hardware. This white paper discusses a legacy environment on a PowerEdge 2850 with SQL Server 2000 32-bit and how it is consolidated onto a PowerEdge 11G R810 server with SQL Server 2008 64-bit.

To begin with, the legacy server PowerEdge 2850 was designed to support up to two Intel Xeon processors (single-core or dual-core) supporting an 800 MHz front side bus and up to 16 GB of DDR-2 memory, along with support for PCI Express. There is a limitation 8th generation servers that if you decide to upgrade your system by installing a second microprocessor, you had to order the microprocessor upgrade kits.

The replacement for the legacy system (the PowerEdge 2850 2U server) is the latest PowerEdge 11G server (the PowerEdge R810 2U server) that supports 2P/4P configurations with 32 DIMMs. This product line of servers increases performance and reduces the amount of energy required to operate workload characteristics. The R810's new features include Intel Xeon (Nehalem-EX) processors, Intel's Boxboro-EX I/O Hub (IOH) with QuickPath Architecture, DDR3 memory, DIMM thermal sensors, PCI Express Generation 2, iDRAC with the integrated video controller, dual-port embedded Gigabit Ethernet controllers, the Internal Dual SD Module with data redundancy, and the iDRAC6 Express.

Need for Consolidation

Most of the IT staff or SQL administrators today, think that consolidation hits application performance, because a consolidated application shares the resources such as CPU, RAM and network devices with other applications residing on the same physical server. But at the same time, allocating legacy resources to multiple database applications is a challenge. To address these challenges Dell Solution Engineering Team's strategy is focused on reducing the complexity of SQL infrastructure which is, dispersed legacy physical server environment.

The 11G servers that come with newer architectures such as multiple CPU cores and a large amount of memory offer additional capability to handle the workload better and to help in scaling-out solutions.

The factors that drive server consolidation are advances in modern hardware and software, for example: the features of SQL Server Enterprise offer high availability to improve performance, reduce licensing costs and provide 64-bit hardware support. Server consolidation enables the organization to reduce infrastructure such as space, power, cooling, network connectivity and storage.

This paper shows the analysis conducted by the Dell solution engineers of how a single 11G server can now host many SQL OLTP workloads that previously were dispersed onto many legacy servers by choosing an appropriate strategy—instance level consolidation. Instance level consolidation is the consolidation of servers with multiple SQL Server instances onto a single physical server where each application contains its own database instance and all other necessary resources.

Multiple Instances Per Physical Server

Instance level consolidation with SQL server 2008 on Windows server 2008 increases the reliability, availability and facilitates scaled-up solutions in SQL infrastructures. Instance level consolidation enables isolation of databases which in turn eliminates the contention of temporary workloads in the tempdb database and this also helps to meet regulatory compliances. SQL Server 2008 Enterprise offers the greatest number of instances per server (up to 50).

Building Block Approach for Workload Consolidation

A building block approach is to optimize database size, user load and server target memory by eliminating all other bottlenecks such as IO and network for databases like OLTP, DSS or Mixed workloads. This is done to achieve compromised SLAs in terms of average response time and CPU utilization.

The Dell Solution Engineering Team has experimented on a legacy server configuration which has an OLTP application to derive a base building block that is optimized for different scales of OLTP requirements. The bottom line of this task was to achieve the SLA of 2 ms Average Response Time on the workload characteristics by stressing the legacy system to its saturation point.

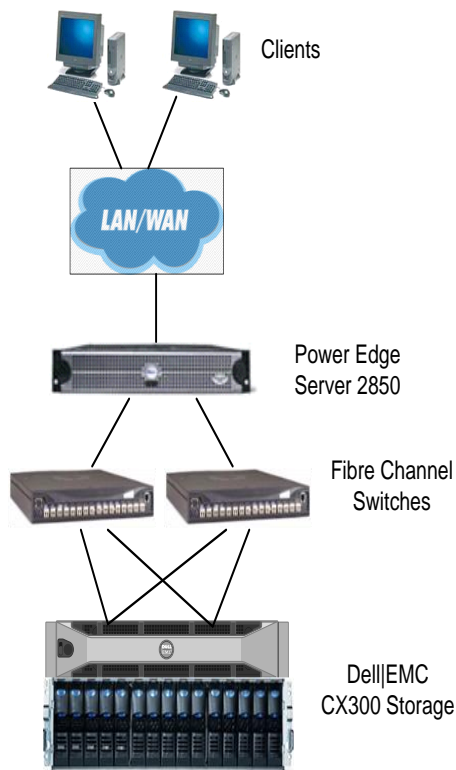
Base Building Block Architecture

The Dell Solution Engineering Team constructed a set of OLTP reference configurations that enable customers to start building an OLTP-based application upon best practices and experiences. The base building block of the legacy server should be a flexible building block for the new generation server, which allows scaling-out additional building blocks. The building block is constructed based on the workload requirements and infrastructure limitations.

The base configuration begins with a reference configuration as follows:

- 2 quad core CPUs PowerEdge 2850 server with 8 GB of RAM
- Dell/EMC CX300 storage system with 14x 146 GB disks
- 2 Brocade DS5300 80 x 4 Gb/s ports
- 2 QLOGIC 2460 single port FC HBAs (4 Gbps ports)

Figure 1. Base Building System Architecture



Scaled-Up Building Block Architecture

When the growth pushes the limits of the initial system block deployed, additional similar system blocks can be readily grouped to support the growth. The key is that each additional building block adds a balanced amount of server processing, memory and storage I/O support capability when

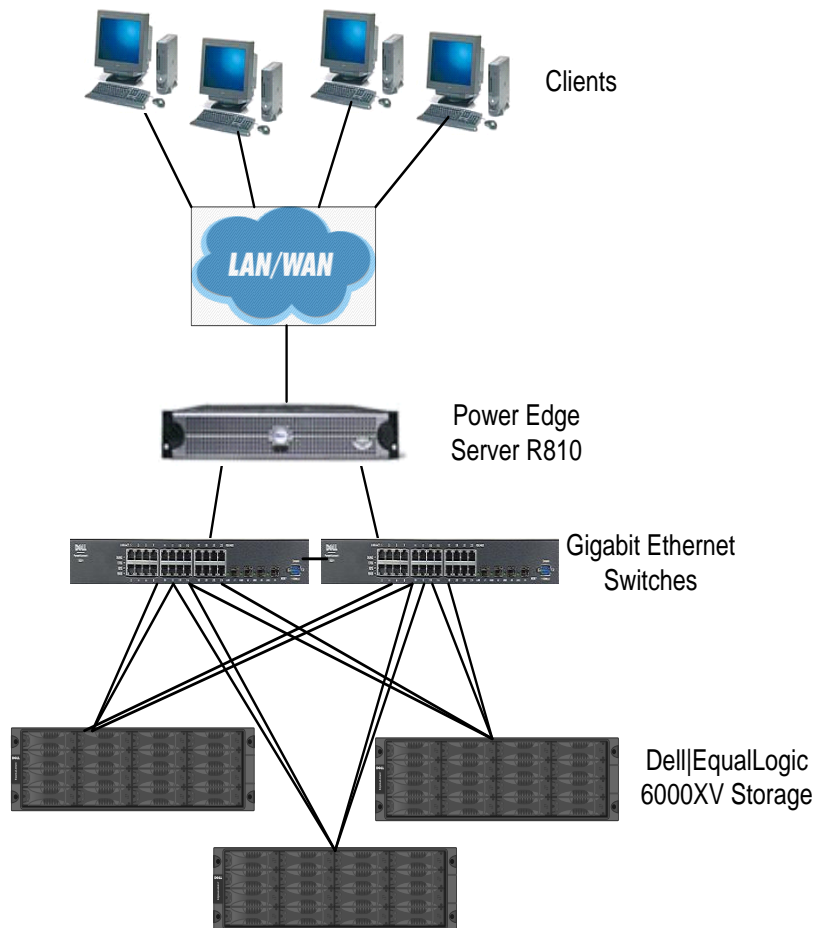
Reduce Costs & Increase Microsoft® SQL Server® OLTP Workload Service Levels

extreme high performance is required. Multiple blocks may be added to mainstream configuration enabling performance scaling that may be pushing the limit of a single block.

The scaled-up configuration begins with a reference configuration derived from the legacy configuration as follows:

- Two six core CPUs PowerEdge R810 servers with 32 GB of RAM
- Dell EqualLogic™ storage system with 28x 146 GB disks
- Two 6224 Gb Ethernet Switches

Figure 2. Scaled-Up System Architecture



Test Configuration

Table 1 describes the complete software and hardware configuration that was used throughout the testing on both the simulated legacy production environment and the 11G test environment.

Table 1. Test Configuration

Component	Legacy Production Environment	Dell PowerEdge 11G Test Environment
Systems	One PowerEdge 2850 2U 2S Server	One PowerEdge R810 2U 2S Server
Processors	Two Intel Xeon CPU 2.80 GHz Dual core CPU Cache: L2=1M	Two Intel Xeon E7530 1.87 GHz Six core CPU Cache: L2=4x256K L3=8 M
Memory	SQL1:2 GB SQL2:2 GB	SQL1:2 GB SQL2:2 GB SQL3:2 GB SQL4:2 GB
Internal disks	Two 73 GB 3.5" SCSI R1	Two 73 GB 2.5" SAS R1
Network	Two Intel 82544EI Gigabit Ethernet	Four Broadcom® NetXtreme II BCM5709 Gigabit Ethernet
External storage	Dell EMC CX300 with 14 x 146 GB FC disks Iteration1: 1 LUN of 200 GB RAID 10 for 1 SQL Instance Iteration2: 2 LUNs of 200 GB RAID 10 for 2 SQL Instances	Two Dell EqualLogic PS 6000XV with 14 x 146 GB SAS disks Iteration1: 2 LUNs of 200 GB RAID 10 for 2 SQL Instances Iteration2: 4 LUNs of 200 GB RAID 10 for 4 SQL Instances Iteration3: 6 LUNs of 200 GB RAID 10 for 6 SQL Instances Iteration4: 8 LUNs of 200 GB RAID 10 for 8 SQL Instances
HBA	Two QLE2460 per node	Two Broadcom NetXtreme II BCM5709 Gigabit Ethernet
OS	Windows Server 2003 x86 (32-bit)	Windows Server 2008 SP2 x64
SQL software	SQL Server 2000 x86 (32-bit)	SQL Server 2008 x64
Workload	Quest Benchmark Factory 6.1.1 TPC- C workload: Scale factor: 800 User connections: 100-2000	Quest Benchmark Factory 6.1.1 TPC- C workload: Scale factor: 800 User connections: 100-10000

Test Scenarios

Test cases were divided into two major categories:

- First, to derive a base building block on the legacy server, the PowerEdge 2850 with the legacy SQL database deployed (SQL Server 2000 32-bit)
- Secondly, to evaluate how many base building blocks the 11G server, the PowerEdge R810, with SQL database deployed (SQL Server 2008 64-bit), will sustain

Note: One Building block = Single SQL Server instance, 100 GB data, 2 GB SQL Server target memory and 1600 user load

Reduce Costs & Increase Microsoft® SQL Server® OLTP Workload Service Levels

Legacy (PowerEdge 2850) Test Scenarios

- **Test Case 1:** A single SQL Server 2000 instance with 100 GB of data and 2 GB of SQL Server target memory
- **Test Case 2:** An additional SQL Server 2000 instance with 100 GB of data and 2 GB of SQL Server target memory was added, totaling two active SQL Server instances with 100 GB of data and 2 GB of SQL Server target memory each

OLTP Workload

The Quest Benchmark Factory has a framework that provides industry-standard TPC-C performance runs. TPC-C was used to load the SQL databases with 100 GB of data (800-scale). Transactions were carried out with increasing user loads from 100 to 3000 (in increments of 100). The performance statistics were collected for each of the user loads.

Results and Analysis

Figure 3. Legacy Server PowerEdge 2850 TPS Plot





Figure 4. Legacy Server PowerEdge 2850 Response Time Plot



The transaction-per-second metrics were captured for two SQL instances on the legacy server. The preceding TPS and Response Time plots are cumulative average values for two building block test runs. From the preceding graphics, it is observed that the two SQL instances could scale up to 1600 user loads with 2 ms Response Time SLA (2 ms is a typical industry-accepted SLA for OLTP workloads) at CPU utilization well below 80 %. But beyond 1600 user loads, the CPU was stressed to its saturation level, provided that I/O, network bandwidth and physical memory did not cause a bottleneck.

Figure 5. Legacy Server PowerEdge 2850 CPU Utilization Plot





Figure 6. Legacy Server PowerEdge 2850 Average DiskQueLength Plot



The preceding CPU Utilization and Average DiskQueLength plots are average cumulative values for two building block benchmark test runs. From the preceding graphs, it is observed that Average DiskQueLength was well below 200, which indicated that disks can still sustain more I/O. In contrast to the preceding TPS and Response Time graphs (Figures 3 and 4), it is observed that at the CPU saturation level, the legacy environment sustained a user load of 1600. From this analysis, a user load of 1600, a database size of 100 GB and 2 GB of SQL target memory is determined to be a *Building Block*. Thus, it can be concluded that the legacy server (the PowerEdge 2850) can sustain two building blocks at well below 80% CPU utilization and can maintain 2 ms Response Time SLA.

11G (PowerEdge R810) Test Scenarios

- **Test Cases1:** Started with two building blocks each with 100 GB of data and 2 GB of SQL Server target memory
- **Test Case 2:** Added two more building blocks totaling up to 4 SQL instances, 100 GB of data and 2 GB SQL target memory each
- **Test Case 3:** Included two more building blocks with up to 6 SQL instances with 100 GB of data and 2 GB of SQL Server target memory each
- **Test Case 4:** Increased SQL Server instances count to eight with two additional building blocks with 100 GB of data and 2 GB SQL Server target memory each

OLTP Workload

The Quest Benchmark Factory has a framework that provides industry-standard TPC-C performance runs. TPC-C was used to load the SQL databases with 100 GB of data (800-scale). Transactions were carried out with increasing user loads from 100 to 8000 (in increments of 100). The performance statistics were collected for each of the user loads.

Figure 7. PowerEdge 11G R810 Server TPS Plot





Figure 8. PowerEdge 11G R810 Server Response Time Plot



Results and Analysis

The saturation point on legacy server PowerEdge 2850 is considered as the starting point for tests on 11G server PowerEdge R810. On the same note started test with two building blocks and continued till eight building blocks in increments of two building blocks.

Figures 7 and 8 show average cumulative values for all the eight building blocks. From the preceding plots, it is observed that the 11G server was able to handle the OLTP workloads much faster as compared to the legacy server, and it sustained very high user loads in the range of 4500 to 5000 for

two building blocks at 2 ms Response Time SLA. By continuing the tests up to eight building blocks in increments of two building blocks for each test iteration, it is observed that the user loads were sustained in the range of 2300 to 2500 at 2 ms Response Time SLA.

Figure 9. PowerEdge 11G R810 Server CPU Utilization Plot





Figure 10. PowerEdge 11G R810 Server Average DiskQueueLength Plot



The preceding CPU utilization and average DiskQueueLength plots were derived from the average cumulative values from the eight building block results. From the CPU analysis plot for all the building blocks tested, the maximum average utilization was well below 30% as shown in the preceding plot as compared to the legacy environment of which 80% was utilized for the two building blocks discussed previously. While these tests were actively running, storage subsystem I/O was monitored closely to

ensure there was no bottleneck. From the average DiskQueLength plot, it is observed that the average DiskQueLength was well below 200 for user loads in the range of 2700 to 2800. But for the eight building block tests, I/O was reaching a bottleneck because the storage disks used (28 146 GB x 15K SAS) were saturated by multiple instances concurrently requesting I/O. However, there is enough CPU processing power available to handle more and more workloads as evident from the plot (CPU was 25% to 30% used). The tests were conducted further to deep dive to avoid an I/O bottleneck by doubling the spindle count (56 146 GB x 15K SAS). Further results and analysis are discussed as follows.

Figure 11. PowerEdge11G R810 Server Average DiskQueLength Plot for Different Disks




Figure 12. PowerEdge11G R810 Server Average TPS Plot for Different Disks


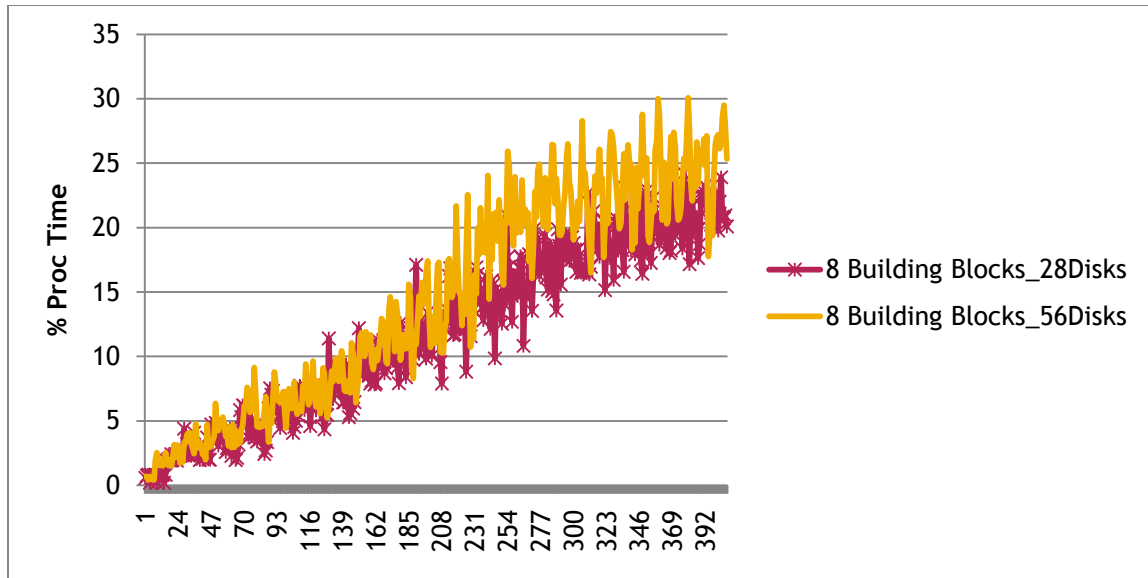


Figure 13. PowerEdge11G R810 Server Average CPU Utilization for Different Disks



Figures 11, 12 and 13 show the cumulative average values of the eight building block benchmark test run results. From the results analysis, we can observe that with additional disks TPS improved by 10%, CPU utilization increased by 5%, average user loads sustained also increased by 10% from 2800 to 3200 and CPU utilization increased. Thus, from the overall test analysis, to take advantage of faster processing power of the PowerEdge 11G server, I/O bottleneck should be eliminated by increasing the spindle count.

Consolidation Factor

After a detailed analysis of the preceding results, the conclusion can be drawn that a PowerEdge 11G server running SQL Server 2008 64-bit handles higher user loads for OLTP workload when compared to a single legacy server running SQL Server 2000 32-bit instances while maintaining an SLA of 2 ms Average Response Time. The consolidation factors derived from the test result analysis are as follows:

One R810 server (< 30% CPU utilization with 8 building blocks) with 2 sockets populated = four * PowerEdge 2850 servers (< 80% CPU utilization with 2 building blocks) with 2 sockets populated

Upon exploration, the consolidation factor can be derived as follows:

- One R810 server with 2 sockets populated = eight * PowerEdge 2850 servers with 2 sockets populated and similar CPU utilization
- One R810 server with 4 sockets populated = 16 * PowerEdge 2850 with 2 sockets populated and similar CPU utilization

A fully populated PowerEdge 11G R810 server with 4 sockets and a 6 core CPU can replace up to 16 legacy PowerEdge 2850 servers with similar CPU utilization, provided that both environments are not constrained by storage subsystem, network bandwidth or memory resources.

Conclusion

Dell's focus is to encourage organizations to improve their SQL IT infrastructure by consolidating legacy database environments on PowerEdge 11G servers equipped with Intel Xeon Nehalem-EX processors. This paper concludes that a PowerEdge 11G R810 server populated with two six core CPUs would be able to consolidate the workload of an eight-node legacy environment, provided that both environments are not constrained by storage subsystem, network bandwidth or memory resources. By extrapolating the results, we can arrive at the conclusion that a PowerEdge 11G R810 server with 4 sockets and a six core CPU can replace up to 16 legacy PowerEdge 2850 servers with similar CPU utilization with sufficient underlying resources. This result is possible due to the better utilization of advances in hardware and software technologies where multiple OLTP SQL databases are consolidated onto one physical Dell PowerEdge 11G server.

References

http://www.dell.com/content/topics/global.aspx/sitelets/solutions/software/db/microsoft_sql_2008?c=us&cs=555&l=en&s=biz

Consolidation guides: SQL on Dell 2650 to Dell 2950, Dell 6650 to Dell R900, Dell 6650 to Dell R905, Dell 2650 to Dell R80