

Dell Virtualization Solution for Microsoft SQL Server 2012 using PowerEdge R820

This white paper discusses the SQL server workload consolidation capabilities of Dell PowerEdge R820 using Virtualization.

Vinod Kumar

Dell™ | Database Solutions Engineering



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.

© 2011 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 Windows 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.

July 2012 | Rev 1.0

Contents

Executive summary.....	4
Introduction	4
Objectives.....	4
Microsoft Hyper-V Overview	5
System Center Virtual Machine Manager.....	5
SQL Server 2012 Virtualization overview	6
SQL Server 2012 Licensing.....	6
Why Dell™ Server & Storage Hardware	7
Dell PowerEdge R820 as a Virtualization Platform.....	7
Dell SQL Server 2012 Reference architecture using PowerEdge R820.....	7
Single Server Reference Architecture	7
Highly Available Reference Architecture	8
Performance and Scalability of Virtual Machines achieved	9
Test Methodology	10
Results and Analysis	12
Conclusion	15
Appendix A.....	17

Tables

Table 1. Test Configuration Details	10
Table 2. Virtual Server Configuration Details	11
Table 3. Observed Consolidation Capability of R820 for the specific OLTP workload	16
Table 4. Hyper-V performance counters(Medium VMs).....	17
Table 5. Hyper-V performance counters(Large VMs).....	17

Figures

Figure 1. SVCMM.....	5
Figure 2. Single Server Reference Architecture	8
Figure 3. HA Reference Architecture	9
Figure 4. SQL Server Storage Layout.....	12
Figure 5. Database Transactions per Second	13
Figure 6. Database Batch Requests per Sec	13
Figure 7. PowerEdge R820 CPU Utilization	14
Figure 8. Impact of Hyper-threading on Large VMs	15

Executive summary

Dell PowerEdge 12th generation servers provide the robustness and reliability for a highly efficient database consolidation platform. You should choose Dell PowerEdge R820—a powerful four socket enterprise server—as an optimal platform for database consolidation because of its robustness, reliability, and performance characteristics.

This white paper describes the advantages of consolidating Microsoft® SQL Server® database application workloads to a virtualization environment using Microsoft Windows® Server® 2008 R2 Hyper-V™ and Microsoft SQL Server 2012 on dell PowerEdge R820. It demonstrates that Dell PowerEdge R820 is capable of delivering the performance, consolidation, and scalability needed to run SQL Server workloads in a virtualized environment.

Introduction

One of the major challenges organizations face is the challenge of ensuring the operating efficiency of the data center. With the emergence of powerful servers, there is a large possibility of not optimally utilizing the hardware capability for a specific workload. In such cases, it is expensive to run workloads on a standalone machine. In addition to the inefficient operation, it leads to several other challenges like over power consumption, manageability, increased data center space usage etc., which becomes even more challenging in scenarios when increasing data growth slows down the system performance and requires physical expansion of the configuration.

Organizations are looking for deploying virtual servers to reduce IT cost and increase data center performance. Microsoft SQL Server is a popular database product that is widely used by business applications. Consolidating servers through virtualization helps us utilize the resources appropriately; eases deployment, maintenance, power saving, reliability, and availability.

The adoption of virtualization in enterprise has increased flexibility in deployment and life cycle management of applications. IT professionals deploy and use virtualization to consolidate workloads and reduce server sprawl. Additionally, they deploy virtualization with clustering technologies to provide a robust IT infrastructure with high availability and quick disaster recovery.

This paper outlines how Dell PowerEdge R820 and Microsoft Hyper-V may be used to deploy multiple SQL Server 2012 databases on a virtualized environment without compromising on the performance.

Objectives

The objectives of this technical paper include:

- Understanding Dell PowerEdge R820 server capabilities to run SQL workload in a virtualized environment.
- Testing to see how many medium and large SQL VM's can be consolidated using R820.
- Using best practices to run SQL server workload on a virtualized environment.

Microsoft Hyper-V Overview

Microsoft Windows Server 2008 R2 SP1 Hyper-V builds on the architecture and functionality of Windows Server 2008 Hyper-V by adding multiple new features that enhance product flexibility.

Some of the key features that are available in Microsoft Hyper-V Server 2008 R2 SP1 include:


- Dynamic memory—takes all memory of the Hyper-V host and distributes it across all the Virtual Machines.
- Live migration—allows you to move running virtual machine from one Hyper-V host to another Hyper-V host within the same failover cluster without a perceived downtime.
- Cluster shared volume support—all the nodes that are part of the failover cluster., can access cluster shared volumes (CSV). Different nodes that are part of failover cluster can host a virtual machine on the same volume.

One of the main highlights of Microsoft Hyper-V is the integration that it provides to the Microsoft System Center Group of products, which brings ease and flexibility of managing the virtual environments.

System Center Virtual Machine Manager

Microsoft System Center is a comprehensive IT infrastructure, virtualization, and cloud management platform. With System Center, you can more easily and efficiently manage your applications and services across multiple hypervisors as well as across public and private cloud infrastructures to deliver flexible and cost-effective IT services for your business.

Figure 1. SVCMM



As shown in figure 1, SCVMM is installed on a dedicated server to manage two hypervisor hosts. In this experiment, we have used SCVMM to deploy a fully configured SQL virtual machine to hyper-v host.

System Center Virtual Machine Manager makes the deploying of SQL virtual machines a very simple process. SCVMM creates the template from a fully configured SQL VM; the created template is stored in its library. This template can be deployed multiple times and anywhere in the network. We have seen deploying fully configured virtual machine in less than 20 minutes.

SQL Server 2012 Virtualization overview

Microsoft SQL Server 2012 builds on the capabilities that were offered in sql server 2008. It is Microsoft's first cloud-ready information platform that will help organizations unlock breakthrough insights across the organization and quickly build solutions to extend data across on-premises and public cloud, backed by mission critical confidence.

SQL Server 2012 added new features, which include the following:

- High availability with AlwaysOn feature
- Improved query performance with ColumnStore Index
- Rapid data discovery via Power View and PowerPivot
- Credible, consistent data via BI Semantic Model and Data Quality Services
- Scale on demand from devices to datacenter to cloud

SQL Server 2012 Licensing

Microsoft SQL Server is increasingly being deployed in virtualized environments, which enable running instances of SQL Server concurrently in separate virtual machines.

SQL Server 2012 offers expanded virtualization rights, options and benefits to provide greater flexibility for customers deploying in virtual environments.

Customers have the choice to license SQL Server 2012 when it is deployed on a virtual environment.

- Licensing an individual SQL Server 2012 Virtual machine:
 1. When licensing an individual virtual machine using per core licensing model, the customer must purchase the core license for each virtual processor.
 2. When licensing an individual virtual machine using Server + CAL licensing model, the customer must purchase one server license for each virtual machine.
- Maximize Virtualization by licensing the entire physical server: Licensing all the physical cores on server can run SQL virtual machine equal to physical cores.

For example, a four processor server with four cores can run sixteen SQL Virtual machines.

Visit <http://www.microsoft.com/sqlserver/en/us/get-sql-server/how-to-buy.aspx> for more information on licensing.

The remainder of this white paper discusses the sample SQL Server 2012 database reference configurations that are possible on Dell PowerEdge servers using Microsoft Hyper-V as the virtualization platform.

Why Dell™ Server & Storage Hardware

Dell 12th generation servers offer an efficient infrastructure for hosting a virtualized platform. Microsoft SQL Server 2012, along with its robust HA and DR capabilities, offer a high performing and reliable software stack good for consolidation scenarios.

Dell PowerEdge R820 as a Virtualization Platform

PowerEdge R820, with the multi-core Intel® Xeon® processor E5-4600 product family support, may be beneficial in delivering outstanding performance for server. The PowerEdge R820, with 2-socket and 4-socket configuration options, can feature up to 32 cores of latest generation server-class processing power. This server supports up to 16 TB internal capacity and integrated PCIe Gen3-capable expansion slots for extra capacity and functionality. Plus, up to two internal RAID controllers allow a range of RAID levels for improved storage reliability, and the front-access PCIe SSDs enable performance-enhancing in-box storage tiering.

Up to 1.5TB of highly scalable memory (48 DIMM slots) lets the R820 efficiently process extremely demanding workloads, such as:

- Enterprise resource planning (ERP)
- Databases and data warehousing
- eCommerce
- Virtual desktop infrastructure (VDI)
- Data-node high-performance computing (HPC)

Visit <http://www.dell.com/us/enterprise/p/poweredge-r820/pd> for more information on Dell PowerEdge R820 Server.


Dell SQL Server 2012 Reference architecture using PowerEdge R820

The sections below discuss the sample single server and multi-server virtualized reference architectures that are possible using Dell PowerEdge R820, SQL Server 2012, and Hyper-V.

Single Server Reference Architecture

Figure 1 shows an example virtualized database reference configuration using PowerEdge R820, Hyper-V, and SQL Server 2012. In the proposed reference configurations and for the follow-on experiments of this white paper, we have used PowerVault MD3620F as the SAN storage, QLogic QLE2562 as the storage interface and Brocade 5100 as the SAN switch.


Figure 2. Single Server Reference Architecture



Highly Available Reference Architecture

As shown in figure 2, we have configured High available virtual machines using Hyper-V Cluster. This enable us to migrate the virtual machine between the host servers with a live migration feature with negligible downtime.

Figure 3. HA Reference Architecture



The two Dell PowerEdge R820s are members of a Hyper-V failover cluster. Each of the VMs would be active on one of the R820s. If the active node fails, all the VMs will failover to the passive node.

For this exercise, we are using pass-through disks for virtual machines. During live migration, the LUN (volume) disconnects and connects back to the virtual machine on the other node. We could observe a slight delay 2 to 5 seconds to get the virtual machine running on the other node. This delay is experienced only when pass through disks are used.

This delay may be avoided using iSCSI storage or fixed virtual hard disk for SQL Data files. This enables virtual machine to migrate between multiple hosts. In addition, this also allows configuring fail over cluster at the Virtual machine level (as an alternative to Hyper-V clustering).

Refer to <http://social.technet.microsoft.com/wiki/contents/articles/440.hyper-v-how-to-add-a-pass-through-disk-on-a-failover-cluster-en-us.aspx> for more information on configuring pass-through disks to virtual machines.

Performance and Scalability of Virtual Machines achieved

We conducted several experiments to analyze the database consolidation capabilities of R820. The tests were focused on the impact on the SQL Server performance due to Virtual Machine consolidation to a single R820 server.

Test Methodology

The test scenarios were implemented to showcase the database consolidation capabilities of Dell PowerEdge R820. Guest virtual machines were configured on a PowerEdge R820 server and Individual SQL Server instances were deployed on each of the Guest VMs. The performance effects of having additional VMs were analyzed using standard OLTP database workloads (TPC-E). TPC-E was chosen to be the test workload because of its CPU intensive nature of the simulated database transactions.

In the test configuration, PowerVault MD3620F was used as the external storage to host the database files. The storage (and the overall configuration) was sized appropriately to reduce the disk IOs to relieve the storage level stress on the configuration. This helped us to make sure that each of the VMs is able to contribute efficiently to the overall R820 server utilization. Table 1 shows the hardware and software component details used for the VM scalability tests.

Table 1. Test Configuration Details

Components	Details
Server Model	PowerEdge R820
Processor	4* Intel® Xeon® Processor E5-4620 @2.20 GHz, 4 Cores per socket; Hyper threading disabled
Memory	256 GB (16x 16GB, DDR3 @ 1600 Mhz)
HBAs	2* QLogic QLE2562 8 Gbps Dual port HBAs
Multipathing Software	Dell MD Series DSM
Storage(FC)	MD3620F, with two MD1220 Expansion Enclosure
FC Switch	Brocade 5100
Network Switch	Dell Power Connect 6348
Virtualization	Hyper-V
Operating System	Windows 2008 R2 SP1 Enterprise Edition
Database Benchmark	TPC-E benchmark(OLTP)

Test VM configuration for SQL Server

It is important to configure Virtual machines to give better performance when you deploy high I/O extension workload on them. For the VM scalability experiments, the below VM configuration (Table 1) was chosen to be the standard (baseline) configuration for the guest SQL server deployments.

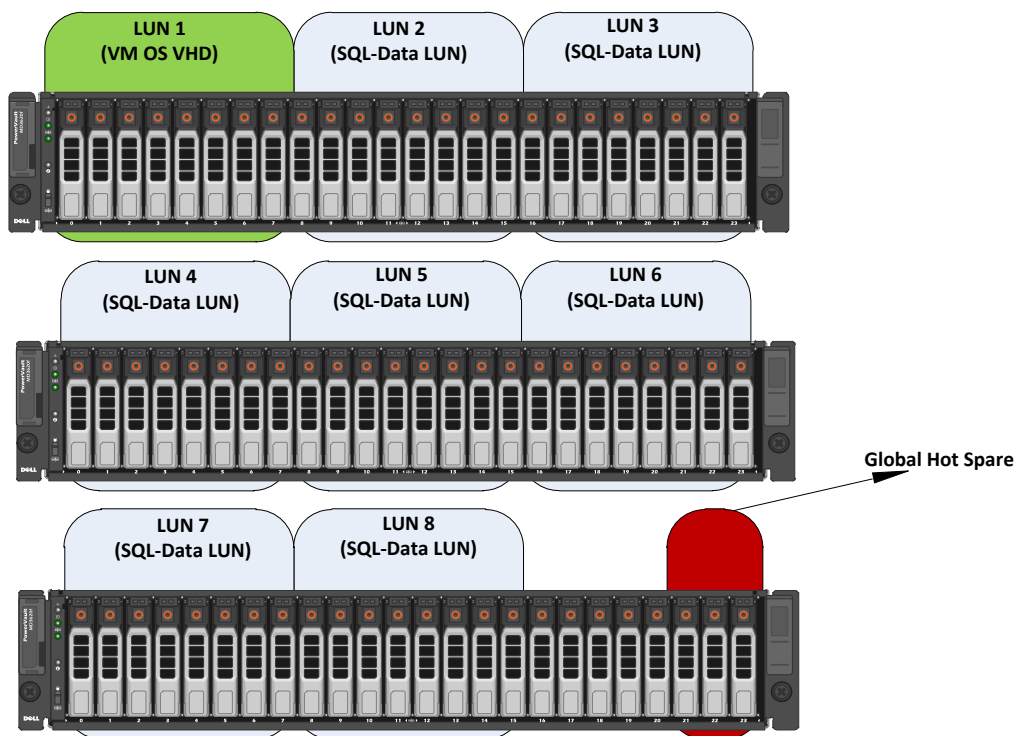
Table 2. Virtual Server Configuration Details

Component	Medium VM	Large VM
Processor	2VCPU	4VCPU
Memory	24 GB	24 GB
Storage	OS: 100 GB VHD Data:1 TB(RAID 10)	OS: 100 GB VHD Data:1 TB(RAID 10)
Operating System	Windows 2008 R2 SP1 Standard Edition	Windows 2008 R2 SP1 Standard Edition
Database size	50 GB	50 GB
Database Software	SQL Server 2012 Standard Edition	SQL Server 2012 Standard Edition
Users	12	24

SQL server enterprise edition may also be used in the guest VMs, in case the customers need additional features like database encryption, database availability features, etc.

Figure 4 shows the storage layout between the VMs that we have used for the VM scalability tests.

Figure 4. SQL Server Storage Layout



As shown in Figure 4, we have created a dedicated disk group of 8 disks (RAID 6) for hosting the VM operating systems; separate fixed size VHDs were created out of the disk and assigned to the individual VMs.

For hosting the database files, we have created dedicated disk groups for each of the VMs on the PowerVault MD3620F storage. Each SQL Database LUN comprised of 8 disk drives configured using RAID 1+0. Using dedicated database LUNs ensures minimal impact on the database performance due to the multiple VMs being stressed at the same time. The databases LUNs were configured as pass-through disks to help reduce the I/O overhead further on the virtual machine.

We observed that each of the medium and large baseline VMs delivered an Avg. Response Time (sec) of 0.007 and 0.015 for a user load of 12 and 24 at average processor utilization of around 80 to 90%. The average disk reads/sec for the VMs was around 0.008. For the further result comparisons, we would be using the user load of 12 & 24 as the evaluation point.

Results and Analysis

Figure 5 and 6 shows the database transactions per second and SQL Server batch requests/sec achieved on a single R820, for medium and large VMs.

Figure 5. Database Transactions per Second

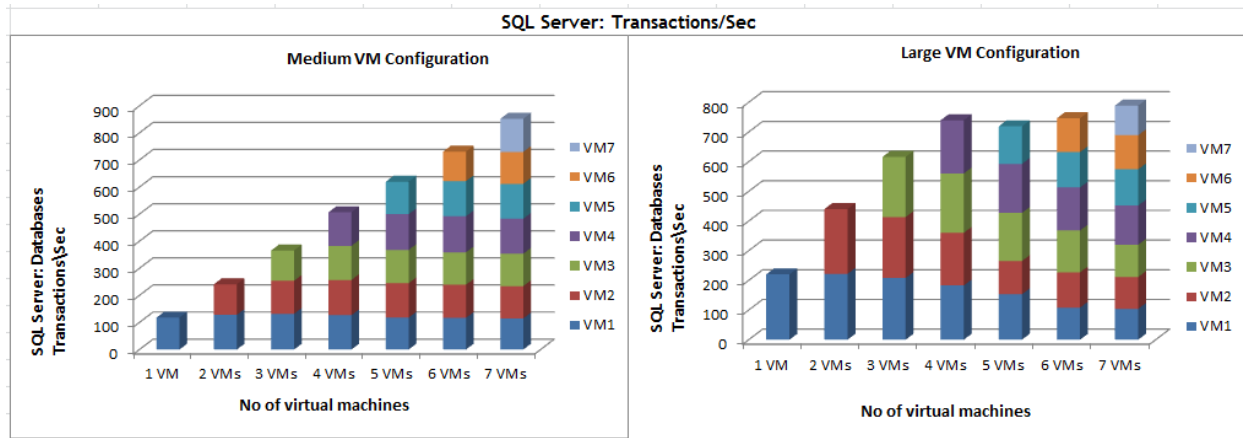
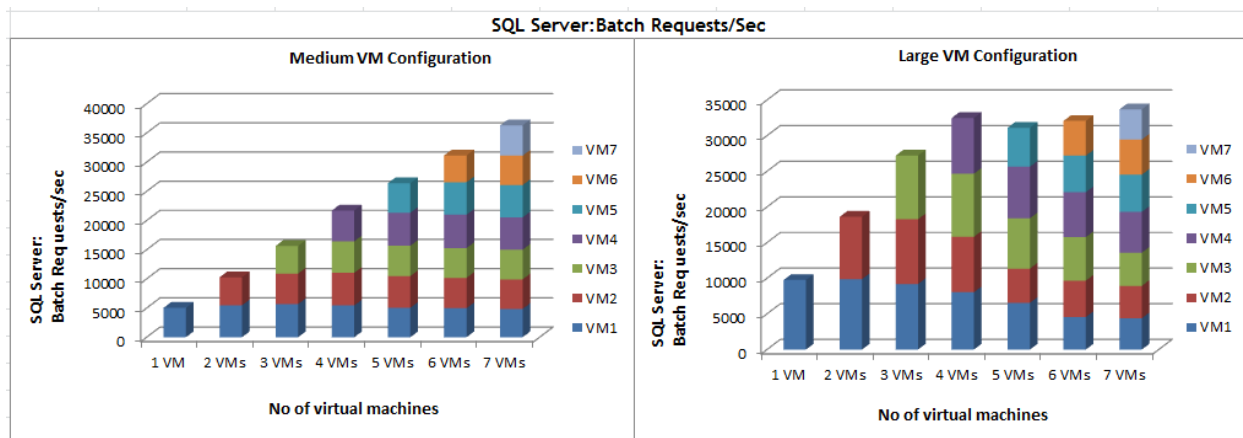


Figure 6. Database Batch Requests per Sec



Medium VMs: We observed a more or less linear scalability of SQL transactions per second (TPS) and SQL batch requests per second on adding multiple VMs. The observed average TPS per virtual machine was around 118, 120, 122, 126, 124, 121 and 121 for 1 VM, 2 VMs, 3 VMs, 4 VMs, 5 VMs, 6 VMs and 7 VMs respectively. Similar trend observed in SQL batch requests per second counter as well.


We could not find any major performance impact on the virtual machines on increasing the number of Virtual machines and running similar workloads on different virtual machines simultaneously.

Large VMs: We observed a more or less linear scalability of SQL transactions per second (TPS) and SQL batch requests per second on adding 1 to 4 VMs; however, adding 5 to 7 VM's resulted in slight performance degradation.

The performance impact was seen due to high processor utilization on the host machine; the observed processor utilization was around 84.9, 93.1, 96.6 and 98.3 for 4VMs, 5VMs, 6VMs and 7VMs respectively.

Figure 7 depicts the Hyper-V host CPU utilization observed on scaling up the VMs on the R820

Figure 7. PowerEdge R820 CPU Utilization



Measuring the processor utilization counter “Processor(*)\%processor time” on the host may not give the accurate CPU utilization results, to measure the total processor utilization on the host machine “\Hyper-V Hypervisor Logical Processor(_Total)\% Total Run Time” counter being used.


Medium VMs: We observed a linear scalability in the CPU utilization as well with the increasing number of database VMs.

Large VMs: Observed linear scalability on adding 1 to 4 large VM's, adding more large VMs impacted overall virtual machine server performance.

We are using 16 physical cores in this configuration, running 4 VM's with each having 4 virtual processors occupying all of the 16 cores. Adding more than 4 VM's with 4 virtual processors would actually sharing the same 16 physical cores; hence, we could see slight performance impact.

Please note that for the above test results, we have enabled only 4 Cores per processor. The hyper threading (HT) feature is disabled on the R820 host machine. On enabling processor HT, we could see that the R820 could accommodate more VMs, with lesser performance impact as compared to the HT disabled configuration.

Figure 8. Impact of Hyper-threading on Large VMs



Until 4 VMs, the performance numbers are similar between the HT enabled and disabled configurations. Beyond 4 large VMs, Figure 8 shows that the batches per sec delivered per VM are better when we have HT enabled in the system. With 7 Large VMs, the server processor utilization was around 86.1% and 98%, on HT enabled and disabled respectively.

Visit <http://msdn.microsoft.com/en-us/library/cc768535%28v=bts.10%29.aspx> for more information on measuring Hyper-V performance.

The Hyper-V counters collected during SQL workload tests are given in Appendix A. It illustrates the difference in the Hyper-V performance counter while running the VM1 to VM7 concurrently.

Conclusion

Dell PowerEdge R820 is a powerful server platform to accommodate a virtual environment due to its large memory support, I/O options and high processing power. The support for the latest Intel® Xeon® processor E5-4600 product family and the high speed(1600Mhz) memory options make it a high performing platform too.

The experiments carried out as part of this white paper showed that the virtual machines scaled up without any major performance impact until up to 85% to 90% of the overall server utilization. Using a 4 processor cores per socket (HT disabled), a single R820 was able to scale up to 4 large virtual machines (each with 4 vCPUs). At the same time, it could accommodate 7 Medium virtual machines (each with 2 vCPUs) without any major impact in the performance.

We observed a performance degradation of individual VMs once the entire server utilization was beyond around 90%. This observation makes it obvious to have a constant monitoring of the overall utilization of the virtualized system and sizing it optimally to meet the workload requirements.

Extrapolating the scalability test results (HT disabled), we can summarize the database consolidation capability of PowerEdge R820 for the specific OLTP workload as given in Table 3.

Table 3. Observed Consolidation Capability of R820 for the specific OLTP workload

Server configuration		Number of Medium VMs that may be consolidated	Number of Large VMs that may be consolidated
Sockets Populated	Processor cores per socket		
4	4	8	4
4	8	16	8

To summarize, SQL Server OLTP workloads can be consolidated onto Dell PowerEdge R820 without any major compromise to the database performance with proper sizing. This enhances the overall cost effectiveness and efficiency of the database environment.

Appendix A

Table 4. Hyper-V performance counters(Medium VMs)

Counter	1 VM	2 VMs	3 VMs	4 VMs	5 VMs	6 VMs	7 VMs
Hyper-V Logical Processor-%Hypervisor Run Time	0.4	0.8	1.1	1.3	1.8	2.1	2.2
Hyper-V Logical Processor-%Total Run Time	11.8	24.2	37	49.8	62.5	75.4	85.5
Hyper-V Logical Processor - %Guest Run Time	11.4	23.4	35.8	48.4	60.6	73.3	83.2
% Processor Time	0.9	1.5	1.9	2.5	3.3	3.7	4
% Privilege Time	0.8	1.4	1.8	2.4	3.2	3.6	3.9

Table 5. Hyper-V performance counters(Large VMs)

Counter	1 VM	2 VMs	3 VMs	4 VMs	5 VMs	6 VMs	7 VMs
Hyper-V Logical Processor-%Hypervisor Run Time	0.8	1.6	2.3	2.4	2.2	2.2	2
Hyper-V Logical Processor-%Total Run Time	23.3	47.4	70.6	84.9	93.1	96.6	98.3
Hyper-V Logical	22.5	45.8	68.3	82.4	90.8	94.4	96.2

Dell Virtualization Solution for Microsoft SQL Server 2012 using PowerEdge R820

Counter	1 VM	2 VMs	3 VMs	4 VMs	5 VMs	6 VMs	7 VMs
Processor - %Guest Run Time							
% Processor Time	1.4	2.3	3.4	3.9	3.9	3.9	3.6
% Privilege Time	1.3	2.3	3.3	3.8	3.8	3.8	3.5