Consolidating Microsoft SQL Server databases on PowerEdge R930 server

This white paper showcases PowerEdge R930 computing capabilities in consolidating SQL Server OLTP databases in a virtual environment.

Dell Engineering
October 2015

A Dell Reference Architecture

Revisions

Date	Description
Oct 2015	Initial release

Table of contents

Executive Summary

Several applications with varying workload behaviors have been hosted on multiple physical servers and this has led to many issues such as server sprawl, high power consumption, increased datacenter footprint, inefficient resource usage. To address these issues, while leveraging the advantages of virtualization technology, IT industry is working toward consolidating several applications on a single, large server. Therefore, IT organizations require servers with dense computing power that enables customers to achieve better consolidation of applications.

Dell has recently launched its highest performing 4-socket server specially designed for mission critical and most demanding enterprise applications such as analytics, In-Memory databases and workload consolidation.

This paper showcases how the PowerEdge R930 server delivers performance under heavy database workload consolidation.

The key observations of the proposed performance study are:

- A linear, scalable and consistent database performance is delivered by the PowerEdge R930 server when several virtual machines, each running specific SQL Server transactional OLTP database workloads, are consolidated on the server.
- More than 25 SQL virtual machines each running a specific transactional OLTP workload can be consolidated on to a single PowerEdge R930 server.
- Reduced licensing costs and increased power and cooling savings can be achieved due to high database consolidation factor provided by the PowerEdge R930 server.

1 Introduction

Virtualizing the workloads is becoming more popular in today's datacenters. Virtualization technology is used to migrate the enterprise applications, such as relational databases, from physical systems to a virtual environment being hosted on powerful servers. To support the growing business requirements and better consolidation factor, there is a constant requirement for powerful servers with huge computing power. In addition, support for large memory footprint, many PCIe slots, server side SSDs will also impact the server selection for the workload virtualization.

Microsoft SQL Server databases are getting popular as relational database backend platform for a wide variety of customer applications. But over a period of time, due to uncontrolled database deployments, database sprawl is becoming one of the common problems for many IT administrators. Hence, IT administrators want to migrate and then consolidate databases used by various transactional applications to a centrally managed and powerful system.

In this paper, we present the Dell PowerEdge R930 server that can be used for consolidating SQL Server databases in a virtual environment, which enables customers to achieve efficient resource utilization, reduce the server sprawl, power and cooling savings and reduce other licensing costs.

1.1 Objective

The scope of this paper is to showcase the computing capabilities, performance and scalability benefits of the Dell PowerEdge R930 server by consolidating Microsoft SQL Server 2014 transactional databases. The paper outlines the reference architecture used for the proposed performance study and leverages the same architecture for consolidating the SQL Server databases.

1.2 Audience

This white paper is intended for IT administrators and architects who are interested in designing and implementing a virtualized infrastructure for consolidating Microsoft SQL Server databases using the Dell PowerEdge R930 server. The reader is expected to have an understanding of Microsoft SQL Server 2014 and benefits of consolidating databases in virtual environments.

2 Dell PowerEdge R930 overview

Dell PowerEdge R930 has been designed and built for performance, scalability, reliability and manageability. It is a 4-socket, 4U rack server ideal for mission critical applications and high performing workloads. The following figures provide more details on the technical specifications of the PowerEdge R930 server:

Figure 1 PowerEdge R930 Front View

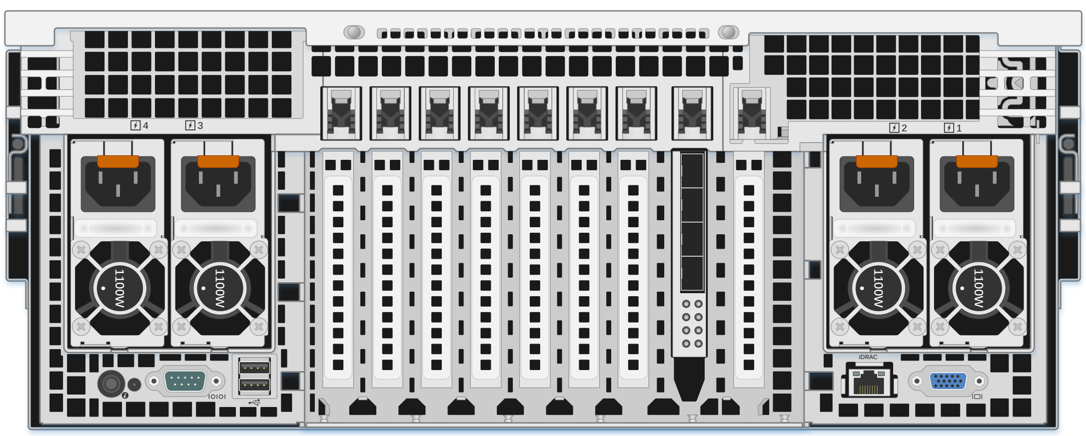

Figure 2 PowerEdge R930 Back View

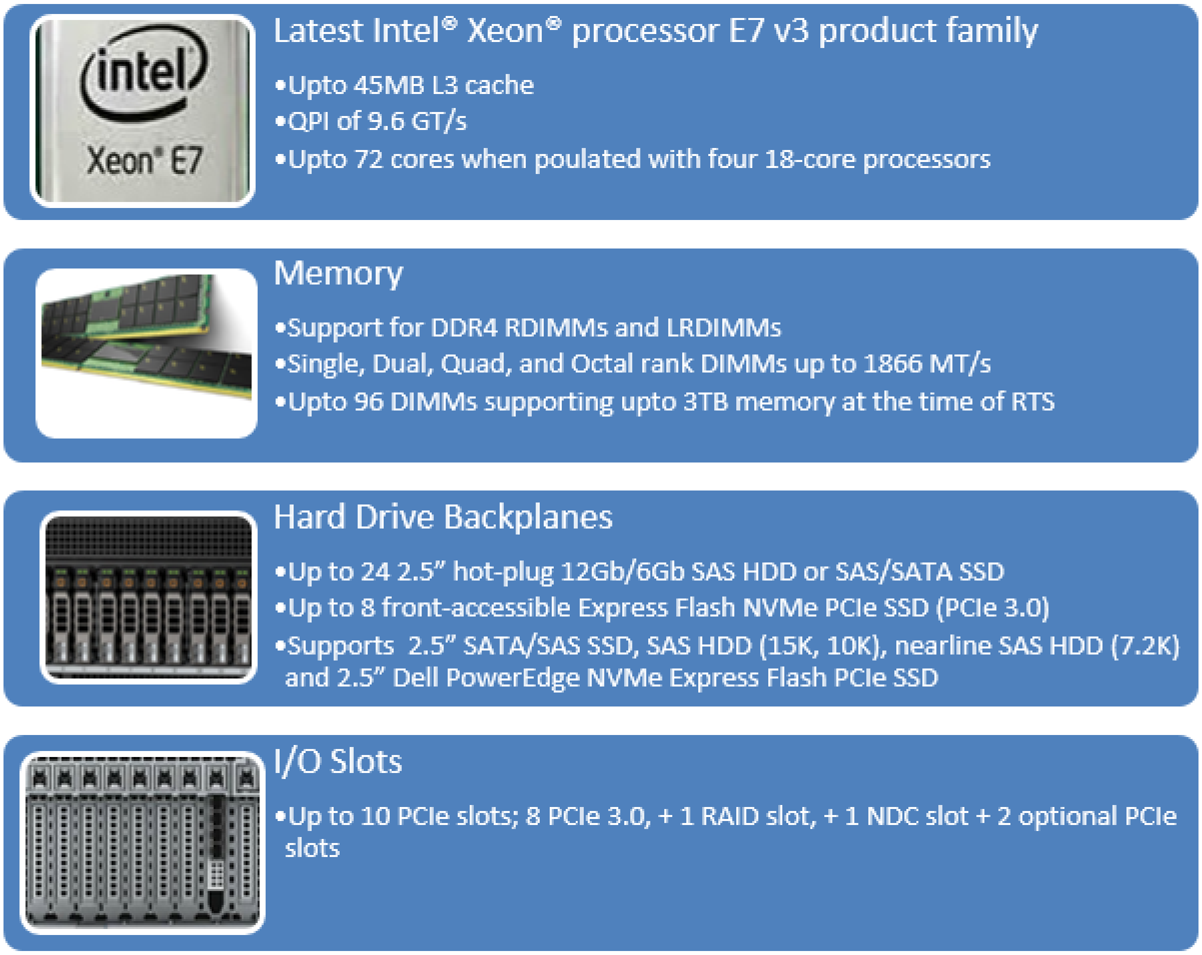

Figure 3 PowerEdge R930 specifications

3 Dell PowerEdge R930 for database consolidation

Several factors impact the selection of a server for database consolidation in a virtual environment. As we consolidate more database virtual machines on the same host, the host must be able to deliver consistent and linear performance within the acceptable response time.

The following sections provide the reasons for selecting the Dell PowerEdge R930 server for the proposed workload consolidation study.

3.1 Huge computing power with latest Intel processors

Powered by the latest Intel processor product family, the PowerEdge R930 server includes more computing cores and more processor cache (L1/L2/L3).

- Database engine can leverage available cores for parallel query processing. The given query tasks are split into multiple sub tasks and executed in parallel on the available processor cores, therefore, achieving faster query execution time.
- The higher processor cache improves query processing while reducing the latencies involved in fetching the data from the main memory.
- When fully loaded, PowerEdge R930 supports up to 72 cores, which enables customers to achieve better consolidation factor in migrating and consolidating many SQL Server databases on to a single host, thereby achieving reduced datacenter foot print and increased power and cooling savings.

3.2 Support for high speed and high capacity DDR4 memory DIMMS

- The capacity and speed of the main memory are crucial to the memory intensive applications such as databases. The DDR4 memory operates at 1866MT/s and provides faster data access to the processors, which leads to the processors being used more efficiently.
- With 96 memory slots, the PowerEdge R930 server supports upto 3TB memory. This massive amount of memory enables the PowerEdge R930 server to support more database virtual machines seamlessly on a single host thereby achieving better database consolidation factor.

3.3 Support for server side SSDs for boosting IO performance

- The PowerEdge R930 server provides multiple disk backplane options offering greater flexibility to the customer while choosing the required disk configuration on the basis of the varying local IO requirements.
- The PowerEdge R930 server supports upto eight 3.2TB NVMe Express Flash PCIe SSDs. These hot pluggable 2.5-inch front loading SSDs provide more IOPS over traditional SSDs with lower latencies.
- The support for NVMe Express flash PCIe SSDs enables PowerEdge R930 to deliver high performing internal storage. This can be used to store the IO intensive SQL server database components such as transaction log files and tempdb files.

3.4 More PCIe slots for scalable networking and storage access

- IO intensive workloads running on the physical host can easily saturate the underlying networking cards. Therefore, it is essential that the host should support enough number of PCIe slots, which enables customers to start with minimum configuration and expand network bandwidth as required.
- The PowerEdge R930server supports up to ten PCIe Gen3 slots in various formats and provides the flexibility to get started with a minimal number of network cards and then scale up as more virtual machines are consolidated on the server. This helps to reduce the network latencies and improve the network bandwidth.

4 Solution Overview

This section describes the solution architecture in detail for the SQL Server database consolidation in a virtual environment.

The following figure illustrates the high-level architecture of the proposed solution:

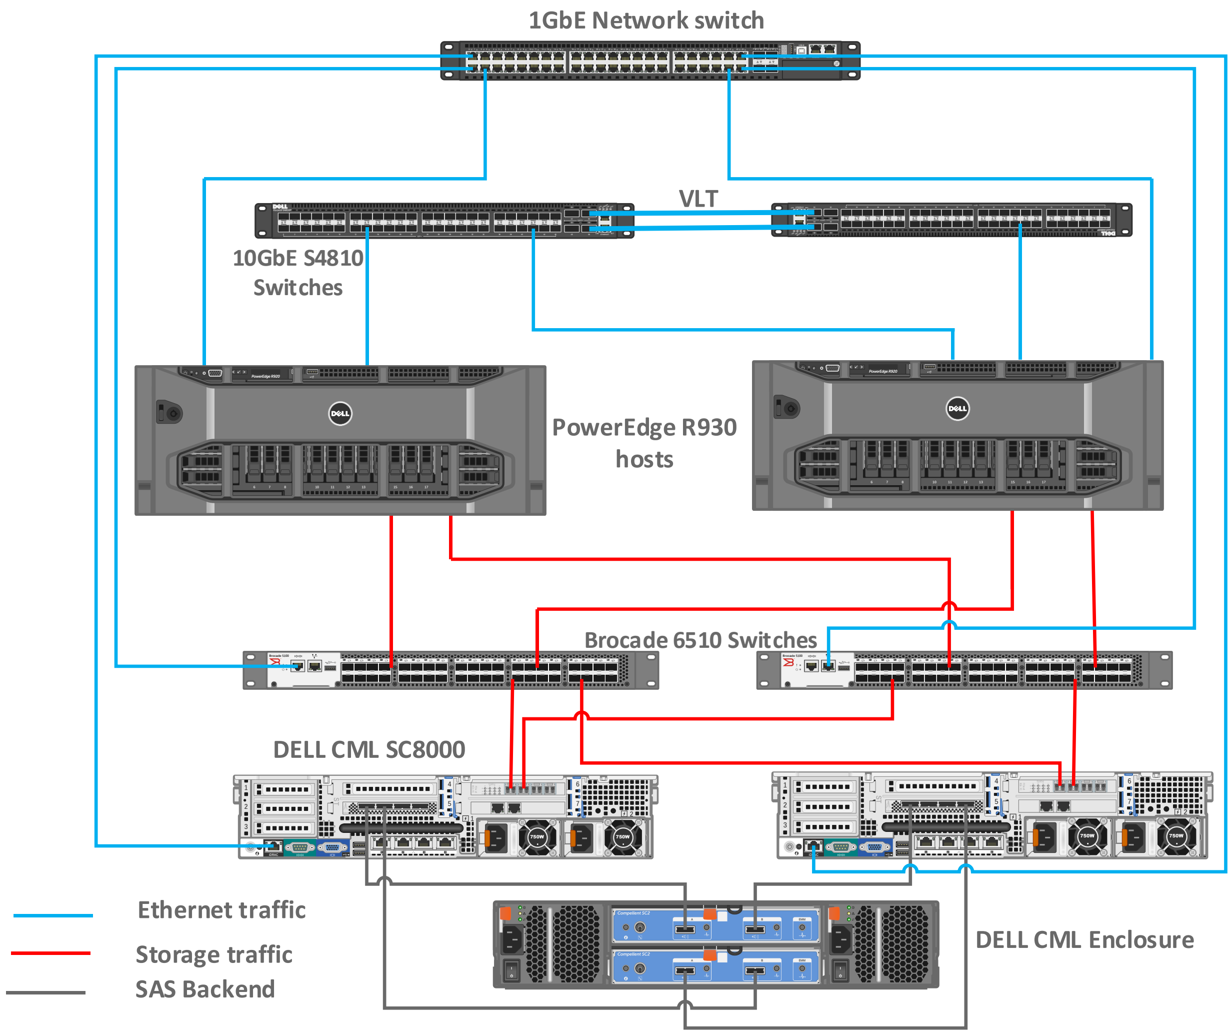

Figure 4 Reference Architecture using PowerEdge R930

In the above reference architecture, two R930 servers are used to provide high availability to the virtual machines running SQL Server databases. Windows Server 2012 R2 has been installed on the 2x2.5-inch 300G internal drives of the hosts. Hyperv-V role has been enabled and two node windows cluster has been created.

4.1 Physical Network Consideration

One on board QLogic 57800 quad port Network Daughter card (NDC) is used on each server. It provides two 10Gb SFP+ interfaces and two 1Gb BASE-T interfaces on a single card. The 1Gb interface(s) is connected to Dell Networking 1GbE switch for Out-Of-Band (OOB) management of the physical server. Two 10G SFP+ interfaces have been used to provide networking access to the server and to connect to two Dell Networking 10GbE Top-Of-Row switches. Windows LBFO NIC teaming has been deployed using these two 10GbE SFP+ interfaces to provide resilient networking connectivity for physical server and virtual machines. The following table describes the settings used to create the LBFO NIC teaming on each host:

Table 1 LBFO NIC teaming settings

Windows Server 2012 R2 team	Configuration
NIC Team Mode	Switch Independent
Load Balancing Algorithm	HyperV Port
Standby Adapter	None
Team Members	Option 1: Two 10G SFP+ interfaces of NDC Option 2: Four 10G SFP+ interfaces; Two interfaces from NDC + Two interfaces from Add-on network Card

Note: To avoid single point of failure on the NDC card, customers can also use additional 10G SPF+ dual port card to create NIC team along with the 10G interfaces of the NDC card. In this case, the NIC team will be created using four 10G SFP+ interfaces: two from NDC and two from additional add-on card.

4.2 Converged Virtual Network Configuration

A converged Hyper-V VM switch is created with minimum bandwidth settings per traffic class for the virtual machine networking. Additional virtual network interfaces have been created for the management server and the Quality of Services (QoS) has been used to configure the minimum bandwidth requirements for each type of traffic. Different type of subnet masks have been used to isolate each type of traffic. The following table provides more details on the virtual network interfaces created and minimum bandwidth requirements for each type of networking:

Table 2 Network Bandwidth requirements

Traffic Class & interfaces	Minimum Bandwidth Weight
Host/VM Management	10
Cluster Private	30
Live Migration	40
DefaultFlowMinimumBandwidthWeight	10

4.3 Storage Requirement and Configuration

In the proposed performance study, our goal is to evaluate the computing capabilities of PowerEdge R930 server through the consolidation of Microsoft SQL Databases. Therefore, it is essential that storage should never become a bottleneck while we are evaluating the computing performance of the server under heavy workloads. Hence, a high performing and low latency Dell Compellent SC8000 flash optimized Fiber

Channel storage is used. The SC220 storage enclosure is equipped with 6x400G Write Intensive (WI) and 6x1.6TB Read Intensive (RI) flash drives to deliver the required storage IOPS and throughput. For more details on the flash optimized Dell SC8000 storage, refer to Dell Compellent SC8000 flash optimized storage.

Two Brocade 6510 16Gbps switches are used to provide the storage connectivity to the virtual machines running on the Hyper-V host. Redundant Brocade switches are used to avoid the single point of failure at the SAN fabric level.

4.3.1 Storage Configuration for Windows Cluster

The following two storage volumes have been created in the storage array and have been provided access to the physical hosts. The following table provides more details on the purpose of the volumes:

Table 3 Storage volume configuration for Windows cluster setup

Volume	Size	Purpose
Quorum	1 GB	For storing the cluster quorum configuration. Configured with a vote to participate in the cluster voting mechanism.
VHD store	1 TB	For storing all .VHDX files and other VM configuration files. Configured with CSV to enable access to all the hosts and for faster fail overs of cluster resources.

4.3.2 Direct Storage Access to Virtual Machines for better performance

The virtual Fiber channel feature of Windows 2012 R2 Hyper-V enables virtual machines to directly access the existing FC storage through virtualized fiber channel adapter. This enables customers to use the existing fiber channel infrastructure to support the virtualized workloads. Virtual fiber channel includes support for other related features such as guest cluster, live migration.

In the proposed solution, two QLogic QLE2662 dual port 16Gbps HBAs are used on each Hyper-V host. On each Hyper-V host, two Fiber channel virtual SANs have been created using two ports one from each HBA such that in case of any HBA failure, virtual machines will still have access to the storage through other HBA ports. Each SQL Server VM is configured to get two virtual FC ports one from each Fiber channel virtual SAN. For each SQL VM virtual FC adapter, a separate zone has been created on each fiber channel switch grouping the WWNs of the corresponding SQL VM and WWNS of the Dell Compellent storage controllers. The following figure shows the configuration used to enable direct storage access to the virtual machines:

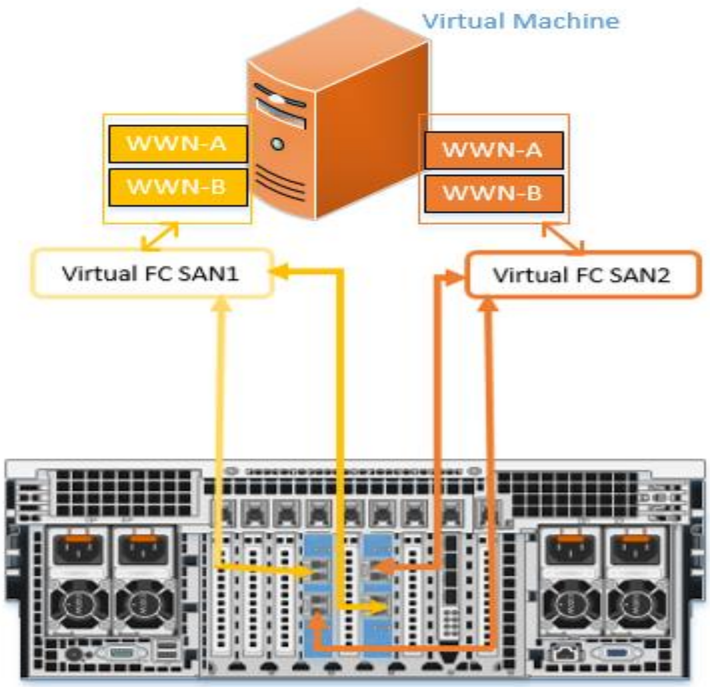

Figure 5 Virtual Fiber Channel SAN

Because the SQL virtual machines are enabled to directly access the SAN storage, dedicated volumes have been created for each virtual machine as described in the following table:

Table 4 Storage Volume configuration for SQL Virtual Machines

Volume	Size	Purpose
SQL_DATA	500GB	For storing test database data files and additional TempDB Data files
SQL_LOG	100GB	For storing test database TransactionLog files

5 Test Configuration

The following table describes the hardware and software components used to build the test configuration for evaluating the computing performance of the Dell PowerEdge R930 server by consolidating SQL Server databases running transactional workload:

Table 5 Hardware and Software components specifications of the RA

Traffic Class & interfaces	Minimum Bandwidth Weight
Server	Two Dell PowerEdge R930 servers
CPUs	Four Intel® Xeon® E7-4850 V3 14C, at 2.20GHz on each host
Total Cores / Logical threads	56/112
Memory	512GB with 32 x 16GB DIMMs running at 1866 MT/s on each host
Network Card	One QLogic 57800 quad port Network Daughter card (NDC) on each host
Host Bus Adapters	Two Dual Port QLogic 2662 16Gbps cards on each host
Management and Network Switches	One Dell Networking S60 1GbE switch and Two Dell Networking S4810 10GbE switches with VLT configured on 40GbE ports
SAN switches	Two Brocade 5610 16Gbps switches
Storage Details	Dell Compellent Flash Optimized Storage
Windows	Windows 2012 R2 Datacenter Edition

The following BIOS setting have been used on each host in the proposed reference architecture:

Table 6 BIOS settings

BIOS Setting	Details
Virtualization Technology	Enables virtualization of the available resources on the host
HyperThreading (HT)	Turns each physical core in to two logical cores; efficient for parallel execution
System Profile: Default Performance-per-watt (DAPC)	Enables efficient power usage of the host

5.1 Test methodology

The objective behind the test was to derive the maximum number of virtual machines that can be run simultaneously on a single PowerEdge R930 server platform and observe how R930 behaves and delivers the consistent performance as we scale up the SQL virtual machines.

The following table describes the virtual machines configuration running SQL Server 2014 database instance:

Table 7 SQL Virtual Machine configuration

Traffic Class & interfaces	Minimum Bandwidth Weight
VCPUs	#6
Memory	40GB (static)
Database Size	350GB data set
Concurrent Users	1500
Vitrual Network interfaces	1
Virtual Fiber Channel ports	2
Operating System	Windows 2012 R2 standard Edition
Database	Microsoft SQL Server 2014 Enterprise Edition
Workload type and Benchmark tool used	TPCC-C OLTP workload style, Dell Quest Benchmark Factory (with 1/20th reduction in Keying and Think times)

On each virtual machine, Windows 2012 R2 Standard Edition and SQL Server 2014 Enterprise Edition have been installed. Each virtual machine is configured with one virtual network port for network access and two virtual fiber channel ports for direct storage access.

A test database of size 350GB is generated and restored on the storage volumes of the SQL virtual machines. Separate storage volumes are used to store the database data and log files. TempDB is configured with multiple data files, and “Max Worker Threads” has been set to 1024 as we observed workerthread starvation during the preliminary workload tests.

To simulate the real time OLTP database on each SQL VM, TPC-C style of workload is selected to simulate the real OLTP workload sql queries. Each SQL VM is stressed with OLTP workload using a dedicated client machine running Dell Quest Benchmack Factory tool. This tool is used to simulate TPC-C style of OLTP workload for databases.

6 Observations and results

The following figure shows the PowerEdge R930 CPU utilization trend as we run multiple SQL virtual machines in parallel. Each SQL virtual machine is stressed upto 75-80% CPU utilization with OLTP database workload.

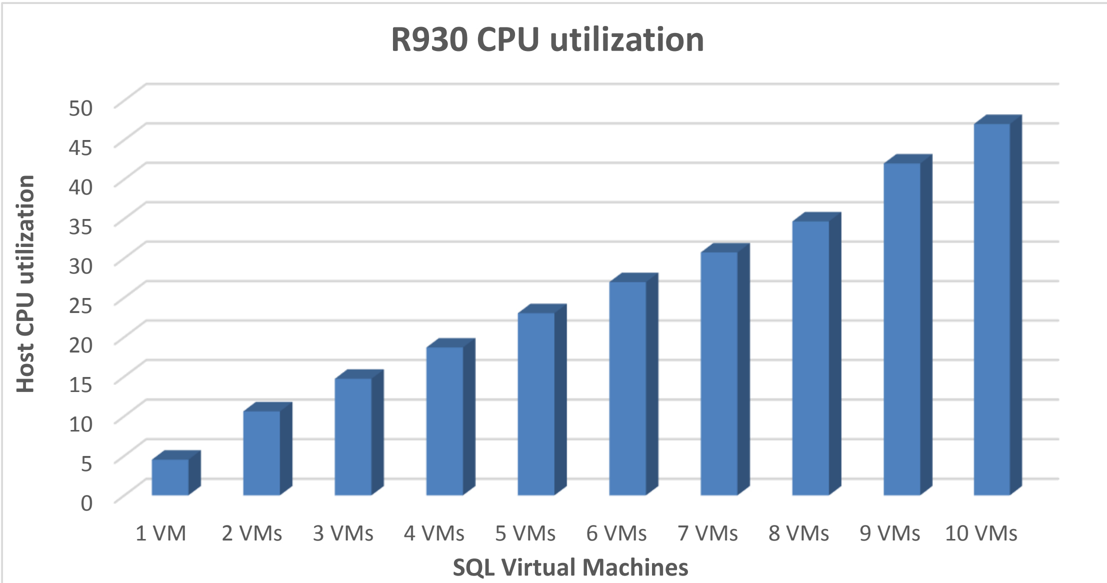

Figure 6 CPU utilization trend of PowerEdge R930 host

As shown in the preceding figure, each instance of SQL VM has resulted in around 4.5% CPU utilization on the R930 host and only around 48% CPU is utilized when the single R930 host is stressed with ten simultaneous database virtual machines. There is still enough space to accommodate more virtual machines.

From the CPU utilization trend of ten SQL VMs on single R930 host with 56 physical cores, it can be extrapolated to be said that the proposed reference architecture using two PowerEdge R930 hosts (each with 56 physical cores) can support up to 20 highly available SQL virtual machines each running specific OLTP database workload. If any R930 host is not available for any reason, the other R930 host will be able to host all the twenty SQL VMs.

Table 8 Number of SQL VMs supported on a single R930 system

When R930 system	# VMs supported
Configured with 4x 14-core CPUs (56 cores)	A single R930 can support up to 20 SQL VMs
Configured with 4x 18-core CPUs (72 cores)	A single R930 can support more than 25 SQL VMs (theoretical/extrapolated)

The following figure describes how the PowerEdge R930 delivered consistent database performance as we scale up the number of virtual machines.

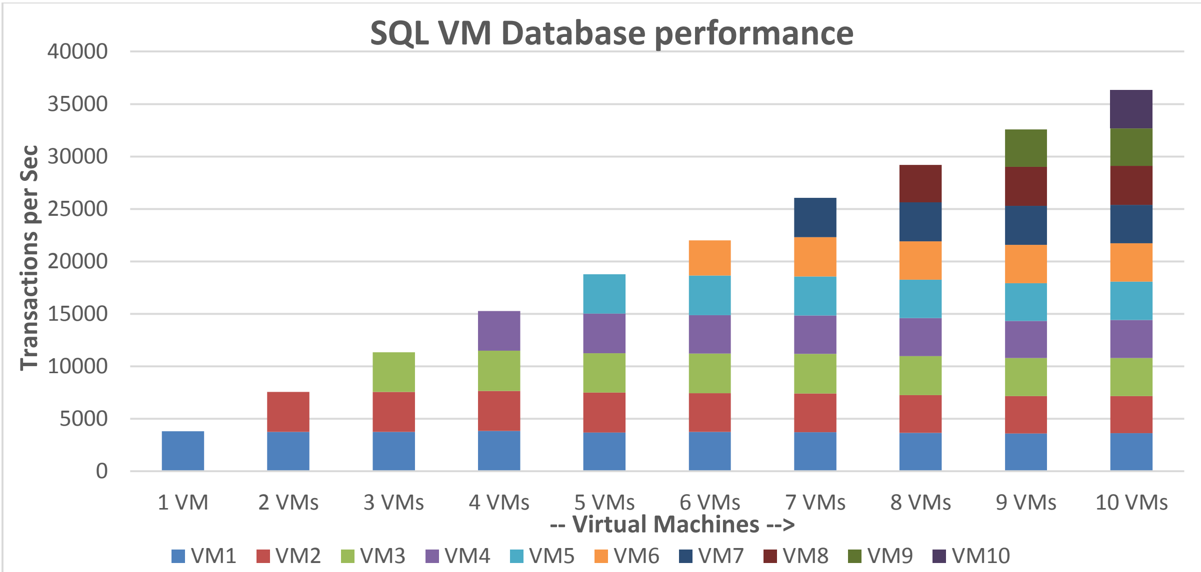

Figure 7 SQL Database Transactions Per Second (TPS)

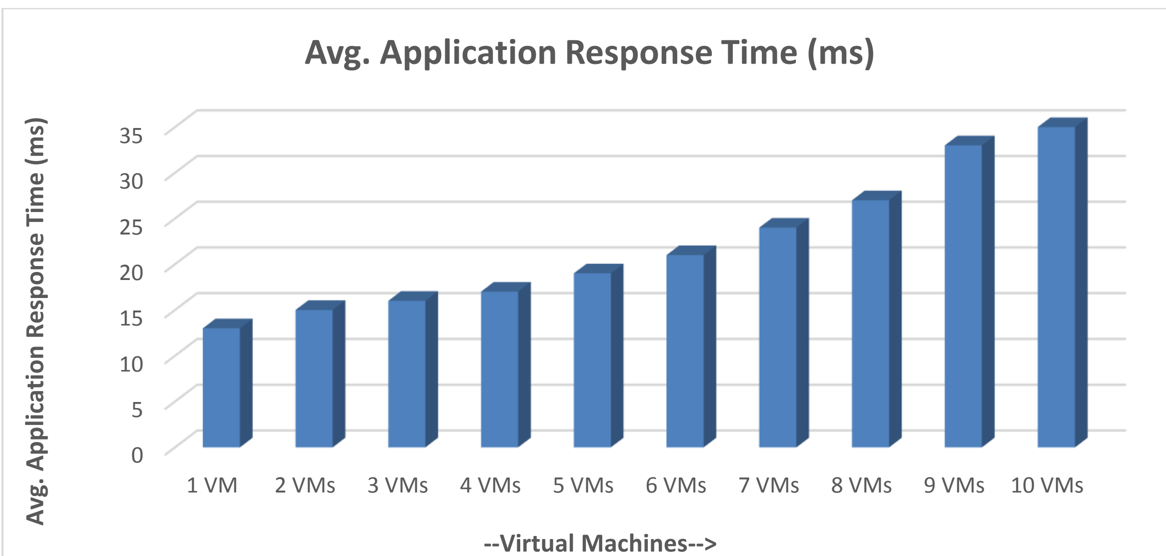

Figure 8 Application Response Time (ms)

Each guest was observed to deliver around 3700 Transactions Per Second (TPS) within the acceptable response time. The same trend was observed from each guest as more guests were run simultaneously on a single R930 host. When 10 SQL VMs are hosted on a single R930 host, the average response time of each VM was around 35ms.

7 Conclusion

The reference architecture using PowerEdge R930 described in this paper provides the best-in-class platform for consolidating critical OLTP database workloads. The PowerEdge R930 server provides huge computing power and features highly scalable processing, memory and internal storage, which are essential to host mission critical workloads.

The key take aways of the proposed solution are:

- Dell PowerEdge R930 is purpose-built and power packed with enormous computing resources and provides a good platform for consolidating mission critical workloads.
- A single fully loaded PowerEdge R930 server can accommodate and run more than 25 virtual machines each running enterprise class SQL Server databases.
- A linear, scalable and consistent database performance is delivered within the acceptable response time by a single PowerEdge R930 server when multiple SQL virtual machines are hosted on it.

A Additional resources

Support.dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com is an IT Community where you can connect with Dell Customers and Dell employees for the purpose of sharing knowledge, best practices, and information about Dell products and installations.

Referenced and recommended Dell publications:

- Dell SQL Server sizing advisor
  http://www.dell.com/sql
- Dell PowerEdge R930 technical specifications
  http://www.dell.com/learn/us/en/08/shared-content~data-sheets~en/documents~poweredge_r930_spec_sheet-final.pdf
- Dell Compellent SC8000 Spec sheet
  http://partnerdirect.dell.com/sites/channel/Documents/Compellent_Storage_Center_Spec_Sheet.pdf

Referenced or recommended Microsoft publications:

- SQL Server consolidation Guidance
  https://technet.microsoft.com/en-us/library/ee819082(v=sql.100).aspx