

---

# Dell™ PowerEdge™ R720xd Scalability for Microsoft® SQL Server® 2008 R2 Databases


---

*A Dell Technical White Paper explains how R720XD can be scaled up from processor and disks point of view for OLTP based SQL Server database deployments.*

**Narasimha Reddy Gopu**

**Dell Database Solutions Engineering Team**

**March 2012**



**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, 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, and SQL Server are either trademarks or registered trademarks of Microsoft 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.

March 2012| Rev 1.0

## Contents

Executive summary.....	4
Audience .....	4
Introduction .....	4
PowerEdge R720xd architecture .....	5
Processors .....	5
Memory .....	5
PowerEdge RAID controllers and massive internal storage .....	7
Why choose the R720xd for SQL Server database deployments serving OLTP workloads.....	8
Achieving high availability and disaster recovery for SQL Databases deployed on R720xd internal storage.....	9
Test methodology.....	9
Storage scalability: half populated versus fully populated R720xd system.....	10
CPU Scalability: Single processor versus dual processor R720xd systems.....	14
Conclusion .....	17

## Tables

Table 1. Memory population per channel and the memory operating speed for RDIMMs .....	6
Table 2. Hardware and software configuration details used for disk scalability test .....	10
Table 3. Hardware and software component details used for CPU scalability test .....	14
Table 4. Performance benefits from PE R720xd scalability .....	18

## Figures

Figure 1. PowerEdge R720xd memory architecture.....	6
Figure 2. DIMMs population in single socket R720xd system .....	7
Figure 3. Front storage chassis of R720xd .....	8
Figure 4. Half populated versus fully populated storage comparison - transactions per second .....	12
Figure 5. Half populated versus fully populated storage comparison - response time .....	12
Figure 6. Half populated versus fully populated analysis: Average Disk Queue Length .....	13
Figure 7. Half populated vs. fully populated analysis: CPU usage .....	13
Figure 8. PowerEdge R720xd with single versus. dual socket comparison - transactions per second. ....	15
Figure 9. PowerEdge R720xd with single versus dual processor comparison - response time .....	16
Figure 10. PowerEdge R720xd with single versus dual processor comparison - CPU usage .....	16
Figure 11. R720xd with single versus. dual processor comparison - disk idle time .....	17

## Executive summary

Dell™ launched the 12<sup>th</sup> Generation PowerEdge™ servers, with support for Intel® Xeon® Sandy Bridge processors for small, medium and enterprise portfolios. The PowerEdge R720xd is a dense storage 2 socket (2S) 2U rack server with emphasis on capacity and performance for entry level customer environments. It was designed to provide an ideal single-platform for the application deployments that need a huge processing power and massive internal storage on a single server. The intent of this paper is to show how R720xd can be efficiently and effectively scaled up from storage and processor point of view.

This white paper focuses on the following aspects for Online Transaction Processing (OLTP) workloads.

- Discuss the R720xd components that make it ideal platform for hosting Microsoft® SQL Server® databases.
- How the R720xd can be scaled up from the storage point of view by comparing the performance results of half-populated and fully-populated server configurations.
- How the R720xd can be scaled up from the processor point of view by comparing the performance results of single socket and dual socket processor systems.

## Audience

The intended audience of this white paper includes database administrators, IT managers, and system consultants.

## Introduction

OLTP based applications have been important applications to the customers to drive their transactional nature of the business. SQL Server databases are one of the most popular databases for OLTP based applications. Since direct attached storage (DAS) enclosures are relatively affordable compared to a storage area network (SAN), it is becoming more common to use DAS storage, with multiple enclosures and multiple RAID controllers, to achieve very high throughput numbers for both OLTP and Online Analytical Processing (OLAP) workloads.

Dell strives to simplify the SQL Server database deployments by introducing a new kind of servers in its 12<sup>th</sup> generation server launch. R720xd was designed to provide an ideal platform for OLTP workloads

by supporting huge processing power, more memory, massive high performing internal storage and optimizations in many areas to improve the overall datacenter efficiency.

With a standard benchmark, we will show how the R720xd could be the best choice for quick SQL Server database deployments and how efficiently the R720xd can be scaled from the storage and CPU perspective. We will also discuss how the R720xd server can be a right choice for consolidation of many small servers hosting small SQL databases.

## PowerEdge R720xd architecture

The PowerEdge R720xd is a multipurpose 2-socket 2-rack unit server offering high performing internal storage and huge processing power. The server is configured with dual Intel Xeon processors and huge memory support. It supports PCIe Gen3 slots for external RAID expansion cards and other host adapters. The PowerEdge R720xd was designed for environments that need higher performing storage with huge processing power which makes it an excellent platform for OLTP deployments.

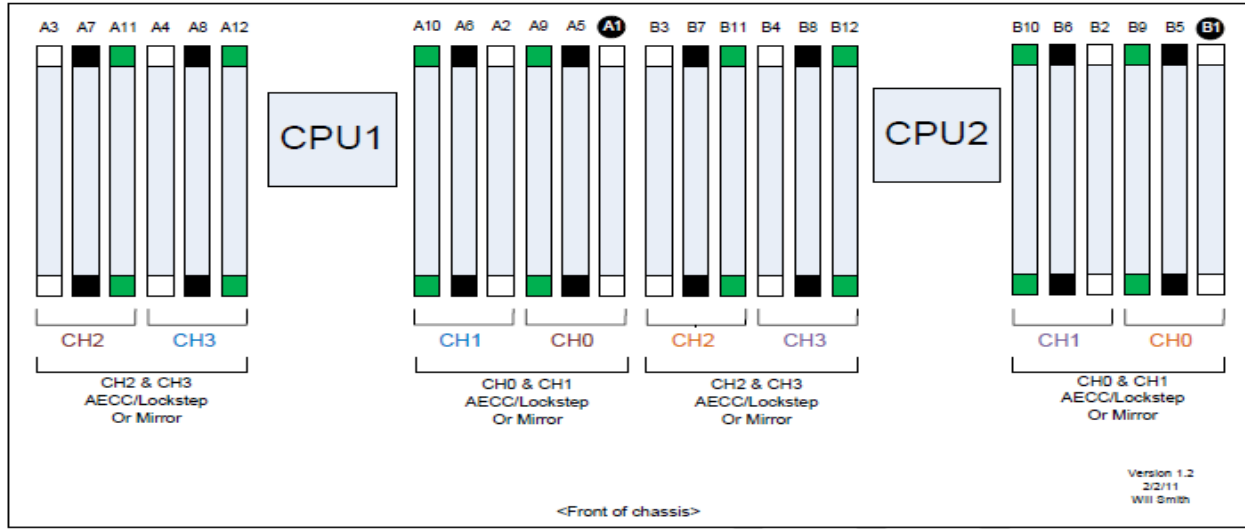
### Processors

The Dell 12th generation servers are designed based on the Intel Xeon Sandy Bridge architecture. The Sandy Bridge architecture delivers superior performance from its predecessors while cutting power costs and improving the integrated graphics performance. Sandy Bridge processors support high Quick Path Interconnect (QPI) links, maximum of 20 MB L3 cache and 1600 MHz memory DIMMS. All these components combined offer high bandwidth for inter-socket communication that delivers high performance. When fully populated, the R720xd can support up to 32 logical processors.

### Memory

The R720xd supports up to 12 DIMMs per processor and 24 DIMMs in a dual processor configuration. Each CPU socket has 4 memory channels with each channel supporting up to 3 DIMMs as shown in the Figure 1. When fully populated with 32 GB DIMMS, the R720xd can support maximum of 768 GB of memory. The operational memory speeds of 1600 MHz, 1333 MHz, 1066 MHz, and 800 MHz are supported. The memory operating speed depends on the DIMM type, the DIMM capability, and the number of DIMMs populated per channel.

Figure 1. PowerEdge R720xd memory architecture



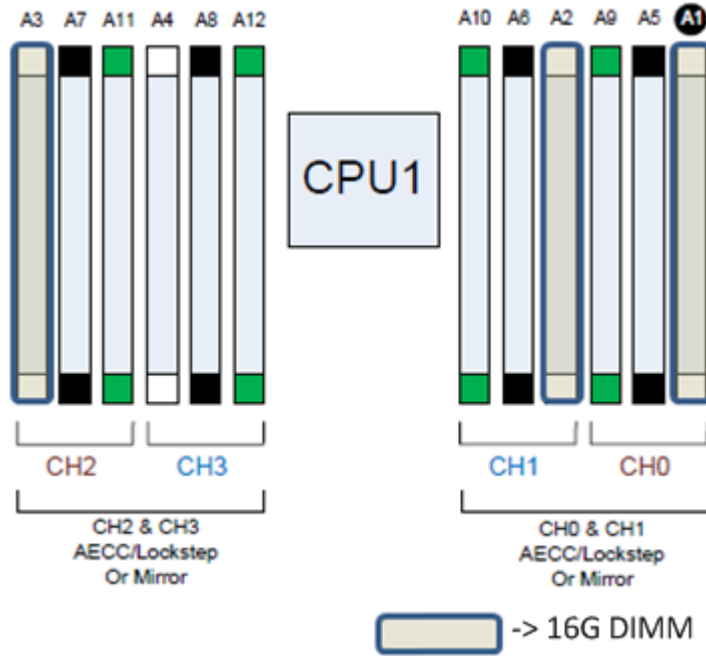
Since the operational speed of memory depends on many factors, some guidelines need to be followed in order to get the memory operating at the intended possible speed. Table 1 shows the memory DIMM populations and the resulting memory operating speeds for 1.5v operating voltage.

Table 1. Memory population per channel and the memory operating speed for RDIMMs

DIMMs Populated per channel	Memory Operating Speed (MHz) for 1.5v operating voltage	Max Ranks for DIMM in Channel
1	1600, 1333, 1066 (Quad Rank 1066,800)	Quad Rank
2	1600,1333, 1066 (Quad Rank 800 )	Quad Rank
3	1066, 800	Dual Rank

For example, the below Figure 2 shows how 16 GB memory DIMMs are populated to get 48 GB of memory with operational speed of 1600 MHz in single socket system.

Figure 2. DIMMs population in single socket R720xd system



## PowerEdge RAID controllers and massive internal storage

Within addition to the 12<sup>th</sup> generation servers, Dell introduced 8<sup>th</sup> series PowerEdge RAID Controllers (PERC) with many improvements and new features to support huge internal and DAS-based storage. The PERC H810 and H710P cards can be used to manage locally attached DAS storage.

In this configuration, we used the H710P Mini Monolithic PERC card which can support all 26 internal disks. The PERC 710P has 1 GB battery-backed cache. It supports online capacity expansion, dedicated and global hot spares, multiple cache policies, various RAID levels and operating systems. For more details on Dell PowerEdge RAID controllers, please refer to [Dell.com/PERC](http://Dell.com/PERC).

The R720XD has a flexible storage chassis that supports 2.5 inch and 3.5 inch SAS drives with various speeds (15K/10K/7.2K RPM) and capacities (73/146/300G etc.). This flexibility allows customers to choose a drive type and the number of drives based on their application requirements. For example, the PowerEdge R720xd supports 24 front-accessible 2.5-inch SAS drives as shown in the below Figure 3. Two additional SAS drives are accessible at the back of the system.

Figure 3. Front storage chassis of R720xd



## Why choose the R720xd for SQL Server database deployments serving OLTP workloads

OLTP workloads are characterized by high number of short and atomic transactions that generate much higher I/O Operations per Second (IOPS). Read and Write activities are random in most of the OLTP workloads. Hence, the OLTP workload performance is impacted considerably by the underlying disk speed, type and quantity. In addition to the high performing storage requirements, OLTP workloads also require huge processing power and more memory.

The PowerEdgeR720xd provides great flexibility to the customers to meet their storage requirements. Especially for OLTP workloads, the flexible design of the R720xd server supports up to 26 X 15,000 RPM, 2.5-inch SAS drives. The H710P PERC cards can be used to efficiently manage all of the internal drives.

By combining the processing power of the dual multi-core Intel sandy-bridge processors, high performing internal storage and huge RAM, the R720xd provides an excellent single-platform for hosting applications that require high processing power, memory and storage. The R720xd is an excellent choice for organizations to host cost-effective single server databases while delivering great overall performance. Microsoft SQL Server, being a high performance database for the small to medium sized business customers, would be a perfect database solution to be deployed on the R720xd.

The R720xd is also a good candidate for consolidating SQL databases and for virtualizing the SQL instances. With huge processing power, high memory and massive high performing internal storage,



the R720xd can support multiple SQL instances with multiple databases on a single-platform. The R720xd is also a good choice for hosting Hyper-V or VMware hypervisors.

The R720xd helps customers to scale up from single socket to dual socket configurations. It also provides the option to increase the internal storage up to 26 drives (Up to 3.79 TB of raw storage). We conducted several scalability tests at Dell lab to determine the R720xd behavior by gradually increasing the user load in OLTP workloads.

The R720xd provides options for storage scale out as well. It supports enough PCIe slots for network cards/HBAs to connect to external storage such as DAS or SAN (both FC, iSCSI).

For the complete feature details of the PowerEdge R720xd, please refer to [Dell.com/PowerEdge](http://Dell.com/PowerEdge).

## Achieving high availability and disaster recovery for SQL Databases deployed on R720xd internal storage

Mission-critical applications require continuous access to data. Microsoft SQL Server 2008 R2 provides various native high availability (HA) and disaster recovery (DR) features to achieve high availability of the databases. The native database HA features help us to recover the database/instances quickly and consistently.

With synchronous database mirroring and with Witness in place, SQL Server 2008 R2 provides excellent high availability for the databases with automatic failure detection and failover and with zero work loss using minimal resources and licensing costs. Log shipping or asynchronous database mirroring features would be right choice for DR solutions for SQL databases deployed with SQL Server 2008 R2.

The following sections of this whitepaper focus on the capability to scale up the PowerEdge R720xd for OLTP database workloads.

## Test methodology

The PowerEdge R720xd is the right choice for the customers who want to start with a minimal server configuration and then later scale up based on the application requirements. In this paper, we considered two such common factors: storage scalability and processor scalability. In the storage scalability scenario, we show how the R720xd performs with half-populated and fully-populated storage. In the latter scenario, we compare how the R720xd performs with single and dual processors.

In both the scenarios, we used internal hard drives on the R720xd server. Two drives were used to host the operating system with RAID1 and two disks were used as global hot spares. The remaining disks were used to store database and log files.

An OLTP testing tool was used to generate TPC-E workload to simulate OLTP stress on the database server. The TPC-E benchmark simulates transactions similar to those in brokerage firms. We chose TPC-E benchmark because it has a more OLTP database-centric workload with enhanced schema complexity. The transactions are more CPU and I/O intensive than TPC-C and are as random as TPC-C transactions.

Note: - TPC-E benchmark is very different from the other benchmarks like TPC-C. Therefore, the performance should not be compared against any other database workloads, including TPC-C.

### Storage scalability: half populated versus fully populated R720xd system

The PowerEdge R720XD provides an option to scale-up the storage by increasing the number of disks. In this section, we describe the performance benefits of increasing the storage in the R720xd from 16 to 26 disks (including OS and hot spares).

The table below describes the hardware and software components used in this scenario.

**Table 2. Hardware and software configuration details used for disk scalability test**

<b>HW/SW Component Used</b>	<b>Half Populated Disks</b>	<b>Fully Populated Disks</b>
<b>Server</b>	1 R720xd(BIOS 0.0.35)	1* R720xd (BIOS 0.0.35)
<b>Processor</b>	1 Intel Xeon E5-2680 8-core processor @2.7 GHz, HT Enabled, 16 Logical Cores	1 Intel Xeon E5-2680 8-core processor @2.7 GHz, HT Enabled, 16 Logical Cores
<b>Memory</b>	48GB (3 x 16 GB DIMMS) operating at 1600 MHz	48GB (3 x 16 GB DIMMS) operating at 1600 MHz
<b>PERC</b>	H710P (Firmware: 3.130.05-1462)	H710P (Firmware: 3.130.05-1462)

<b>Disks</b>	<b>10</b> X 15K RPM, 2.5" SAS disks for Data (RAID 10) <b>2</b> X 15K RPM, 2.5" SAS disks for Log (RAID 1) files <b>2</b> for OS deployment (RAID 1) <b>2</b> global hot spares	<b>20 x</b> 15K RPM, 2.5" SAS disks for Data (RAID 10) <b>2 x</b> 15K RPM, 2.5" SAS disks for Log (RAID 1) files <b>2</b> for OS deployment (RAID 1) <b>2</b> global hot spares
<b>Raw Data Storage Capacity</b>	1.360 TB total with 10 x 146 GB 15K RPM Disk drives	2.720 TB total with 20 X 146 GB 15K RPM Disks Drives
<b>Operating System</b>	Microsoft Windows® 2008 R2 SP1	
<b>RDBMS &amp; DB size</b>	Microsoft SQL Server 2008 R2 SP1 & 500 GB DB (simple recovery model)	
<b>Bench Mark</b>	TPC-E Workload	

As explained in the above table, in the half configuration, we used 10 hard drives for data files and 2 drives for Log files. Single RAID 10 virtual disk was created over 10 drives to host all the data (including Tempdb) files and RAID 1 virtual disk was created over the 2 drives to host Log files.

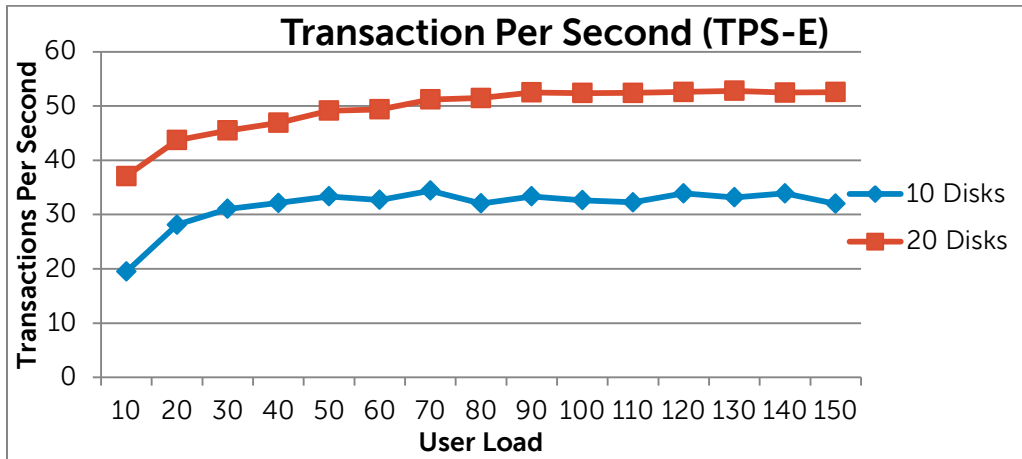
For the full configuration, we have extended the storage capacity by adding 10 more drives. A second RAID 10 virtual disk was created over new 10 drives and the database was rearranged to make sure that database was distributed equally between the two virtual disks.

In both cases, we gradually increased the user load to observe the I/O behavior of the two different storage configurations.

### Results and analysis

Figure 4 shows the Transactions per Second (TPS) delivered by fully populated system in comparison with half populated system.

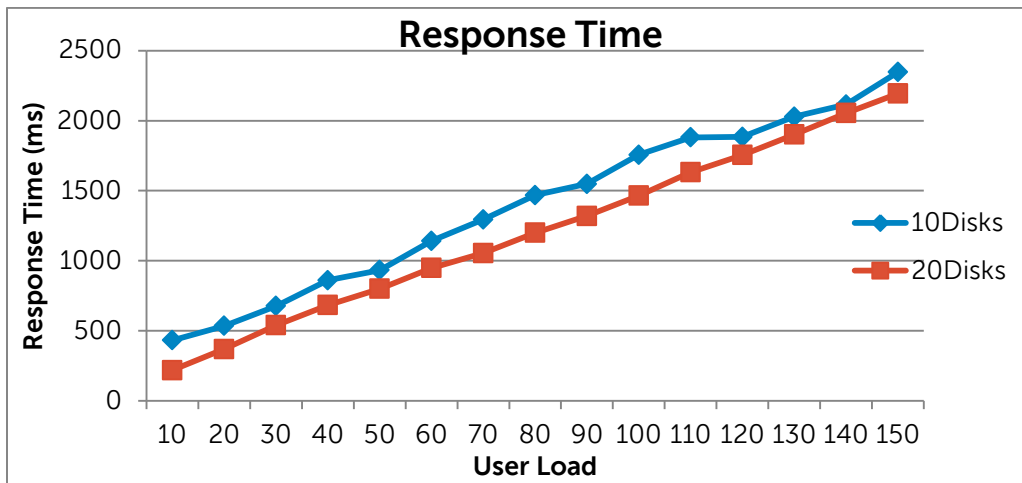
Figure 4. Half populated versus fully populated storage comparison – transactions per second



The fully populated R720xd system consistently delivered a TPS of 1.5 times greater than the half populated system even under maximum user load.

Figure 5 represents the response times observed during the workload.

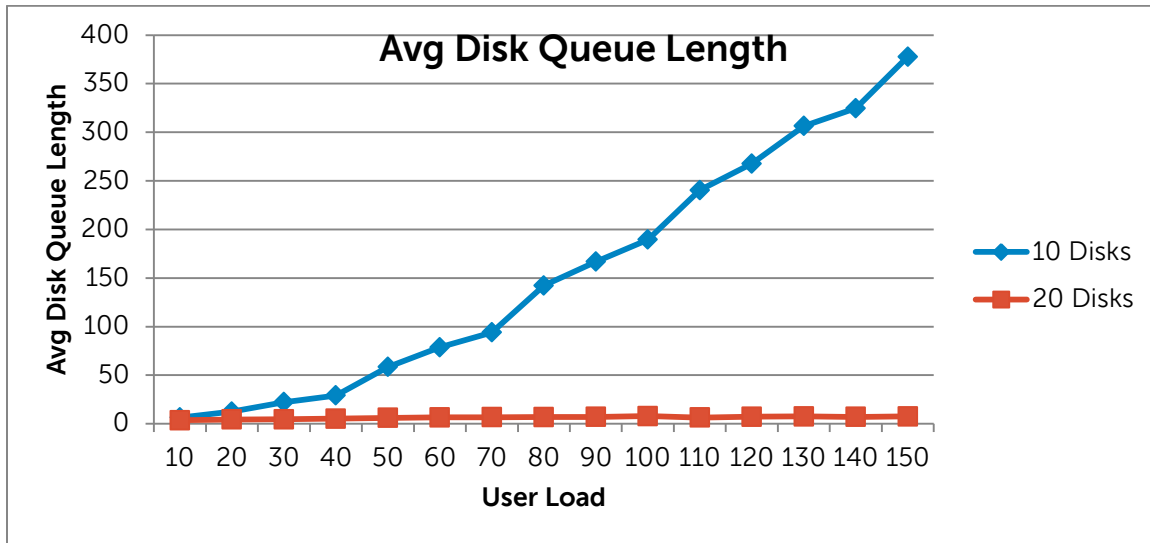
Figure 5. Half populated versus fully populated storage comparison – response time



The figures above shows that the fully populated storage system delivered 1.5 times higher TPS with 18 percent lower response time when compared to half-populated system.

Figure 6 represents the average disk queue length observed during the workload.

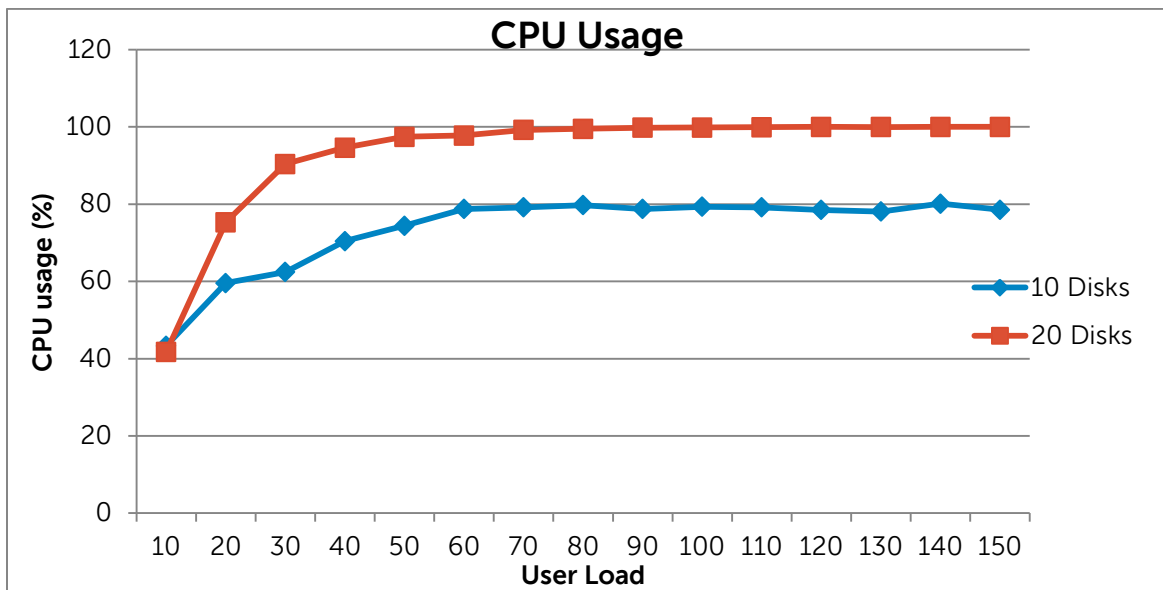
Figure 6. Half populated versus fully populated analysis: Average Disk Queue Length



In the half populated system, the average disk queue length values were observed to increase at a much greater rate than the fully populated scenario.

Figure 7 shows the CPU behavior in the half versus fully populated systems

Figure 7. Half populated vs. fully populated analysis: CPU usage



We observed a 20 percent better CPU utilization in the fully populated scenario, compared to the half populated scenario. The high disk I/O pressure in the half populated system caused the CPU to wait for disks to finish the requested I/O operations. In other words, in the half populated system the

workload bottlenecked at the disks causing the CPU to be underutilized. On the other hand, since the I/O pressure is alleviated by adding more physical disks in the fully populated system, storage subsystem was able to process more I/O requests thereby causing the CPU to be utilized to the maximum, based on the CPU-intensive nature of the TPC-E workload. This made CPUs able to process more transactions and thereby deliver a higher TPS with reduced response times.

### CPU Scalability: Single processor versus dual processor R720xd systems

The PowerEdge R720xd provides an option to scale-up its processing power from a single to a dual processor system. In this section, we describe the performance comparison of the single versus dual processor R720xd system.

The below table describes the hardware and software components used in the single versus dual processor configuration.

**Table 3. Hardware and software component details used for CPU scalability test**

HW/SW Component Used	Single Socket System	Dual Socket System
<b>Server</b>	1 x R720xd (BIOS 0.0.35)	
<b>Processor</b>	1 x Intel Xeon E5-2680 8-core processor @2.7 GHz, HT Enabled, 16 Logical Cores	2 x Intel Xeon E5-2680 8-core processor @2.7 GHz, HT Enabled, 32 Logical Cores
<b>Memory</b>	96GB (6 x 16 GB DIMMS) operating at 1600 MHz	
<b>PERC</b>	H710P (Firmware: 3.130.05-1462)	
<b>Disks</b>	20 X 15K 2.5" SAS disks for Data (RAID 10) files 2 X 15K 2.5" SAS disks for Log (RAID 1) files 2 for OS Deployments (RAID 1) 2 for Global Hot spares	
<b>Operating System</b>	Microsoft Windows 2008 R2 SP1	
<b>RDBMS &amp; DB size</b>	Microsoft SQL Server 2008 R2 SP1 & 500GB DB (Simple Recovery Model)	
<b>Bench Mark</b>	TPC-E workload	

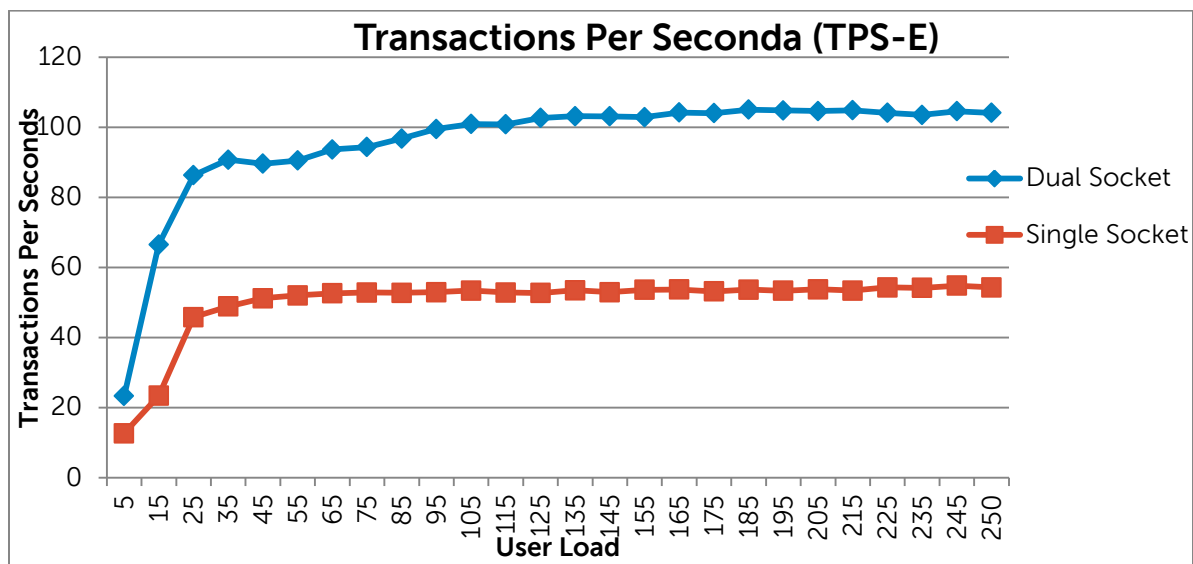
In the single socket configuration one Intel E5-2680 @2.7Ghz 8-core processor is used. In the dual socket configuration, two processors of the same type as in the single socket configuration are used. In both the cases, Memory and storage were kept constant which allowed us to focus on the CPU behavior. We used 96 GB of memory operating at 1600 MHz and 22 disks for storage. Two RAID 10 (10 Disks for each) virtual disks are carved out from 20 disks to store the database files and one RAID 1 virtual disk from the two drives for log files.

In the both the cases, we gradually increased the user load to see the CPU behavior of the two configurations.

### Results and Analysis

Figure 8 shows TPS delivered by dual processor system in comparison with single processor system.

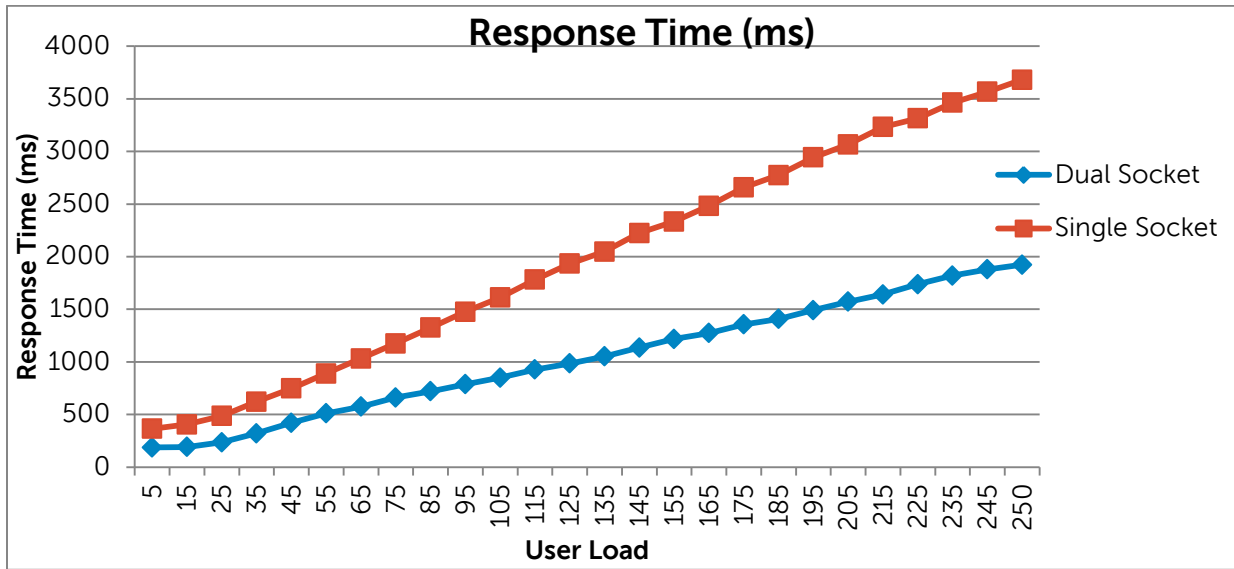
**Figure 8.** PowerEdge R720xd with single versus. dual socket comparison – transactions per second



The dual processor R720xd system constantly delivered a TPS of two times better than the single socket system even under high user load.

Figure 9 represents the response time observed during the workload.

Figure 9. PowerEdge R720xd with single versus dual processor comparison – response time

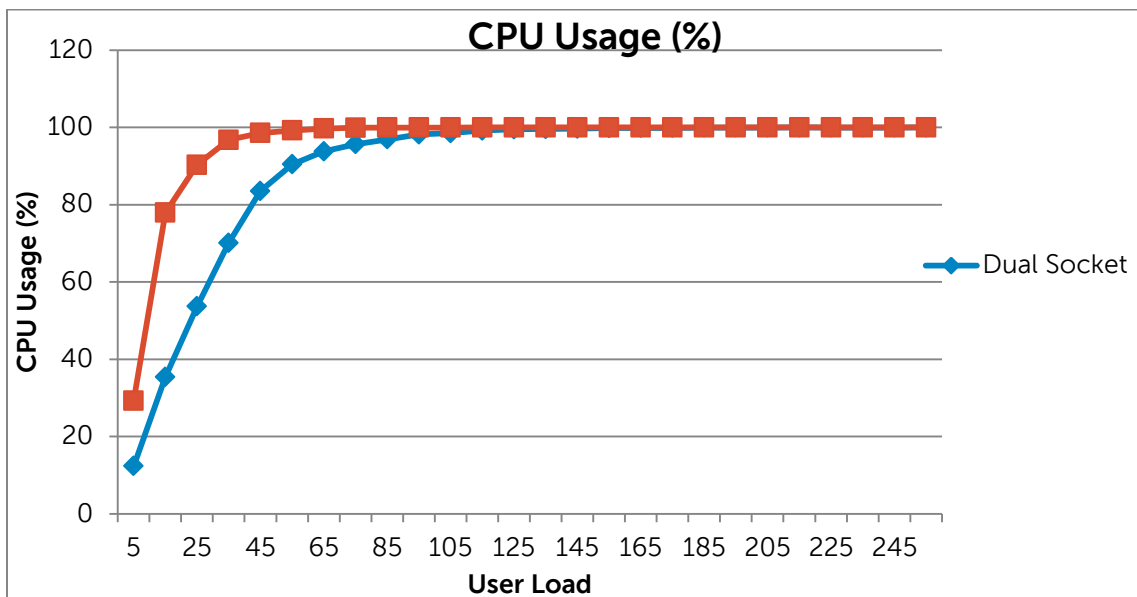


The dual processor R720xd system delivered two times higher throughput with 80 percent improvement in the response time over the single processor system.

NOTE: - In both the configurations, we used a constant memory configuration of 96 GB. You may observe better TPS and response times as you increase the memory. You can allocate maximum of 384 GB for each processor.

Figure 10 shows the CPU usage in for each system configuration tested.

Figure 10. PowerEdge R720xd with single versus dual processor comparison – CPU usage

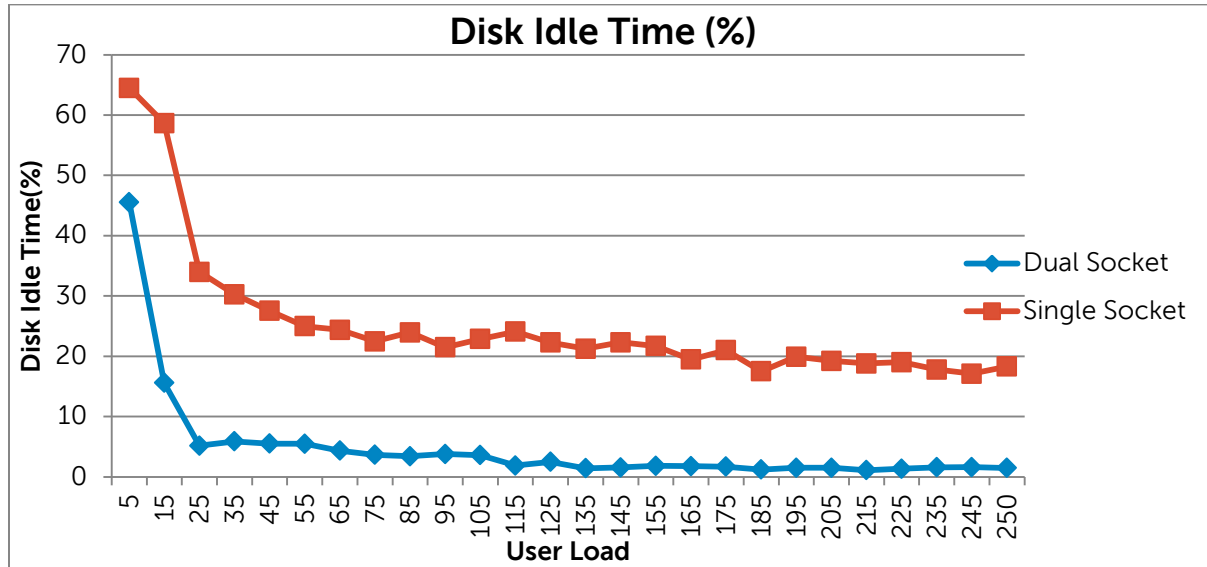




The single processor system reached 100 percent CPU utilization with a user load of 55; whereas the dual processor system reached 100 percent CPU utilization at a user load of 110. This clearly indicates that when the customers scale up from a single to dual processor configuration on the R720xd system they may see up to a 2 times improvement in the overall performance.

Figure 11 shows the disk idle time (%) comparison of the single versus dual processor configurations.

Figure 11. R720xd with single versus. dual processor comparison – disk idle time



At higher user load, the single processor system is bottlenecked at CPU level, and it is not able to take full advantage of the high performing storage backend. Consequently, we observed more disk idle time in the single processor system. On the other hand, in the dual processor system, the CPU bottleneck is alleviated by adding another CPU. In this case, the dual processor system makes complete use of high performing storage, which reduces the disk idle time. The dual processor system to delivers two times the TPS with an 85 percent improvement in the response time.

## Conclusion

Combining the huge processing power, high performing massive internal storage and huge memory support in a single server makes Dell PowerEdge R720xd an excellent single-platform to host core business OLTP-based applications with Microsoft SQL Server databases. With its flexible storage chassis, it can also provide an excellent platform for other business applications using Microsoft Exchange, Hyper-v, and VMware hypervisors.

Especially for OLTP-based SQL database deployments, the R720xd comes with 6 Gbps backplane chassis to support up to 24 hot swappable front loaded 15,000 RPM 2.5-inch SAS drives. In addition, the R720Xd can be easily scaled up in both storage and CPUs. All of this makes the R720xd able to provide an excellent entry-level single server platform to host SQL Server databases.

Table 4 describes the performance benefits that were observed during the SQL Server OLTP Database scalability tests using PowerEdge R720xd.

**Table 4. Performance benefits from PE R720xd scalability**

<b>Configuration</b>	<b>TPS improvement</b>	<b>Response Time Improvement</b>
<b>Full vs. Half populated storage</b>	1.5 times	18%
<b>Dual vs. Single Processor System</b>	2 times	85%

To summarize, with the support for high performing internal Storage, Intel Xeon processors and rich memory configurations, the R720xd would be the optimal choice for a high performance, entry-level SQL Server database deployment. The benefits of using the R720xd for a single server Microsoft SQL Server database deployment include:

- Enables the customer to start with a minimal configuration and scale up as needed by adding more memory, disk drives and an additional CPU
- Improves overall database performance with powerful processors, flexible memory configurations, and large configurable internal storage.