Microsoft® SharePoint® Server 2010

Virtualized Small Farm with Dell PowerEdge 12th Generation Servers – A Performance Study

Ravikanth Chaganti
and Quocdat Nguyen

Dell™ SharePoint Solutions Engineering

March 2012

March 2012| Rev 1.0

Contents

Tables

Figures

Executive Summary

A Microsoft® SharePoint® Server 2010 farm hosts the core platform services and applications that provide many different functions for users. With the multi-tier architecture, sizing of each of the tiers of a SharePoint farm requires a comprehensive study of the workload requirements and performance capabilities of each hardware component. Dell's priority is to provide accurate guidance to customers when recommending infrastructure elements of a SharePoint implementation.

Dell's SharePoint engineering team developed a load generation framework to perform SharePoint load testing to provide guidance on choosing the best farm architecture to increase performance and keep client response times to less than one second. This performance data is provided to our customers to help them understand the impact of SharePoint collaboration workloads, and how to size and design the best farm architecture to support these workloads.

This white paper includes results of a comprehensive study and describes how a small virtualized SharePoint farm, built using the new 12th generation Dell PowerEdge™ servers and Dell EqualLogic™ storage arrays, performed under load testing. The key findings from this study are:

- The recommended farm architecture was able to support 3500 concurrent users.
- The farm architecture used in this study had an average farm response time of 0.31 seconds at the maximum supported user load, which was well below the one-second response time target.
- The Dell EqualLogic storage backend was capable of supporting the above concurrent user load with an average disk response time of ~7ms.
- The enhanced processing capabilities of the new Dell PowerEdge R720 servers resulted in 133% increase in the number of concurrent users supported, as compared to a similar performance study[1] using Dell PowerEdge R710 servers.

The paper also details information on how the farm was configured, factors considered while designing the farm, how Dell performs SharePoint load testing, and several performance metrics of various farm components.

These additional papers are available from www.dell.com/sharepoint:

- *Sharepoint Server 2010: Virtualized Small Farm Performance*, describes how a SharePoint small virtualized farm was built and configured on Dell PowerEdge 11th generation servers. The white paper also details the performance results of the virtual server SharePoint small farm under the load test.
- *Sharepoint Server 2010: An Introduction*, offers an overview of SharePoint Server 2010 and provides common concepts and definitions that form a good basis for understanding the reference architectures presented in this paper.
- *SharePoint 2010: Designing and Implementing a Small Farm*, provides the reference architecture and infrastructure best practices for implementing a SharePoint 2010 small farm. These reference architectures formed the basis of the performance study described in this paper.

---

[1] http://content.dell.com/us/en/enterprise/d/business~solutions~whitepapers~en/Documents~virtual-small-farm-performance.pdf.aspx

Introduction

Microsoft SharePoint Server 2010 builds on the capabilities that were offered in Microsoft Office SharePoint Server 2007 to provide a rich platform for collaboration, information sharing, and document management. SharePoint 2010 adds several new features, and introduces important architectural changes and product improvements.

Capacity planning for a SharePoint farm deployment requires a thorough study of the existing requirements, future growth and application use. These choices introduce a complexity factor while sizing servers and storage for a SharePoint implementation.

There are six pillars[2] or the six facets of SharePoint 2010 that organizations can apply to create clarity around how SharePoint is used. This performance study paper provides performance capacity details of a virtualized SharePoint 2010 small farm configured with the 12th generation Dell PowerEdge rack servers and EqualLogic iSCSI storage in the context of SharePoint collaboration[3] workload.

SharePoint 2010 Farm Topologies

A SharePoint server farm is a collection of servers that collectively provide the services required by a SharePoint deployment. Some of these services, or sets of services, comprise predefined roles, and must be configured within the solution. Other services and components are optional, but they provide additional features and functionality that are often desirable. The optional components may include service applications, such as managed metadata and Excel services. Some constraints and best practices help determine which components should be located on each server in the farm. In addition, by considering how the components are distributed, you can design the farm to more easily accommodate later growth.

NOTE: In SharePoint Server 2010, components generally provide functionality for a given service application. As a result, this paper may use the terms “role” and “component” interchangeably. In this context, SharePoint roles refer to one or more components that provide a farm service, and should not be confused with Windows Server roles, which generally include one or more Windows services to provide operating system functionality.

The size and capacity of a SharePoint 2010 implementation can vary based on several factors, such as number of concurrent users, service applications in the farm, the expected uptime service-level agreement (SLA), etc. These factors dictate how many servers are needed in the SharePoint farm and how the overall farm architecture looks. Based on the these factors, SharePoint 2010 farm implementations can be classified as small farm, medium farm[4] or large farm[5] deployments.

---

[2] SharePoint 2010 Capabilities: http://sharepoint.microsoft.com/en-us/product/capabilities/Pages/default.aspx

[3] SharePoint 2010 Collaboration Capabilities: http://sharepoint.microsoft.com/en-us/product/capabilities/communities/Pages/default.aspx

[4] SharePoint 2010 – Designing and implementing a medium farm: http://www.dell.com/downloads/global/services/dell_large_sharepoint_farm.pdf

[5] SharePoint 2010 – Designing and implementing a large farm http://www.dell.com/downloads/global/services/dell_large_sharepoint_farm.pdf

Small Server Farm Topology

A typical SharePoint small server farm[6] consists of two tiers: Web front-end/application and database. Dedicated servers are used to host each tier to provide process isolation and to allow for future growth. A server farm deployment model helps ensure that the solution infrastructure is scalable, flexible and resilient to hardware failures. To achieve these goals, a small farm implementation uses multiple servers at all tiers of the farm deployment. The farm model uses a dedicated database server, and generally distributes Web front-end server roles and application server roles across multiple hosts at the front-end tier. This performance study paper used the SharePoint 2010 small farm architecture to understand how several components of a farm perform at incrementing user loads.

Within the scope of this paper, a virtualized farm configuration was used to study the performance characteristics of SharePoint 2010 on Dell servers and storage. This virtualized configuration describes a SharePoint deployment on virtual machines using Microsoft® Hyper-V virtualization. The following sections explain how the virtual farm was configured and how it performed during the load test.

As shown in Figure 1 and Figure 2, the virtual farm configuration used in the performance study employed Dell PowerEdge R720 rack servers and an EqualLogic PS6000XV storage array for the database backend. The following sections describe the server choice for the Hyper-V host and provide a technical overview of the servers/storage array used in this performance study.

Figure 1. Virtual Farm Configuration - Front View of the Farm

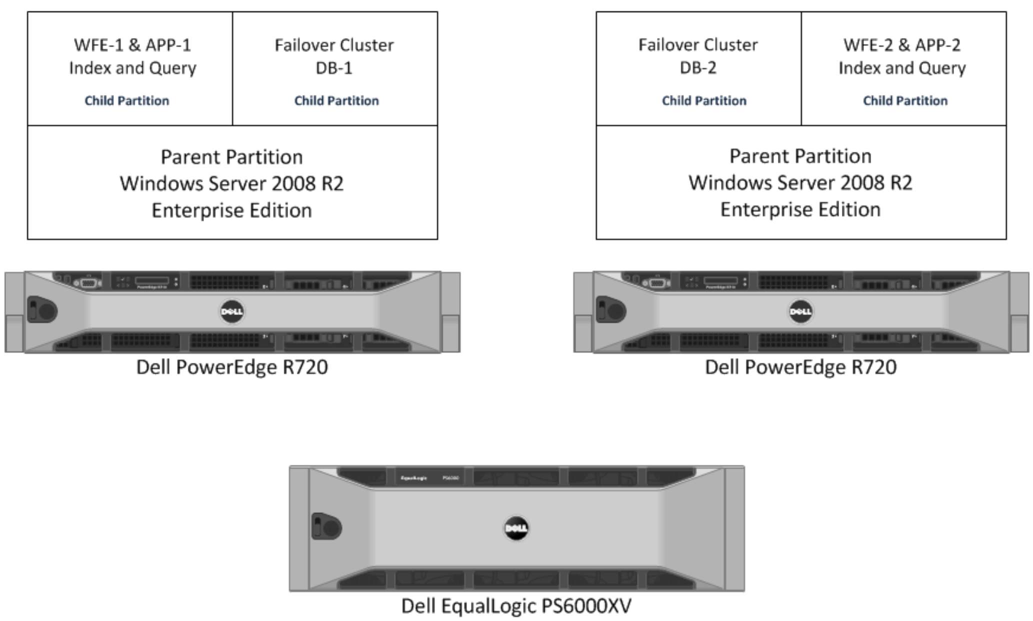

[6] SharePoint 2010 – Designing and implementing a small farm: http://www.dell.com/downloads/global/services/dell_small_sharepoint_farm.pdf

Figure 2. Small Virtualized Farm - Cabling Details

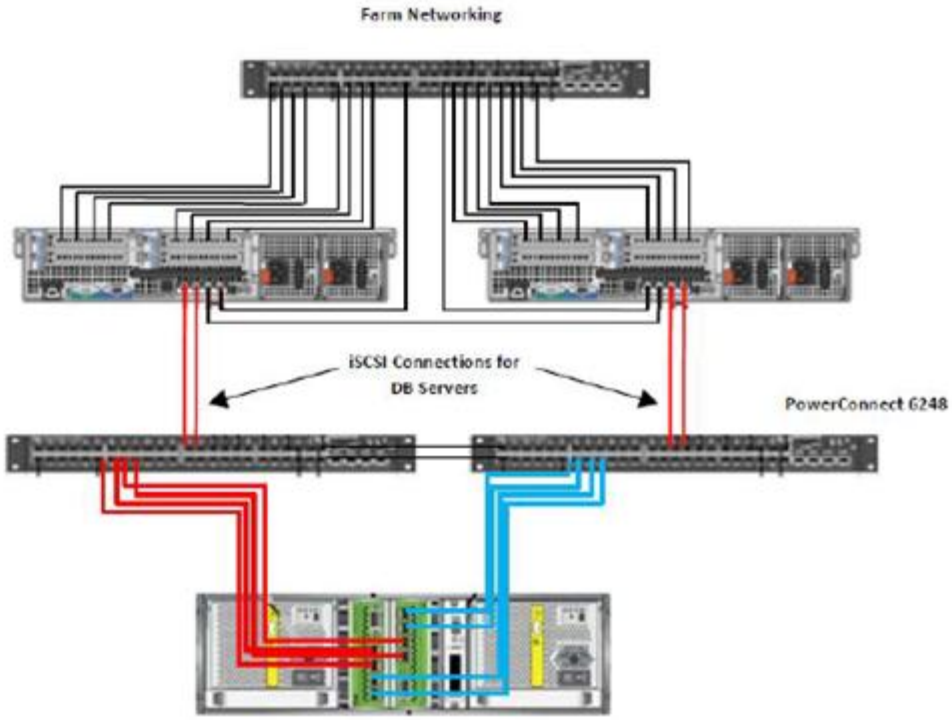

Dell PowerEdge R720 Server

The PowerEdge R720 server is part of Dell's 12th generation rack server family. The PowerEdge R720 is a two-socket, 2U rack server with support for up to 768GB of physical RAM and the latest six-core and eight-core Intel® Xeon® E5-2600 series processors.

The R720 offers a wider range of internal storage controllers compared to its predecessor, the Dell PowerEdge R710. This server also supports increased internal disk capacity of up to 16TB when using 16 x 1TB near-line SAS disks. It has four PCIe expansion slots that enable future hardware upgrades. In addition, the R720 has four internal embedded 1Gbps network ports. The increased processing and memory capacity makes the Dell PowerEdge R720 an appropriate choice for the virtualization platform.

The Virtual Farm section describes the PowerEdge R720 hardware configuration and its sizing to implement the virtualization configuration for the SharePoint small server farm in more detail. It also explains why a PowerEdge R720 server is a suitable choice to implement this virtualization farm.

EqualLogic PS6000XV Storage Array

The EqualLogic PS6000XV is a virtualized iSCSI SAN that combines intelligence and automation with fault tolerance to provide simplified administration, enterprise performance and reliability, and seamless scalability.

A PS-Series array provides the following features:

- No-single-point-of-failure hardware
- Redundant, hot-swappable hardware components (disks, control modules, fans, and power supplies)
- Component failover and disk sparing occur automatically without user intervention or disrupting data availability
- RAID technology provides data protection in each array
- High-performance control modules (the PS6000 control module has four 1 Gigabit Ethernet interfaces)
- Support for standard Gigabit Ethernet networks.

Virtual Farm Configuration

As a virtualization solution, Hyper-V allows user to take maximum advantage of server hardware by providing the capability to run multiple operating systems (on virtual machines) on a single physical server.

Figure 3. Hyper-V Architecture Block Diagram

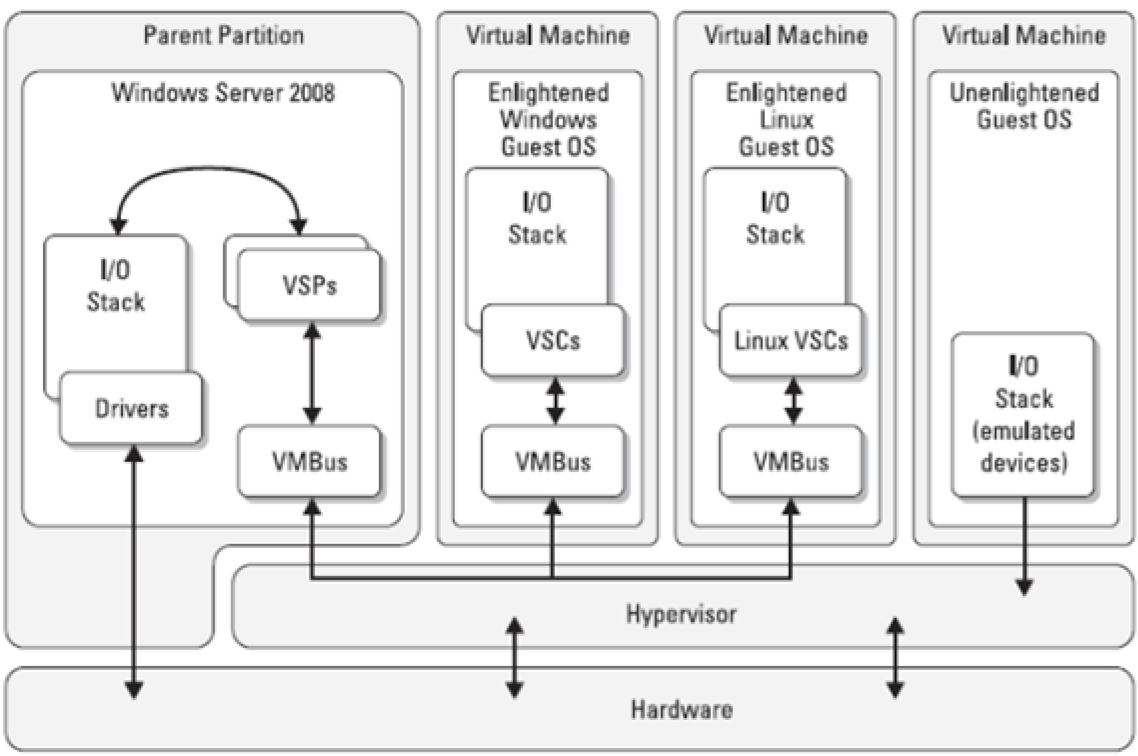

Hyper-V is implemented as a role in Windows Server 2008 R2 and is responsible for creating and managing isolated execution environments called "partitions."

Hyper-V requires the virtualization extensions in the processors (Intel VT or AMD-V™). When the Hyper-V role is enabled in Windows Server 2008 R2 and loaded for the first time, it creates a partition called the "root" or "parent" partition. The virtual machines created on this hypervisor are referred as "child partitions" and the operating system installed within a virtual machine (VM) is commonly called the "guest operating system." Figure 3 shows a high-level architecture of Hyper-V in Windows Server 2008 R2.

The root/parent partition has direct access, and controls all hardware devices such as network, storage, and memory allocation to the partitions. Unlike the parent partition, child partitions do not have access to physical hardware, but rather have virtualized devices. I/O requests from VMs are routed through the parent partition to the physical adapters on the system. Devices presented to the child partitions are either considered synthetic or emulated virtual devices. An emulated virtual device is a software implementation of a typical PCI-device, and requires additional processing power from a host server. Synthetic virtual devices, which are also implemented in software, use the high-performance VMBus channel as the communication mechanism between partitions.

Hyper-V also includes Integration Services, which provides better integration between child and parent partitions, device drivers for synthetic devices, and other enhancements such as time synchronization and support for VSS-based backup, and so on. The Hyper-V Integration Services automatically install after Windows Server 2008 R2 guest operating system installation completes.

Dell recommends a two-server solution using Dell PowerEdge R720 servers for the virtualized environment in the SharePoint Small server farm. The solution also employs high-availability technology to ensure continuity of service. As mentioned earlier, this SharePoint 2010 small server farm consists of two tiers: Web front-end/Application (WFE/APP) and database server (DB).

Deployment of server roles is on virtual machines instead of on physical servers. VM nodes of WFE/APP and DB are installed across two physical R720 servers using Windows Server 2008 R2 with Hyper-V enabled. In addition, Microsoft SQL Server® 2008 R2 is deployed in an active/passive Microsoft failover cluster environment to provide high availability for the farm database content. Figure 1 provides a high-level view of the server roles in this virtualized configuration.

The two PowerEdge R720 servers are equipped with two 8-core 2.40GHz Intel® Xeon® E5-2600 series processors, which can provide a total of 16 CPU cores per physical server (multi-threading was not enabled in this deployment). In this configuration, each VM is allocated four virtual processors (vProcs) and appropriate memory. This leaves adequate hardware resources for the host server to support the VMs and any hypervisor and operating system overhead. Table 1 and Table 2 show the sizing for both physical and virtual servers to support the SharePoint small server farm.

Table 1. Physical Server Configuration Details for a Virtual Farm

Host Platform	Processor	RAM	Storage	Network	OS
2 x R720	1. 2 x Eight-Core 2.40GHz Intel® Xeon® E5-2600 series processors	64GB	Internal (SAS 15K) 4 x 146GB RAID 1 Volumes OS, VHD Volumes, and Index on Pass-through Disk	2 x network for WFE, DB public, DB VM cluster, and host management: 4x NIC ports and 2x (Quad-port Broadcom 5709 NIC)	Windows Server 2008 R2 Enterprise with Hyper-V Role enabled Applied MS hotfix KB2517374[7]

Table 2. Virtual Machine Configuration Details for a Virtual Farm

Virtual Machine	Processor	RAM	Storage	Network
WFE & APP VMs	4 vProc	8GB	• OS Fixed VHD • Index Pass through volume	1 Dedicated Virtual Network: 4 Port NIC Teaming
SQL Server DB (Failover Cluster) VMs	4 vProc	16GB	• OS Fixed VHD • 4 LUNs for SQL Cluster DB using iSCSI connections	• 1 Dedicated Virtual Connection for Cluster / Farm network: 4 port NIC teaming • 1 Dedicated Virtual Connection for private cluster network • 2x dedicated virtual network connection for iSCSI

[7] You cannot start virtual machines on a computer that is running Windows Server 2008 R2 and on which a CPU is installed that supports the AVX feature: http://support.microsoft.com/kb/2517374

NIC teaming was used for the farm network connections on WFE and SQL DB servers to provide sufficient network bandwidth for the farm networking. Two iSCSI connections were used to provide redundancy and load balancing to the back-end storage array. All virtual networks are implemented using isolated dedicated virtual networks to ensure the isolation in the connectivity from each other, as well as from a VM to a parent/root partition. To achieve this, 2 x quad-port Broadcom add-on network adapters and 4 x internal LOM NICs were used in the R720 server. Figure 4 provides details on how the virtual networks were implemented.

Figure 4. Hyper-V Virtual Switch Configuration

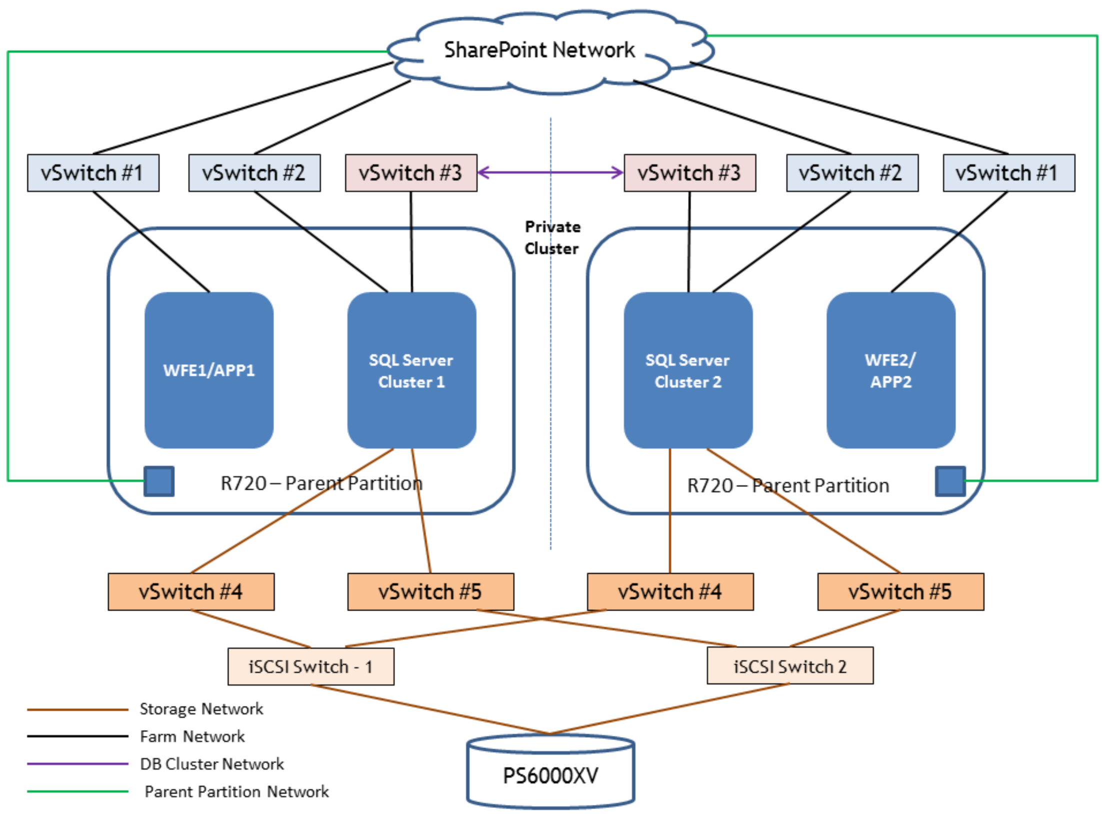

The PowerEdge R720 storage consists of eight HDD bays which provide 4 x 146 GB RAID-1 volume to host the two VMs. Three local volumes hosted the root operating system partition and two VHDs for WFE/APP and SQL Server DB nodes. The fourth volume was deployed as a pass-through drive for better performance of the farm's indexing feature.

EqualLogic iSCSI PS6000 storage is a suitable solution for this virtualized farm configuration. Hyper-V supports iSCSI connections initiated directly from within a guest operating system to an iSCSI storage array. This means that any iSCSI LUN that is provisioned directly to a guest operating system is not visible to the parent partition. In addition, the EqualLogic solution provides many advantages for supporting of high-availability VMs, VM Quick Migration, and additional storage capacity for the server.

For additional information on Dell value propositions, guidance on best practices to implement Hyper-V in your environment, and other design considerations including networking, storage, and high availability solutions, see the References section of this paper.

Farm Architecture and Configuration of Farm Roles

The experimental farm configuration, as shown in Figure 1 and Figure 2, included two Web Front End/Application Servers (WFE/APP) and two database servers in a failover cluster. Because there is negligible requirement in processing power on an application server role in this SharePoint small farm, the application server role was enabled on the same server as the WFE. This farm was configured to use Microsoft Windows® authentication and hence all the requests during the load test were authenticated requests.

The farm servers at all tiers used teamed network connections to provide load balancing and failover capabilities.

Table 3 lists the operating system and software editions used in the farm configurations. The rationale for choosing this matrix is explained in the later sections of this paper.

Table 3. OS and Software Editions Used in this Study

	Web Front-ends	Database Servers
Operating System	Windows Server 2008 R2 Enterprise Edition	
SharePoint Server	SharePoint Server 2010 Standard Edition	NA
Database Server	NA	SQL Server 2008 R2 x64 Standard Edition

NOTE: Step-by-step instructions for installing and configuring a SharePoint farm and any service applications used in this performance study are outside the scope of this paper. For more information and resources, refer to the References section of this paper.

Configuration of Web Front-end Servers

This SharePoint 2010 farm design included two Web front-end servers. The software matrix for these web front-end servers is shown in Table 3. SharePoint 2010 Standard Edition was used as the performance study, including only out-of-the-box features used for collaboration workloads. As a part of the collaboration workload, only the search service application was deployed' no other service applications such as Excel and Visio were deployed.

All of the Web front-end servers were configured in a Network Load Balancing (NLB) cluster[8]. Using NLB, stateless applications such as SharePoint Web front-end can be made scalable by adding additional servers when the load increases. As mentioned earlier, because the network connections were made redundant by using Broadcom NIC teaming, NLB has been configured to use multicast mode to avoid IP address conflicts[9] in the farm.

[8] NLB Deployment Guide: http://technet.microsoft.com/en-us/library/cc732855(WS.10).aspx
[9] Using teamed adapters with network load balancing may cause network problems: http://support.microsoft.com/kb/278431

HTTP Request Throttling

SharePoint 2010 offers resource throttling features that can be configured to help increase server performance and protect server resources during peak usage times. SharePoint 2010 has a default timer job that checks server resources compared to configured throttle levels. By default, server CPU, memory, requests in queue, and request wait time are monitored. After three unsuccessful checks, the server enters a throttling period and will remain in this state until a successful check is completed. Requests that were generated prior to the server's entering throttling mode will be completed. Any new HTTP GET and Search Robot requests will generate a 503[10] error message, and will be logged in the event viewer.

The throttle settings can be modified to increase the overall load supported by the farm servers. However, a complete study is required to determine accurate throttle setting recommendations for any given user load or requests per second. The default HTTP throttle monitor settings will prevent an extensive load testing to find out the real capacity of the farm servers. As a result, HTTP request throttling was turned off during the load testing of SharePoint.

Search Service Application Configuration

SharePoint 2010 changed the search architecture and introduced high availability at the application tier or crawler. The new search service application architecture in SharePoint 2010 includes greater redundancy. The new design provides flexibility and lets the query and crawler roles be scaled-out separately on an as-needed basis. Search crawlers are now stateless; they do not store a copy of the index. However, the index does still propagate and is stored locally on the query servers. Two virtual machines hosting both crawler & WFE roles were used in this performance study.

The small farm configuration used the search service application configuration shown in Figure 5. The search crawler and query roles were hosted on both of the Web front-end servers to enable high availability and provide improved search performance. Both the crawlers were associated with the same crawl database, and indexed the same content source. Two index partitions were created and, for redundancy, a copy (or mirror) of the index partition was placed on each Web front-end. On both of the Web front-ends, a dedicated pass-through RAID 1 volume was used to store the index contents.

The search architecture shown in Figure 5 is a logical presentation of how SharePoint Search service was configured. This configuration provides complete redundancy for both search crawler and query roles.

---

[10] Throttling starts alert – Events 8032 and 8062: http://technet.microsoft.com/en-us/library/ee513044.aspx

Figure 5. Search Architecture Deployed in This Study

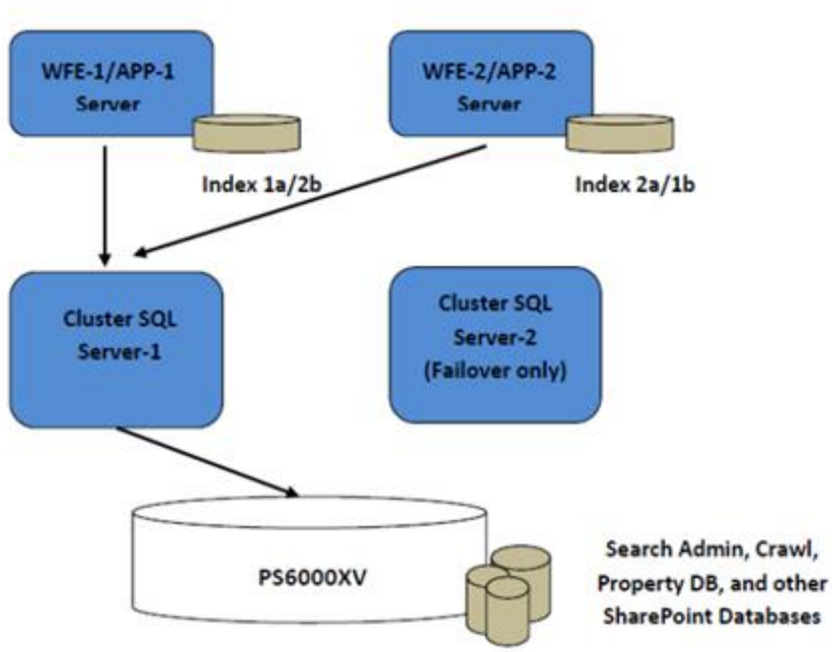

Network Configuration

For both the Web front-end and application servers and SQL Database Servers, teamed network connections were used in the SharePoint farm's network. These teamed connections (shown in Figure 2) were configured to be in the smart load balancing mode (SLB), which supports both load balancing and failover.

Configuration of the Database Server

As shown in Figure 1 and Figure 2, this performance study paper deployed virtual machines hosted on a PowerEdge R720 server as the database servers in the virtualized farm configuration. Two database servers were deployed in a highly available fail-over cluster, with redundant data paths at the database tier of the SharePoint farm.

A SharePoint farm's performance depends on the performance of the database server and the database backend. As the storage resources on the host server hosting the virtual machine are limited, the SQL instance was connected directly to an external EqualLogic iSCSI storage array. In this virtualized configuration of the SharePoint small farm, one EqualLogic PS6000XV storage array was used to provide sufficient storage capacity and I/O bandwidth to the SQL database contents. This array provides 16 x 300GB 15K RPM SAS drives configured in a RAID 10 for storing the SharePoint content.

Figure 6. DB Layout for the SharePoint Databases

SQL Server Memory Configuration

By default, SQL Server uses all available physical memory[11]. This is because SQL Server dynamically grows and shrinks the size of the buffer pool depending on the physical memory reported by the operating system. However, this behavior can be adjusted to limit the amount of physical memory used by SQL Server. Within the scope of this paper, SQL Server memory was limited to 80 percent of the actual physical memory available in the system. For example, on the virtual clustered SQL database server, out of 16GB of physical memory, 12.8GB was allocated to SQL Server.

Database Server Network Configuration

Similar to the Web front-end and application tiers, the database tier also used teamed network connections for the farm network as shown in Figure 2. For the iSCSI storage network, two virtual network connections were dedicated, and MPIO was configured to provide load balancing and failover capability.

Performance Study of a Small Farm

SharePoint 2010 is a versatile platform that can be used in a variety of ways. Some SharePoint workloads work almost out of the box; others require or allow significant customization; and still others are the result of completely custom-developed applications. This flexibility results in many possible ways to use SharePoint, leading to significant challenges in accurately sizing servers and storage for a SharePoint farm.

Given this difficulty and the lack of a standard benchmark for sizing SharePoint workloads it is important to provide accurate guidance to customers when recommending the infrastructure elements of a SharePoint implementation. These factors led to the development of the Dell SharePoint Load Generation framework used to perform load testing of a SharePoint farm.

[11] SQL Server memory options: http://msdn.microsoft.com/en-us/library/ms178067.aspx

Dell SharePoint Load Generation Framework

An internally developed load generation framework was used to understand the performance characteristics of a SharePoint farm. This framework includes load testing of SharePoint out of the box usage profiles, such as collaboration and publishing.

The Dell SharePoint load generation framework has two components - a content population tool and Microsoft Visual Studio® Team Suite (VSTS) Web test framework.

Content Population Tool

The content population tool is designed to prepare the SharePoint farm for load testing. This content population tool was designed to distribute the SharePoint content across multiple site collections.

Figure 7. LoadGen Tool Framework

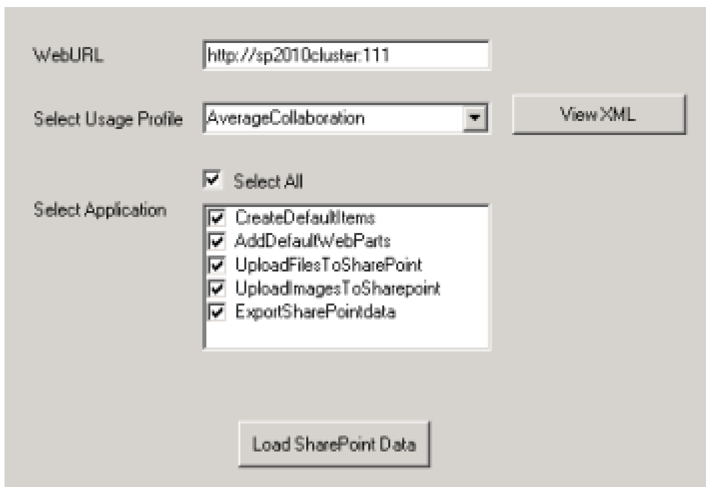

The content population tool was developed to

Create SharePoint Web applications

- Create site collections
- Add Web components to home pages
- Create document libraries
- Create SharePoint list items
- Upload documents, images, and other files

This utility is capable of populating hundreds of gigabytes of SharePoint content in a few hours. The size of the SharePoint content database and other aspects such as number of site collections vary based on the usage profile selection. A usage profile is a collection of use cases closely mapped to realistic SharePoint usage. To some extent, these usage profiles were mapped in to SharePoint Capacity

Planner[12] and used other Microsoft recommendations. Although the SharePoint Capacity Planning Tool was intended for Microsoft Office SharePoint Server 2007, there are several aspects of these recommendations[13] that still apply to SharePoint 2010 out of the box workloads. The content generated and uploaded by the content population tool serves as a baseline for SharePoint 2010 load testing using Visual Studio test framework.

Visual Studio Load Testing (VSTS) Framework

Dell's SharePoint load generation framework uses VSTS 2010 to perform load testing. Within Visual Studio, each load test directly maps in to a SharePoint usage profile, and each usage profile defines a list of use cases and how many use cases are run per hour per connected user. Using VSTS 2010 helps to more rapidly create use cases and parameterize those use cases. SharePoint load testing is performed using a test rig, as shown in Figure 8, of several physical test agents. The results are captured in to a SQL database on the test controller.

Figure 8. Visual Studio Test Rig

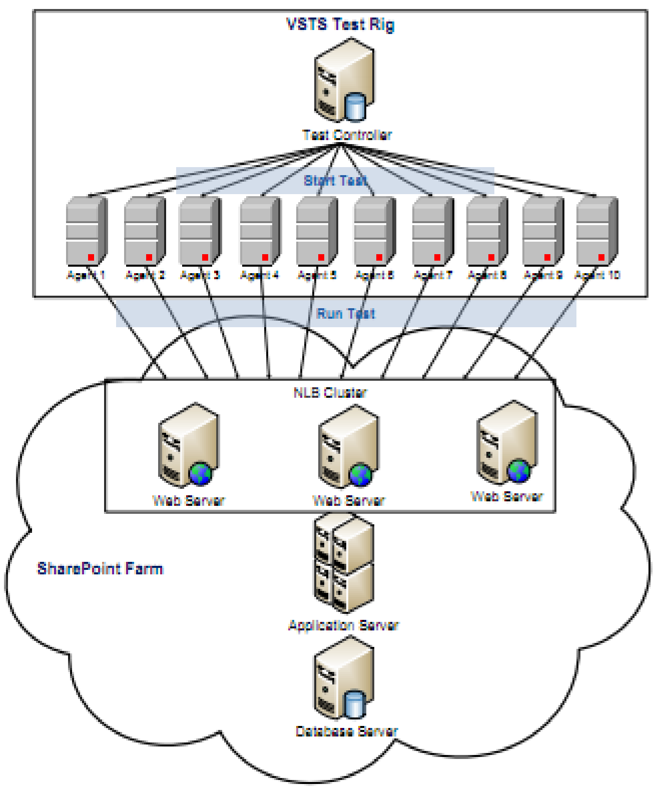

[12] SharePoint capacity planner - http://www.microsoft.com/downloads/details.aspx?FamilyID=dbee0227-d4f7-48f8-85f0-e71493b2fd87&displaylang=en

[13] Microsoft SharePoint 2010 performance and capacity management - http://technet.microsoft.com/enus/library/cc262971.aspx

Load Testing Workload Test Mix

As mentioned earlier, the load test usage profiles were based on the SharePoint Capacity Planner (SCP) and other Microsoft recommendations for SharePoint 2010. System Center Capacity Planner defines several usage profiles for both collaboration and publishing workloads. These usage profiles are categorized as light, medium, and heavy-usage profiles. These categories define several aspects of a usage profile, such as how many requests are sent per hour per connected user, what use cases constitute a load test, and what percentage (test mix) of each use case is used within each load test. Within the scope of this performance study paper, a heavy collaboration usage profile was used. Table 4 shows the heavy collaboration test mix as suggested by System Center Capacity Planner for SharePoint.

Table 4. SCP Collaboration Profile Mix

SCP Usage Profiles	Heavy Collaboration
Home Page Access (%)	30
List Page Access (%)	20
Document/Picture Download (%)	15
Document/Picture Upload (%)	8
Search (%)	15
Total Requests / hour / connected user	60

As shown in Table 4, SCP defines only a high-level test mix for each usage profile. Table 5 shows a more granular translation of this SCP heavy collaboration usage profile. Several use cases were mapped into each category, and the number of use cases per hour per connected user has been assigned.

Table 5. Dell Heavy Collaboration Test Mix

Heavy Collaboration Test Mix	Number of tests/hour/user
Home Page Access	
Read Site Home Page	18
List Page Access	
Read Survey	6
Read List	6
Document/Picture Download	
Read Document Library	2
Read Home to Document Library	1
Read Wiki Page	2
Read Picture Library	1
Read Home to Wiki Page	2
Read Home to Picture Library	1
Document/Picture Upload	
Create Wiki Page	3
Upload Document	2
Search	
Search Site	10
List Item Insertion / Deletion	
Respond to Survey	2
Reply to Discussion Topic	1
Edit Wiki Page	2
Comment home to blog post	1
Total tests/hour/connected user	60

Dell's test mix, shown in Table 5, is not a one-to-one mapping into the above SCP and Microsoft recommendations. For example, SCP defines total requests per hour per connected user. However, within Dell's test mix for the heavy collaboration profile, this translates into more than 60 requests per hour, as the usage profile uses 60 ***tests*** per hour for each connected user. Additionally, one test could mean more than one request. Hence, the results published in this paper may or may not map directly to SCP recommendations, and are specific to the workload mix defined in Table 5.

Test Methodology

The intent of the experiments conducted as a part of this performance study was to understand the capacity of a small virtualized SharePoint farm, as shown in Figure 1 and Figure 2, with the configuration described in Table 2. Several load test iterations were conducted with incremental user load. For example, an initial user load of 250 virtual users was used, and then incremented by 250 users until the farm resources reached an optimal level of usage. The overall goal of the load test was to ensure that the processor utilization was below 70% and the average farm response time is less than one second.

The data set used to build the content database included several different types of files. This includes Microsoft Office documents, Adobe® PDF documents, and several image formats. Table 6 shows a distribution of file content sizes used in this performance study.

Table 6. Dataset Categorization

Average File Size	Number of Files
1KB to 500KB	34240
500KB to 1MB	5223
1MB to 10MB	13003
10MB to 70MB	125

The aggregated SharePoint content database size was around 53GB. During the load test duration, this content database grew by almost 20 percent. This performance study involved load testing of out-of the-box SharePoint deployment using a test mix shown in Table 5. A full content crawl was performed once at the beginning of the load tests. There were no subsequent crawls after load testing or during the load test duration.

The performance data shown in this paper was a result of load testing on the final configuration of a SharePoint farm as described in Table 2. The following sections of this paper described the performance data and how several components within the farm performed at incremental user loads.

Performance Results & Analysis

As a part of this performance study, several performance metrics were collected and analyzed. Based on the results, the farm configuration was modified to reach the final farm configuration shown in Figure 1 and Figure 2. This section describes the performance data and how various components of the SharePoint farm performed under incremental user load.

As mentioned earlier, this study included only a collaboration workload and the test mix shown in Table 5. Therefore, all results shown here are relative to the workload used and may differ with any other implementation outside of the test mix shown.

The following table shows, at a high level, how the SharePoint 2010 farm performed in terms of maximum concurrent user load achieved and requests per second.

Table 7. High Level Overview of the Small Farm Performance

	SharePoint 2010 Virtualized Small Farm
Maximum concurrent user load achieved	3500
Requests per second at max concurrent user load	173 / sec

The virtualized SharePoint small farm shown in Figure 1 could support up to 3,500 concurrent users with an average farm response time of 0.31 seconds at the maximum user load tested on the final configuration. This indicates that the farm architecture used for this performance study could support faster sub one-second farm response times, even at the maximum concurrent user load.

Figure 9. Farm Response Time

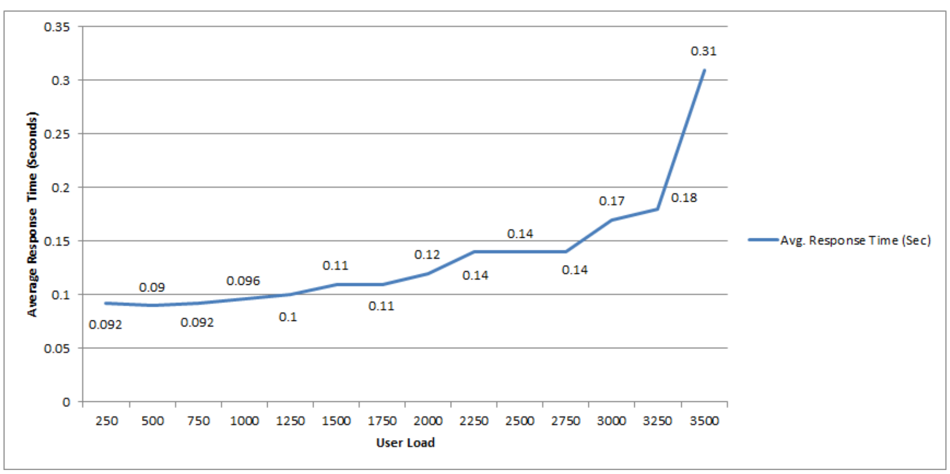

The average processor utilization on the virtual machines used as web front-end servers in this farm configuration was around 78 percent at the maximum user load. As expected because NLB was used at the Web front-end tier, all WFEs were more or less equally loaded.

Figure 10. Average Processor Utilization

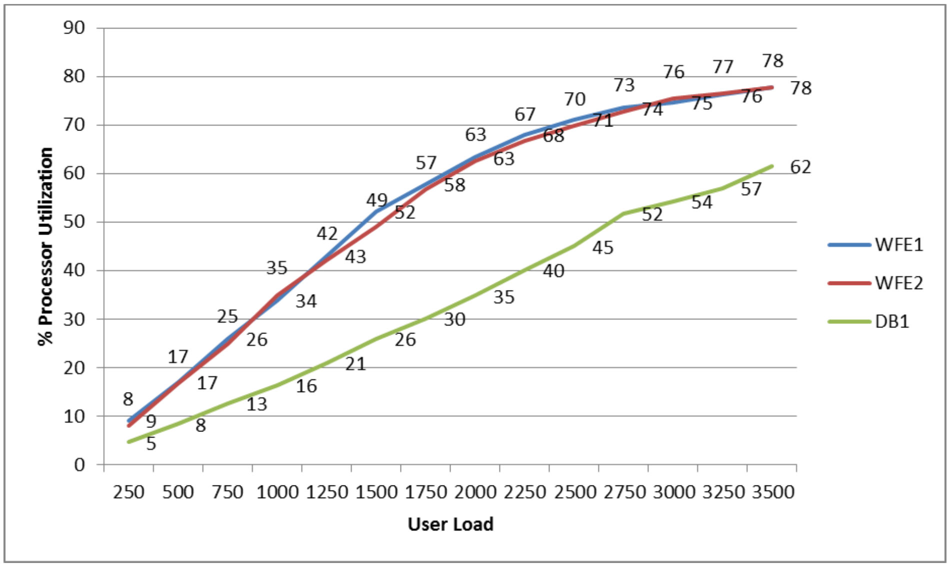

In a heavily loaded scenario, even though the average processor utilization is below 80%, the SharePoint 2010 farm (Figure 1) may not support more users than what is shown in Table 7. This constraint is mainly because of the ASP.NET and IIS request queue length limitations. The out of the box IIS and ASP.NET queue length settings can be tweaked to go beyond the concurrent user load shown in this performance study paper. However, this analysis is outside the scope of this paper and may require an in-depth study in itself.

The overall network utilization was well below 50% of the total available bandwidth, 4GB of a four-NIC teaming in this virtualized farm configuration. The following chart captures the network utilization at Web front-end and database tiers of the farm, and shows the aggregated performance numbers of the teamed NICs used on all servers in the farm. The database network utilization numbers in the following chart are an aggregation utilization of all network channels used in MPIO.

Figure 11. Network Utilization on Farm Servers

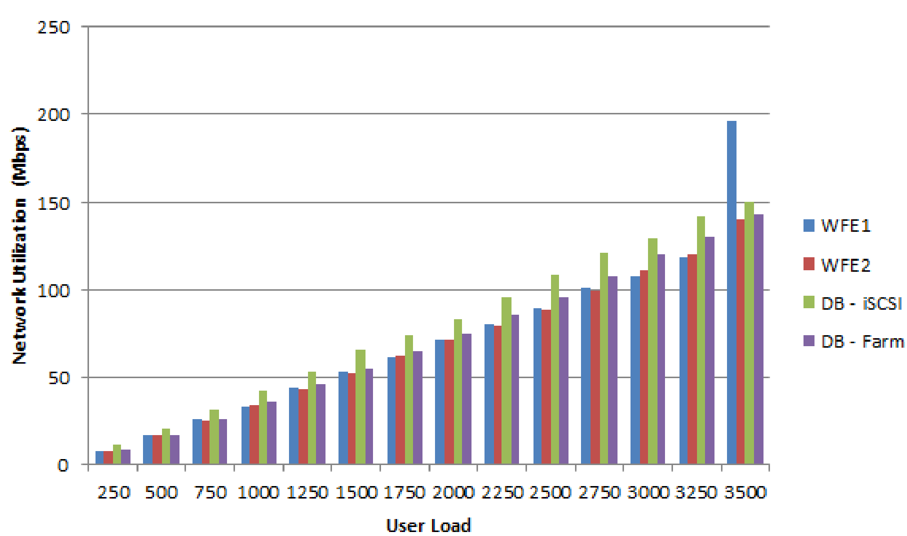

In addition, performance results show that maximum memory utilization on the Web front-end servers at the maximum user load was within 60% of the available physical memory. This indicates that there is enough room for future growth, while providing high availability for all SharePoint roles hosted on Web front-end servers.

As shown in Figure 1 and Figure 6, one EqualLogic PS6000XV array was used, and four SQL server databases were placed on the iSCSI LUNs provisioned on this array. The following charts show how these iSCSI LUNs performed during the load test with incremental user loads. The following table shows the I/O read-write statistics and overall IOPS achieved at the maximum supported concurrent user load.

Table 8. Storage Performance Metrics

IO Performance Metric	SharePoint Virtualized Small Farm
Avg. Disk Transfers / Second	~440
Avg. Disk Writes / Second	265
Avg. Disk Reads / Second	176
Avg. Disk Queue Length (_Total)	2.38
Avg. Disk Bytes/Transfer	42609 or 41KB
Avg. Disk Seconds / Transfer	0.0070 or 7ms

The I/O generated by the test mix shown in Table 5 was relatively small in size (~100KB), and was random I/O. This resulted in approximately 440 disk transfers per second to the LUNs provisioned on the EqualLogic iSCSI enclosure.

The above storage metrics indicate that the workload had an I/O mix of ~40 percent reads and ~60 percent writes with an approximate I/O size of 41KB. The following table provides the detailed storage usage statistics for the farm configuration used in this performance study.

Table 9. Detailed Storage Performance Data

Database Name	Avg. Disk Queue Length	Disk Transfers / Second	Disk Seconds / Transfer
Content DB	1.8	221	0.0089
Content Logs	0.1	42	0.0002
Search DB	0.35	37	0.0089
Temp DB	0.0062	125	
Usage DB	0.029	3.4	

The EqualLogic PS6000XV array with a 16-disk backend provided optimal performance for the SharePoint databases. All the performance metrics were seen to be well within the acceptable limits.

The following chart shows the processor utilization at the Hyper-V host.

Figure 12. Hyper-V Host's Processor Utilization

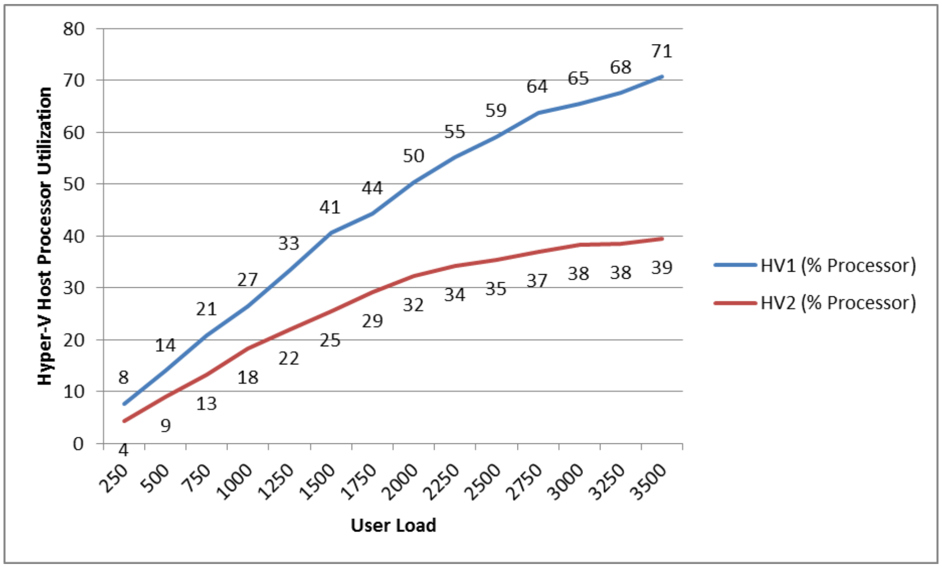

Allocated with 4 virtual processors each VM, WFE1/WFE2 and SQL Server® database, DB1-- at the peak user load supported--utilized almost 80 and 50 percent, -- – respectively, of their processing power. The Hyper-V hosts which hosted the farm virtual machines had 70 and 40 percent of processor utilization respectively. The second Hyper-V host (HV2) utilized less processing power due to the containment of a passive cluster SQL machine. The processor utilization of the Hyper-V hosts as seen at the maximum user load shows that the farm is in a healthy state. The total memory availability of Hyper-V hosts and VMs are sufficient. There is enough available memory for each node to perform under a maximum workload.

An earlier performance study[14] -- with the Dell PowerEdge R710 as the virtualization host in a small virtualized SharePoint farm -- showed that the maximum supported concurrent users was 1500. This was due to the limitation with the number of processing cores available on the Dell PowerEdge R710. With the increased processing and memory capacity, the Dell PowerEdge R720 servers are capable of supporting higher work load than the 11th generation R710 server.

This performance study with Dell PowerEdge R720 Servers showed that there is enough room to increase the number of virtual machines on each of the virtualization hosts to support a higher number of concurrent users than what was achieved with Dell PowerEdge R710 servers as the virtualization hosts.

---

[14] http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/virtual-small-farm-performance.pdf

The following table shows the difference in number of concurrent users achieved in these two performance studies.

Table 10. Difference in SharePoint Farm Performance Between R720 and R710

	SharePoint 2010 Farm on Dell PowerEdge R720	**SharePoint 2010 Farm on Dell PowerEdge R710**
Maximum concurrent user load achieved	3500	1500
Requests per second at max concurrent user load	173 / sec	73.7 / sec

Summary

A SharePoint 2010 farm consists of multiple servers, each provisioned with different SharePoint components. In general, a smaller SharePoint farm is the best choice for small and medium businesses with relatively lower concurrent user load. A small SharePoint farm employs two-tier architecture with all SharePoint Web front-end and application roles hosted on the Web front-end servers, and, at the database tier, one or more SQL Servers in a failover cluster are used to enable high availability of SharePoint databases. The reference architecture used in this performance study enables high availability at all tiers of the farm, and provides complete search service application redundancy by hosting two crawlers and mirroring the index partitions.

SharePoint 2010 can be used in many different ways, and each implementation needs an in-depth study of requirements, such as expected user load, requests per second and future growth. This performance study paper was intended to help customers understand the performance capacity of a virtualized small SharePoint 2010 farm built on Windows Server 2008 Hyper-V R2, using PowerEdge rack servers and EqualLogic iSCSI storage.

This study showed that the virtualized farm configuration could support approximately 35000 users, with a minimum concurrency of 10 percent. Also, the average farm response time was well below one second. The EqualLogic PS6000XV array and the Dell PowerEdge R720 servers provided highly optimal performance for the SharePoint 2010 deployment used in this performance study.

References

Dell solutions for SharePoint: Http://dell.com/sharepoint

Microsoft SharePoint page: http://sharePoint.Microsoft.com

Microsoft SharePoint TechNet: http://technet.microsoft.com/en-us/sharepoint/default.aspx

Dell EqualLogic: http://www.dell.com/EqualLogic

SharePoint 2010 operations framework and checklists:
http://technet.microsoft.com/enus/library/gg277248.aspx

SharePoint 2010 health monitoring: http://technet.microsoft.com/en-us/library/ee681489.aspx

Dell solutions for Hyper-V: www.dell.com/hyper-v
http://support.dell.com/support/edocs/software/HyperV/en/

- Dell Solution Overview Guide for Microsoft Hyper-V
- Dell Networking Solutions Guide for Microsoft Hyper-V
- Dell Storage Solution Guide for Microsoft Hyper-V
- Dell High-Availability Solution Guide for Microsoft Hyper-V