

---

# Implementing Microsoft SharePoint Server 2010 on Dell 12<sup>th</sup>-Generation Blade Servers and VMware vSphere 5.0


---

*A reference architecture and performance study for a SharePoint  
Server 2010 farm up to 10,000 concurrent users*

Dell Global Solutions Engineering

Revision: A00

September 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, Compellent, Force10, 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. VMware, vSphere, and vSwitch are either trademarks or registered trademarks of VMware. 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.

September 2012 | Rev A00

## Contents

Executive Summary .....	6
Introduction .....	7
Microsoft SharePoint Server 2010.....	7
VMware vSphere 5.0.....	8
Solution Architecture and Configuration .....	9
Infrastructure Components .....	9
Dell PowerEdge M1000e .....	10
Dell PowerEdge M620 .....	11
Dell PowerEdge M820 .....	11
Dell PowerConnect M8024-K 10Gb Ethernet Switch .....	11
Brocade M5424-N 8Gbps Fibre Channel Switch .....	12
Brocade 5100 Fibre Channel Switch .....	12
Dell Force10 S55 Network Switches .....	12
Dell Force10 S4810 Network Switches .....	13
Dell Compellent Series 40 Storage Center .....	13
BIG-IP Local Traffic Manager (LTM) Virtual Edition .....	16
Virtual Infrastructure Configuration.....	17
Network architecture .....	20
VMware ESXi Host (M620) Networking.....	22
Microsoft SQL Server (M820) Networking .....	25
Storage architecture .....	26
Storage Frontend Architecture .....	26
Storage Backend Architecture .....	28
SharePoint Farm Configuration.....	28
Configuration of Web Front End (WFE) Servers.....	30
Search service application configuration .....	30
HTTP Request Throttling .....	31
Configuration of the Database Server .....	32
Memory configuration .....	33
Performance Study of a SharePoint Farm .....	34
Dell SharePoint Load Generation Framework.....	34
Content Population Tool.....	34
VSTS Load Testing Framework .....	35
Load Testing Workload Test Mix .....	36

Test Methodology .....	38
Performance Results and Analysis .....	39
Summary .....	42
References .....	43

## Tables

Table 1. High-level overview of the farm infrastructure components.....	9
Table 2. Dell Compellent hardware capabilities.....	14
Table 3. Hardware configuration of blade servers .....	16
Table 4. WFE, APP, and BIG-IP virtual machine configuration .....	17
Table 5. Virtual machine to cluster host mapping .....	19
Table 6. NPAR bandwidth allocation.....	22
Table 7. Network partitions to vSwitch mapping.....	24
Table 8. VLAN to virtual switch mapping - LAN.....	24
Table 9. SQL network connections overview .....	25
Table 10. Trunk and N_Port mapping .....	27
Table 11. Operating systems and software editions.....	29
Table 12. SharePoint database LUN details .....	32
Table 13. Dell's Collaboration test mix .....	37
Table 14. Dell's Publishing test mix .....	37
Table 15. Data set used in the performance study .....	38
Table 16. Validated farm capacity .....	40
Table 17. Storage performance metrics .....	42

## Figures

Figure 1. High-Level Overview of Farm Infrastructure Components .....	10
Figure 2. Virtual Machine Placement .....	18
Figure 3. VMware DRS Rules.....	19
Figure 4. High-Level Network Architecture .....	21
Figure 5. VMware ESXi Host Networking .....	23
Figure 6. High-Level Network Design on the Database Hosts .....	25
Figure 7. High-Level Storage Frontend Architecture .....	26
Figure 8. High-Level Storage Backend Architecture .....	28
Figure 9. SharePoint Farm Logical Architecture.....	29
Figure 10. High-Level Search Service Architecture.....	31
Figure 11. SQL Service Disk Dependency Configuration .....	33

Figure 12. LoadGen Content Population..... 34

Figure 13. Load Test Deployment ..... 35

Figure 14. Average Farm Response Time..... 40

Figure 15. Processor Utilization on Farm Servers..... 41

## 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.

To help customers with this sizing task, the Dell Global Solutions Engineering team developed a load generation framework to test the SharePoint application's performance and provide guidance on the best farm architecture to keep client response times below one second. This performance data is provided to our customers to help them understand the impact of different SharePoint workloads, such as SharePoint collaboration and publishing, 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 partially-virtualized SharePoint farm, built using Dell PowerEdge™ blade servers, Dell Compellent™ Series 40 Storage Center storage arrays with Fibre Channel controllers, and VMware vSphere 5.0 performed under load testing. The key findings from this study are:

- The recommended farm architecture was able to support more than 100,000 users with 10 percent concurrency.
- The farm architecture used in this study had an average farm response time of 69ms or 0.069 seconds at the maximum supported, which was well below the one-second response time target.
- The Dell Compellent storage backend was capable of supporting a 2TB SharePoint content database<sup>1</sup> with an average disk response time of 2ms and ~1454 average disk transfers per second.

This 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. For more information on SharePoint Server 2010 reference architectures and performance studies, see <http://dell.com/sharepoint>.

---

<sup>1</sup> SharePoint 2010 content database limits: <http://technet.microsoft.com/en-us/library/cc262787.aspx#ContentDB>

## Introduction

Microsoft SharePoint Server 2010 builds on the capabilities that were offered in Microsoft Office SharePoint Server (MOSS) 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 business application and user needs, as well as plans for future growth. SharePoint can be deployed for a wide range of services, from website development to content management to custom-developed applications and process flows. Use cases can dramatically change sizing requirements.

We chose a typical mixed collaboration<sup>2</sup> and publishing<sup>3</sup> workload for this performance study. This paper provides performance capacity details of a partially-virtualized SharePoint 2010 farm configured with Dell PowerEdge blade servers, Dell Compellent Storage Center Series 40 storage arrays, and VMware vSphere 5.0 for this mixed workload.

## Microsoft SharePoint Server 2010

SharePoint is a family of products and technologies that interact with Microsoft SQL Server<sup>®</sup> and Internet Information Server (IIS) to provide a web-based engine and a platform for deploying a wide range of business services. The most common solutions deployed using this platform are collaborative sites, content management systems, publishing intranets, business intelligence systems, and Web portals. SharePoint solutions are usually deployed in a farm environment that provides scalability by distributing database, application, and presentation roles across different groups of servers. Some of the SharePoint services, or sets of services, comprise predefined roles, and must be configured within the solution. These predefined roles include server roles, such as web front end (WFE) and application (APP) service roles. Other services and components are optional, but they provide additional features and functionality that are often desirable. These optional components – deployed usually as a part of APP server role – 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. Also, by considering how the components are distributed, a SharePoint farm can be designed to 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 Server 2010 implementation vary based on several factors, such as the number of concurrent users, service applications in the farm, the expected uptime service-level agreement (SLA), and others. These factors dictate how many servers are needed in the SharePoint

---

<sup>2</sup> SharePoint collaboration capabilities - <http://sharepoint.microsoft.com/en-us/product/capabilities/communities/Pages/default.aspx>

<sup>3</sup> SharePoint publishing features - <http://technet.microsoft.com/en-us/library/ff628963.aspx>

farm and how the overall farm architecture looks. Based on these factors, SharePoint Server 2010 farm implementations can be classified as small farm<sup>4</sup>, medium farm<sup>5</sup> and large farm<sup>6</sup> deployments.

For more information on Dell SharePoint reference architecture guidance, performance sizing and best practices, refer to the white papers available at <http://www.dell.com/sharepoint>.

## VMware vSphere 5.0

VMware vSphere 5 includes the ESXi™ hypervisor as well as vCenter Server™ which is used to configure and manage VMware hosts. The vSphere 5 release includes several enhancements and new features compared to its predecessors. Some key capabilities include:

**vSphere High Availability (HA):** VMware HA provides automated detection of failures and movement of virtual machines (VM) from one host to another. A VMware HA cluster can include up to 32 nodes for virtual machine high availability.

**vSphere Distributed Switch:** VMware's vNetwork Distributed Switch spans across many ESXi hosts and aggregates networking to a centralized cluster level. The vNetwork Distributed Switch abstracts configuration of individual virtual switches and enables centralized provisioning, administration, and monitoring through VMware vCenter server.

**VMware Distributed Resource Scheduler (DRS):** VMware DRS technology enables vMotion to automatically achieve load balancing according to resource requirements. When VMs in a DRS cluster need fewer resources, such as during nights and weekends, DPM consolidates workloads onto fewer hosts and powers off the rest to reduce power consumption.

**VMware vMotion:** VMware vMotion technology provides real-time migration of running virtual machines (VM) from one host to another with no disruption or downtime.

**VMware Storage vMotion:** VMware Storage vMotion enables real-time migration of running VM disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.

**VMware vCenter Update Manager:** VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.

For more information on enhancements and new features in vSphere 5.0, see <http://www.vmware.com/files/pdf/products/vsphere/vmware-what-is-new-vsphere5.pdf>.

In this performance study, VMware vSphere 5.0 update 1 was used on the hypervisor hosts for deploying SharePoint web frontend and application roles as virtual machines. The VMware HA and DRS features

---

<sup>4</sup> SharePoint 2010 - Designing and implementing a small farm:  
<http://i.dell.com/sites/content/business/solutions/whitepapers/en/Documents/dell-small-sharepoint-farm.pdf>

<sup>5</sup> SharePoint 2010 - Designing and implementing a medium farm:  
[http://www.dell.com/downloads/global/services/dell\\_medium\\_sharepoint\\_farm.pdf](http://www.dell.com/downloads/global/services/dell_medium_sharepoint_farm.pdf)

<sup>6</sup> SharePoint 2010 - Designing and implementing a large farm  
[http://www.dell.com/downloads/global/services/dell\\_large\\_sharepoint\\_farm.pdf](http://www.dell.com/downloads/global/services/dell_large_sharepoint_farm.pdf)



were deployed to enable VM failover and resource optimization. Other features such as Distributed vSwitch and Storage vMotion were not used or validated.

In addition to the VMware vCenter management capabilities, the Dell PowerEdge plug-in for vCenter and Dell Compellent plug-in for vCenter enhance the management of both physical servers and Compellent storage and provide a single pane of management.

## Solution Architecture and Configuration

In a SharePoint farm architecture, dedicated servers were used to host each tier to provide process isolation and 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 SharePoint farm implementation uses multiple servers at all tiers of the farm deployment. In a very large SharePoint deployment, service applications such as search service are hosted in a central farm.

In this performance study paper, a SharePoint 2010 farm architecture was used to understand how several components of the farm perform at incrementing user loads. The farm roles such as WFE and APP server roles were implemented as virtual machines in this reference architecture.

The following sections describe several infrastructure components used in this farm architecture.

## Infrastructure Components

Within the scope of this paper, a single SharePoint Server 2010 farm was used to study the performance characteristics of SharePoint Server 2010 on Dell servers and storage.

Table 1. High-level overview of the farm infrastructure components

	SharePoint Farm Components
Blade Chassis	One Dell PowerEdge M1000e with: <ul style="list-style-type: none"><li>• M8024-K 10GbE Ethernet Switch modules in Fabric A1 and A2</li><li>• M5424-N 4/8Gbps FC SAN aggregator modules in Fabric B1 and B2</li></ul>
VMware Virtualization Hosts	Four Dell PowerEdge M620 Blade Servers
Database Server Physical Hosts	Two Dell PowerEdge M820 Blade Servers
Storage Array	Dell Compellent Series 40 Storage Center
Storage Fabric	Brocade 5100 FC Switches
Network Switches	Force10 S4810 10GbE switches
Load Balancer	F5 BIGIP Load Balancer - Virtual Edition (VE)

Figure 1. High-Level Overview of Farm Infrastructure Components

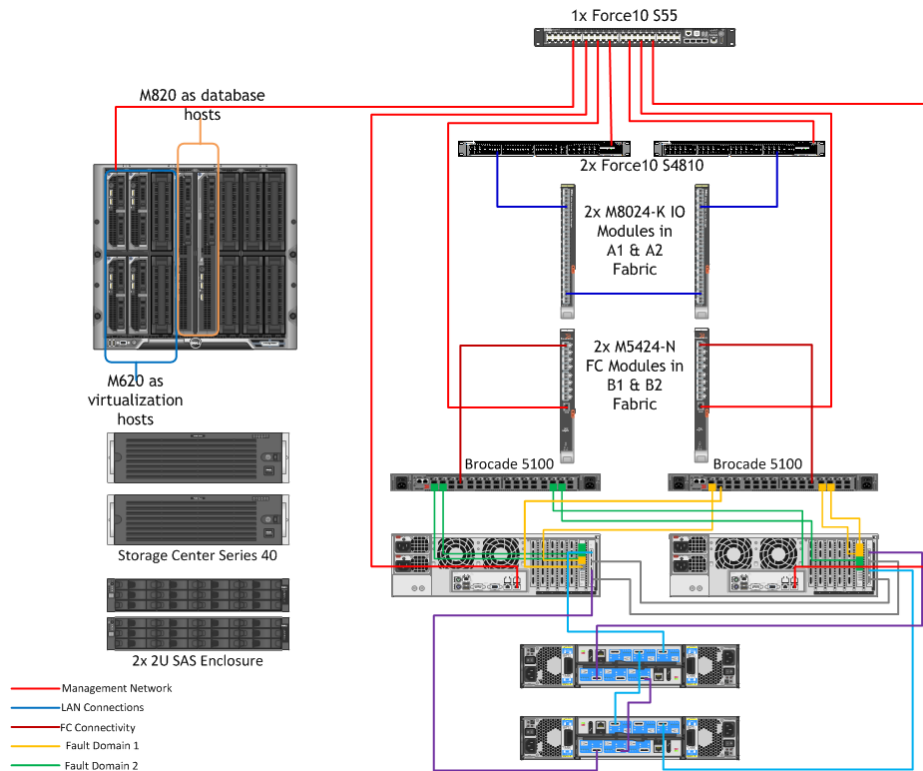


Figure 1 depicts the solution infrastructure components and the connectivity at a high level.

**NOTE:** The SharePoint farm implementation requires prerequisites such as Active Directory Domain Services (AD DS) and Domain Name Server (DNS). These components are not shown in this paper as a part of the reference architecture. Customers implementing SharePoint must have these components before starting SharePoint deployment.

The following sections describe the choice of the components shown in Figure 1 and describe in detail how these components are used in this performance study.

### Dell PowerEdge M1000e

The PowerEdge M1000e modular blade enclosure is the foundation for Dell's blade server architecture, providing one of the most energy-efficient, fully redundant, flexible, and manageable blade server platforms in the market for building any IT infrastructure. Flexible and scalable, the M1000e is designed to support future generations of blade technologies regardless of processor/chipset architecture. The M1000e is optimized for use with all Dell PowerEdge Blades including the M620 and M820 blades servers. Features include:

- **Energy Efficiency**  
M1000e is built on Dell's energy smart technology, which can help you increase capacity and lower operating costs while delivering better performance per watt.
- **Effortless Scalability**  
With scale-on-demand switch design and additional I/O slots and switch options, the M1000e provides the flexibility to meet the increasing demand for I/O consumption. Plus, Dell's FlexIO

modular switch technology offers great scalability to provide additional uplink and stacking functionality.

- **Powerful Management Tool**

M1000e includes centralized management controllers, dynamic power management, and real-time reporting service, which allow IT administrators to manage and monitor multiple enclosures and blades from a single console. The multi-chassis management feature enables enterprise administrators to monitor and manage multiple blade chassis from a single console without any additional cabling or software agent requirements.

For more information on Dell PowerEdge M1000e blade enclosure, see

<http://www.dell.com/downloads/global/products/pedg/en/poweredge-M1000e-Spec-Sheet.pdf>

## Dell PowerEdge M620

Dell PowerEdge M620 is the 12<sup>th</sup> generation, half-height blade server that supports up to 2 socket Intel Xeon E5-2600 family processors and 768GB of physical memory. This server supports up to 2TB of internal storage when using 2x 2.5 inch 1TB SATA hard drives and supports a wide variety of IO options to connect to an external iSCSI or Fibre Channel storage enclosure. The unprecedented memory density and the processing capabilities make this server one of the best choices for a virtualization host.

Within this performance study, four Dell PowerEdge M620 servers were used as the virtualization hosts to deploy SharePoint 2010 web frontend (WFE) and application server (APP) roles as virtual machines. For network connectivity, Broadcom 57810-K dual port 10GbE KR blade adapters were used. For FC SAN connectivity, QME2572 8Gbps Fibre Channel I/O mezzanine cards were used.

For more information on Dell PowerEdge M620, see

[http://www.dell.com/downloads/global/products/pedg/m620\\_spec\\_sheet.pdf](http://www.dell.com/downloads/global/products/pedg/m620_spec_sheet.pdf).

## Dell PowerEdge M820

Dell PowerEdge M820 is the 12<sup>th</sup>-generation, full-height blade server that supports up to 4 socket Intel Xeon E5-4600 family processors and 1.5TB of physical memory when using 48x 32GB memory DIMMs. This server supports up to 2TB of internal storage when using 2x 2.5 inch 1TB SATA hard drives and supports a wide variety of IO options to connect to an external iSCSI or Fibre Channel storage enclosure.

The amount of processing and physical memory capacity available in a Dell PowerEdge M820 makes it a good choice for the database hosts in this reference architecture. In this performance study, two Dell PowerEdge M820 servers were used as the database servers to support SharePoint 2010 databases. Although, the entire capacity is not used, the choice of M820 for the database tier provides enough room for growth to support future needs. For network connectivity, Broadcom 57810-K dual port 10GbE KR blade adapters were used, and QME2572 8Gbps Fibre Channel I/O mezzanine cards were employed for fibre channel SAN connectivity.

For more information on Dell PowerEdge M820, see

[http://www.dell.com/downloads/global/products/pedg/m820\\_spec\\_sheet.pdf](http://www.dell.com/downloads/global/products/pedg/m820_spec_sheet.pdf).

## Dell PowerConnect M8024-K 10Gb Ethernet Switch

Dell PowerConnect M8024-K 10Gb Ethernet switch is a 10GbE Ethernet switch IO module for the Dell PowerEdge M1000e Blade enclosure. This switch provides 10Gb Ethernet connectivity for up to sixteen M-Series blade servers equipped with the latest KR-based 10Gb network daughter cards (NDCs) or

Implementing Microsoft SharePoint Server 2010 on Dell 12<sup>th</sup>-Generation Blade Servers and VMware vSphere 5.0

mezzanine cards. Dell M8024-K is designed to support cutting-edge networking features such as NIC partitioning (NPAR), SR-IOV, and iSCSI offload.

Dell PowerConnect M8024-K provides 16 internal ports for blade server connectivity and up to 8 SFP+ external ports to connect to the top-of-rack (TOR) network connectivity. This switch can operate in either a simple mode or a normal mode. In the simple mode, this switch acts as an IO aggregator and provides an overall bandwidth of 80Gbps to the TOR network fabric. In the normal mode, M8024 acts as a usual 10GbE Ethernet switch.

In this performance study paper, this switch was used in normal mode to connect to the top-of-rack Dell Force10 S4810 switches.

For more information on Dell PowerConnect M8024-K, see [http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/PowerConnect\\_M8024\\_K\\_10Gb\\_Ethernet\\_Switch.pdf](http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/PowerConnect_M8024_K_10Gb_Ethernet_Switch.pdf)

### Brocade M5424-N 8Gbps Fibre Channel Switch

The Brocade M5424-N 8Gbps Fibre Channel switch for Dell PowerEdge M1000e blade enclosure connects the blade servers to the top of the rack fibre channel fabric or fibre channel-capable devices such as Dell Compellent Series 40 Storage Center controllers. On M5424-N FC Switch, Inter-Switch Link (ISL) Trunking allows all eight external SAN ports to be combined to form a single, logical ISL with an aggregate bandwidth of 64Gb delivering scalable I/O bandwidth utilization and load balancing.

In this performance study paper, the Brocade M5424-N FC switches in the blade enclosure were connected to the top-of-rack Brocade 5100 switches using a 4-port ISL trunk.

For more information on Brocade M5424 FC Switch, see <http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/brocade-m5424-data-sheet.pdf>

### Brocade 5100 Fibre Channel Switch

Brocade 5100 is a 40-port Fibre Channel (FC) switch that offers 8Gbps of throughput per port and offers Adaptive Networking services, including adaptive fabric quality of service (QoS), traffic isolation and ingress rate limiting, enabling organizations to maintain optimized application performance in virtual environments.

This performance study paper implemented 2x Brocade 5100 FC switches to connect to the Brocade M5424-N blade switches and the Dell Compellent Series 40 Storage Center controllers.

For more information on Brocade 5100 FC switches, see <http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/switch-brocade-5100-datasheet.pdf>

### Dell Force10 S55 Network Switches

The Dell Force10 S-Series S55 1/10 Gigabit Ethernet (GbE) top-of-rack (ToR) switch is optimized to lower operational costs while increasing scalability and improving manageability at the network edge. Designed for high-performance data center applications, the S55 leverages a non-blocking architecture that delivers line-rate, low-latency L2 and L3 switching to eliminate network bottlenecks. The Dell Force10 S55 switch provides 48 Gigabit Ethernet ports and supports up to four 10GbE uplinks when using an additional rear module.

In this performance study, a Dell Force10 S55 switch was used for the management network. The infrastructure components' management interfaces were connected to this switch as shown in Figure 1.

For more information on Dell Force10 S55 Ethernet switches, see

[http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell\\_Force10\\_S55\\_Spec\\_Sheet.pdf](http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell_Force10_S55_Spec_Sheet.pdf)

## Dell Force10 S4810 Network Switches

The Dell Force10 S4810 is a high-density 10/40GbE ToR switch with 48 dual-speed 1/10GbE (SFP+) ports, four 40GbE (QSFP+) uplinks, and up to 64 10GbE ports with breakout cables. This switch provides an Open Automation Framework<sup>7</sup> that adds VM-awareness as well as automated configuration and provisioning capabilities to simplify the management of a virtual-network environment.

In this performance study, two Dell Force10 S4810 Ethernet switches were used as the ToR switches. The Dell M8024-K blade switches were connected to these TOR switches to enable farm connectivity at the TOR. The Force10 S4810 switches provide higher bandwidth uplinks and are a perfect choice for connectivity to the datacenter core switches.

For more information on Dell Force10 S4810 Ethernet switches, see

[http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell\\_Force10\\_S4810\\_Spec\\_sheet.pdf](http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell_Force10_S4810_Spec_sheet.pdf)

## Dell Compellent Series 40 Storage Center

Dell Compellent Series 40 Storage Center manages data based on the Fluid Data architecture. Storage center optimizes the placement of data at different tiers of the storage based on data access patterns and performance needs. Storage Center also enables data protection with the help of features such as instant replays and replication. Dell Compellent Storage Center allows mixing of SSD, FC, and SAS drive technologies as well as iSCSI, FCoE, and FC frontend interconnects.

Dell Compellent Storage Center provides a fully-virtualized storage platform that includes:

- Storage Virtualization that abstracts and aggregates all resources across the entire array, providing a high-performance pool of shared storage.
- Thin Provisioning and Automated Tiered Storage to deliver optimum disk utilization and intelligent data movement.
- Space-Efficient Snapshots and Thin Replication for continuous data protection without wasted capacity. Built-in automation and unified storage management to streamline storage provisioning, management, migration, monitoring, and reporting

---

<sup>7</sup> Open Automation Framework: [http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell\\_Force10\\_Open\\_Automation\\_Framework\\_Spec\\_Sheet.pdf](http://i.dell.com/sites/docontent/shared-content/data-sheets/en/Documents/Dell_Force10_Open_Automation_Framework_Spec_Sheet.pdf)

The following table gives a high level overview of Dell Compellent Storage Center arrays hardware features.

**Table 2. Dell Compellent hardware capabilities**

Feature	Dell Compellent Storage Center SAN
<b>Storage Capacity</b>	739TB of raw storage capacity with FC drives or 1,920TB with SAS drives per system.
<b>Hard Drives</b>	Each 3.5" FC enclosure holds up to 16 FC or 14 SSD drives (with a FC drive in bays 1 and 16). Each 3.5" SAS enclosure holds up to 12 SAS drives. Each 2.5" SAS enclosure holds up to 24 SAS drives.
<b>3.5" Drive Performance and Capacities</b>	15,000 RPM SAS drives available in 450GB and 600GB 15,000 RPM FC drives available in 300GB and 600GB 10,000 RPM FC drives available in 450GB 7,200 RPM SAS drives available in 1TB, 2TB, and 3TB
<b>2.5" Drive Performance and Capacities</b>	15,000 RPM SAS drives available in 146GB and 300GB 10,000 RPM SAS drives available in 600GB and 900GB 7,200 RPM SAS drives available in 1TB SAS Solid State Drive (SSD) available in 200GB
<b>Expansion Capabilities</b>	Each system supports up to 1,232 FC drives or 960 SAS drives
<b>Host Connectivity</b>	Maximum of 22 frontend ports per controller. A switch is recommended for all configurations and required for clustered controller configurations.
<b>RAID Levels</b>	Support for RAID levels 0, 5, 6, and 10. Any combination of RAID levels can exist on a single Storage Center. Multiple RAID levels can exist on the same storage tier within an array. Storage is virtualized and pooled across arrays without RAID group limitation.

Within the scope of this performance study, Dell Compellent Storage Center Series 40 Fibre channel storage arrays were used. As part of this storage deployment, only tier 1 storage was used with a RAID10 and RAID5-9 storage profile with no replay configuration.

The Storage Center features such as Data progression, FastTrack, and Instant Replays were not used as part of the performance study. However, these features enable a great deal of long-term savings and reduce the total cost of ownership. Using data progression and the dynamic block architecture, the cold data that has not been accessed or modified for a long time automatically moves to tier 3 of the storage. The tier 3 storage is the inexpensive near-line SAS storage that provides higher capacity. This

Implementing Microsoft SharePoint Server 2010 on Dell 12<sup>th</sup>-Generation Blade Servers and VMware vSphere 5.0

automatic data movement between tiers reduces the time required to access the most critical data while storing the old and rarely-accessed voluminous data on inexpensive tier 3 storage, thereby reducing the total cost of ownership.

For more information on Dell Compellent Series 40 Storage Center, see [http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/DE5ENTC15917\\_DS8\\_Storage\\_Center\\_TS\\_12.pdf](http://i.dell.com/sites/content/shared-content/data-sheets/en/Documents/DE5ENTC15917_DS8_Storage_Center_TS_12.pdf)

## BIG-IP Local Traffic Manager (LTM) Virtual Edition

Within the scope of this paper, the farm configuration used F5 Networks<sup>®</sup> BIG-IP<sup>®</sup> Local Traffic Manager™ (LTM) virtual edition load balancers to enable load balancing across Web frontend nodes. The BIG-IP load balancers offer additional features – such as SSL acceleration, protocol optimization, rapid deployment templates, etc – that are unavailable in the native software network load balancing (NLB) clusters. Hence, two F5 BIG-IP virtual load balancer virtual machines were deployed in an active - standby configuration mode.

As a part of this study, no custom load balancer profiles were defined. The farm configuration used the out-of-box acceleration functionality and BIG-IP Application (iAPP) templates, thereby simplifying the administrative tasks and shortening the required set-up time. In this performance study, OneConnect profile and strict updates were disabled and no other application profiles were deployed.

For more information on BIG-IP LTM virtual edition, see <https://devcentral.f5.com/HotTopics/VirtualEdition/tabid/1082200/Default.aspx>

The following sections describe the hardware configuration for the hypervisor hosts and the database servers and describe the choices made in configuring the virtual infrastructure for high availability and optimal resource utilization.

The hardware configuration of the blade servers – Dell PowerEdge M620 and M820 – is shown in Table 3.

**Table 3. Hardware configuration of blade servers**

	Dell PowerEdge M620	Dell PowerEdge M820
<b>Processor</b>	2x Intel <sup>®</sup> Xeon <sup>®</sup> E5-2660 2.2GHz 8 core, 20MB L3 Cache, Turbo, Hyper Threading	2x Intel <sup>®</sup> Xeon <sup>®</sup> E5-2660 2.2GHz 8 core, 20MB L3 Cache, Turbo
<b>Memory</b>	64GB (8x 8GB) DDR3 1600Mhz RDIMMS	96GB (8x 16GB) DDR3 1600Mhz RDIMMS
<b>Network controller</b>	Broadcom 57810 10GbE Blade NDC Ethernet Controller	Broadcom 57810 10GbE Blade NDC Ethernet Controller
<b>Storage Controller</b>	Dell PERC H310P mini	Dell PERC H310P mini
<b>Fibre Channel HBA</b>	Qlogic QME2572	Qlogic QME2572
<b>Out-of-band (OOB) management</b>	iDRAC7 Enterprise	iDRAC7 Enterprise
<b>Purpose of these servers in the solution</b>	Hypervisor hosts - VMware vSphere 5.0 update 1	Database Server - SQL Server 2008 R2 SP1 CU7



## Virtual Infrastructure Configuration

The SharePoint 2010 farm web frontend (WFE) and application (APP) server roles were implemented as virtual machines on top of the hypervisor servers in this reference architecture.

Table 4. WFE, APP, and BIG-IP virtual machine configuration

	Web frontend	Application	BIG-IP LTM VE
Number of Processors	8 vProcs	4 vProcs	2 vProcs
Memory	16GB	8GB	4GB
Network	4x VMXNET3 <sup>8</sup> vNIC	3x VMXNET3 vNIC	2x VMXNET3 vNIC
Storage Controller	LSI Logic SAS	LSI Logic SAS	LSI Logic Parallel
Virtual Disk for OS	250GB Thick Eagerly Provisioned VMDK	250GB Thick Eagerly Provisioned VMDK	250GB Thick Eagerly Provisioned VMDK

The configuration of these SharePoint farm roles along with BIG-IP LTM virtual edition VM configuration details is as shown in Table 4.

Four Dell PowerEdge M620 blade servers were used to deploy VMware ESXi and configured to be a part of a single vSphere cluster. The SharePoint and BIG-IP virtual machines — as described in Table 4 — were deployed in this cluster. These virtual machines were placed on the VMware ESXi hosts in an optimal way to ensure equal load distribution and failover or high availability configuration. VMware High Availability (HA) and Distributed Resource Scheduler (DRS) features were enabled to ensure continuous availability and optimization of the virtual infrastructure resource utilization.

---

<sup>8</sup> VMXNET3 virtual adapters provide enhanced traffic management capabilities including multiqueue support, IPv6 offloads, and 10GbE support.

Figure 2. Virtual Machine Placement

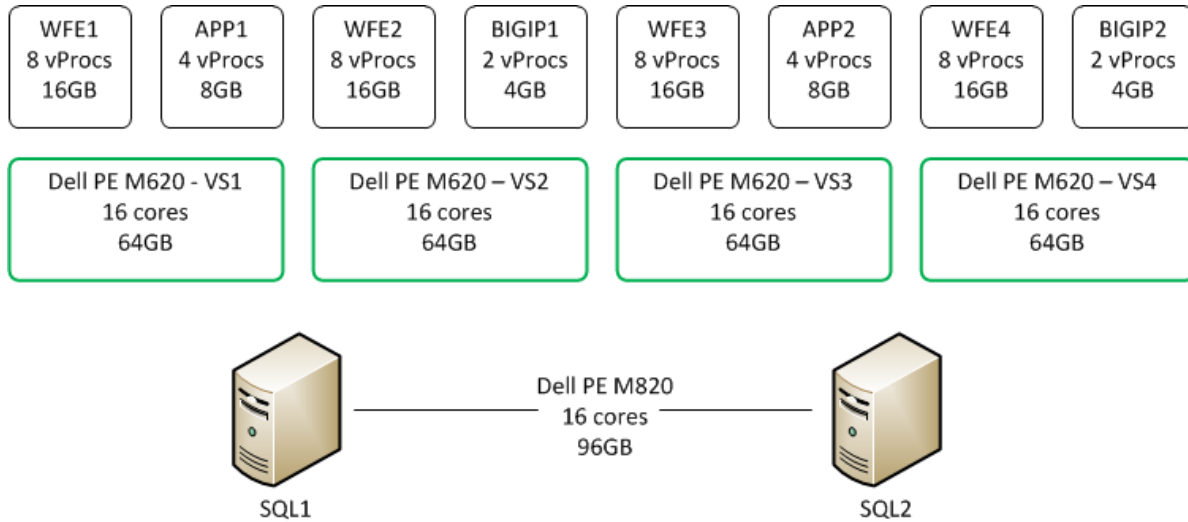


Figure 2 depicts the way in which different SharePoint roles were deployed within this performance study. This recommended placement of the web frontend and the application roles ensures equal load distribution among hosts. The preferred node configuration and HA or failover configuration was managed using VMware Distributed Resource Scheduler (DRS) affinity rules. These affinity rules govern the preferred owner allocation for virtual machines and virtual machine separation in the event of a failover. Within DRS affinity rules, “Run VMs on Hosts” and “Separate Virtual Machines” rules were deployed to achieve the VM allocation on the hypervisor’s hosts and to control VM separation.

The DRS rule “Run VMs on hosts” ensures that the virtual machines always start on preferred clustered hosts for optimizing resource utilization during normal operation and the “Separate Virtual Machines” rule ensures that the virtual machines that are either clustered or deployed in a highly-available configuration – such as the two SharePoint application VMs and the BIG IP load balancer VMs – do not get migrated to the same host during VMware HA.

Figure 3. VMware DRS Rules

Name	Type	Defined by
<input checked="" type="checkbox"/> VM4-Host4-Mapping	Run VMs on Hosts	User
VM Group 4	Cluster VM Group	
Cluster Host 4	Cluster Host Group	
<input checked="" type="checkbox"/> VM3-Host3-Mapping	Run VMs on Hosts	User
VM Group 3	Cluster VM Group	
Cluster Host 3	Cluster Host Group	
<input checked="" type="checkbox"/> VM2-Host2-Mapping	Run VMs on Hosts	User
VM Group 2	Cluster VM Group	
Cluster Host 2	Cluster Host Group	
<input checked="" type="checkbox"/> VM1-Host1-Mapping	Run VMs on Hosts	User
VM Group 1	Cluster VM Group	
Cluster Host 1	Cluster Host Group	
<input checked="" type="checkbox"/> BIG-IP VM Separation	Separate Virtual Machines	User
SP-Q3-BIGIP1		
SP-Q3-BIGIP2		
<input checked="" type="checkbox"/> APP VM Separation	Separate Virtual Machines	User
SP-Q3-APP1		
SP-Q3-APP2		

As shown in Figure 3, SharePoint WFE, APP, and BIG-IP VMs were grouped into VM groups and the DRS rules were used to assign VM groups to different cluster hosts. The following table shows the DRS group configuration in the vSphere cluster.

Table 5. Virtual machine to cluster host mapping

	Virtual Machines in the VM Group	Cluster Host mapped to the VM Group
VM Group 1	SP-Q3-WFE1	SP-Q3-VS1
	SP-Q3-APP1	
VM Group 2	SP-Q3-WFE2	SP-Q3-VS2
	SP-Q3-BIGIP1	
VM Group 3	SP-Q3-WFE3	SP-Q3-VS3
	SP-Q3-APP2	
VM Group 4	SP-Q3-WFE4	SP-Q3-VS4
	SP-Q3-BIGIP2	

As shown in Figure 2, the SQL server roles required for the SharePoint content databases were hosted on physical servers (Dell PowerEdge M820 blade servers). This choice was made to enable higher concurrent user load support. Within this performance study paper, Microsoft SQL Server 2008 R2 was

deployed in a failover cluster configuration to enable high availability for the SharePoint content databases.

For hosting the virtual hard disk files (vmdk) for each of the virtual machines, a single 3 terabytes (TB) volume was created and assigned to all the hosts in the vSphere cluster as the VM data store.

**NOTE:** Backup and recovery procedures or tools are not a part of this infrastructure or discussed in this paper. The Dell Compellent storage array is a fully-redundant storage system with many available safeguards, but Dell highly recommends regular backups of the SharePoint environment to protect against failures and allow restoration or recovery of content. For assistance in planning backup policies, contact your Dell representative or visit [www.dell.com/storage](http://www.dell.com/storage) for more information

## Network architecture


The network traffic in this Sharepoint 2010 solution is comprised of six distinct types:

- Management and Hypervisor Networks
  - VM & ESXi Host Management Network
  - vMotion Network
- VM networks for SharePoint
  - SharePoint Farm public network
  - SQL Cluster Private communication
  - WFE/APP to SQL communication Network
  - WFE to APP communication Network

The following sections include an in-depth discussion of various networks deployed in this farm, configuration, and design of network for ESXi hosts and the SQL Server hosts.

The following high-level network diagram depicts various network components used as a part of this performance study, connections among these components, and virtual networks defined for achieving traffic isolation.

Figure 4. High-Level Network Architecture



Within this architecture, the Dell PowerConnect M8024-K 10Gb blade Ethernet switches – placed in A1 and A2 IO fabrics of Dell M1000e Blade enclosure – were used for inter-blade and top-of-rack (TOR) connectivity. These switches were configured to be in the normal operation mode and were connected together using a Link Aggregation Group (LAG) created using two 10GbE external ports. On the Dell PowerConnect M8024-K blade switches, multiple spanning tree<sup>9</sup> (MSTP) was configured to channel VLAN traffic over different interfaces – in this architecture, the port groups 1 and 2 as shown in Figure 4.

The Dell Force10 S4810 10GbE Ethernet switches were used for the top-of-rack network. As shown in Figure 4, the M8024-K switches in the blade chassis were connected to the TOR switches by creating a LAG using two external ports on the M8024-K switches. The two S4810 switches were connected together using two 40GbE ports to provide 80Gbps of aggregated bandwidth between the switches. These TOR switches were used to provide uplinks to the core network containing other infrastructure elements required for the SharePoint farm such as Active Directory and DNS.

<sup>9</sup> MSTP configuration on M8024-K-  
<http://support.dell.com/support/edocs/network/m8024k/en/ucg/html/spantree.htm>

As shown in Figure 4, the traffic on the LAN is segregated into six virtual LANs (VLANs); one VLAN each for:

- Virtual Machine and ESXi host management traffic (VLAN 10)
- SharePoint farm public network traffic (VLAN 20)
- WFE to APP communication traffic (VLAN 30)
- WFE/APP to SQL communication traffic (VLAN 40)
- SQL Cluster Private traffic (VLAN 50)
- VMware vMotion traffic (VLAN 60)

VLAN tagging was performed through switch tagging as well as host-level tagging. Traffic isolation among various roles in the SharePoint farm architecture was enforced using HOSTS file. For example, for all web frontend (WFE) to application server communication was restricted to use the WFE to APP VLAN only. Similarly, all traffic between WFE/APP servers and SQL servers roles was restricted to flow through the WFE/APP to SQL virtual network only.

Also, the LAG – port group 1 in Figure 4 – between the blade switches was configured to allow only inter-chassis VM and host traffic only. This ensures that the network traffic related to vMotion or communication among the SharePoint server roles does not reach top-of-rack switches.

### VMware ESXi Host (M620) Networking

As shown in Table 3, Broadcom 57810-K 10GbE Ethernet Network Daughter Cards (NDC) were used in the Dell PowerEdge M620 blade servers. These network adapters provide two 10GbE Ethernet ports that connect to the internal ports of the M8024-K switches in fabrics A1 and A2. The Network Partitioning (NPAR) feature of these adapters was used to create multiple virtual functions each with a specific bandwidth allocation. NPAR allows splitting the 10GbE bandwidth to divide and reallocate bandwidth and resources as needed to a maximum of four partitions.

Within this reference architecture, the network controllers on the ESXi hosts were configured to create three network partitions on each port of the NDC. The bandwidth and weight allocations are as follows:

Table 6. NPAR bandwidth allocation

Broadcom NDC Ports	Network Parition Name	Bandwitdh %	Weightage %
NIC_Port-1	NIC 0	100	10
	NIC 2	100	20
	NIC 4	100	70
	NIC 6	1	0
NIC_Port-2	NIC 1	100	10
	NIC 3	100	20

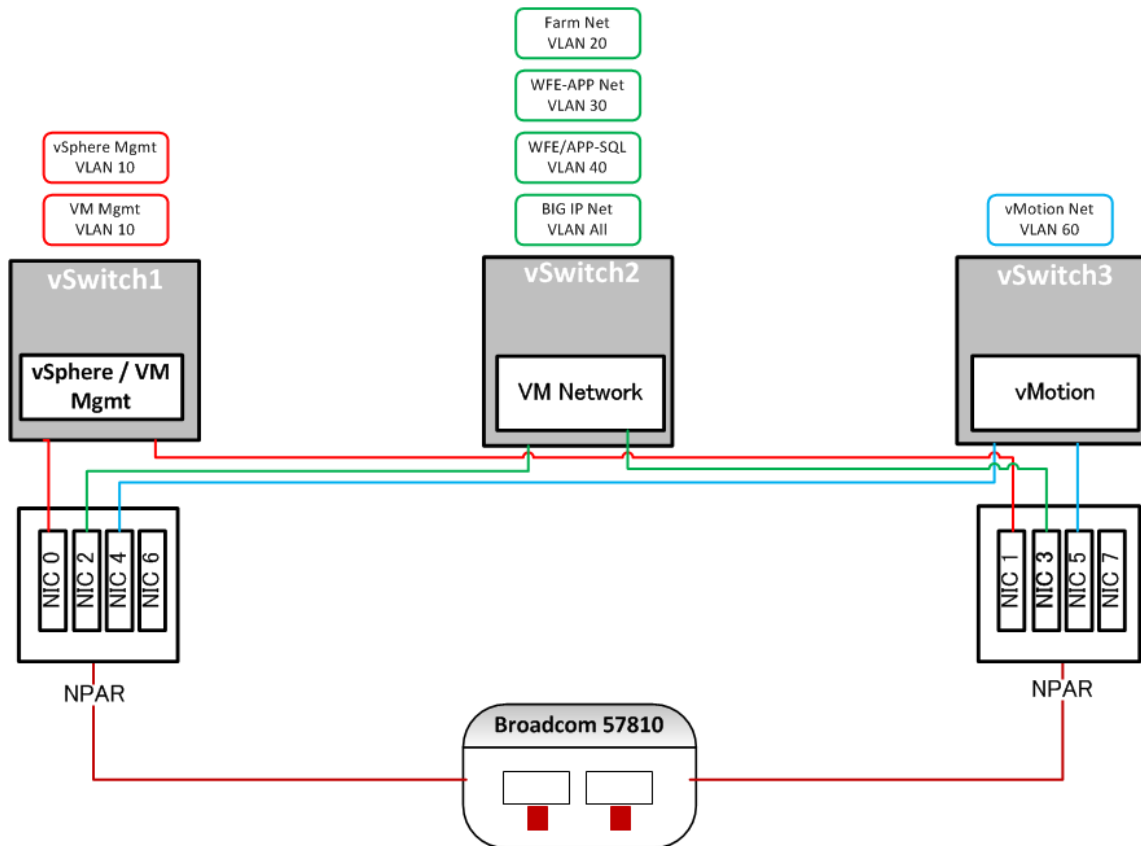
Broadcom NDC Ports	Network Parition Name	Bandwitdh %	Weightage %
	NIC 5	100	70
	NIC 7	1	0

Although there were only three partitions used in this study, the configuration of NPAR requires us to specify bandwidth allocations for the fourth partition as well. For the ESXi host network deployment, the 4<sup>th</sup> partition (NIC 6 and NIC 7 in Table 6) on each physical port of the network adapter (Broadcom 57810-K) was allocated zero bandwidth as it was not intended to be a part of any vSwitch configuration.

**NOTE:** Step-by-step guidance to configure NPAR is outside the scope of this paper. For more information on NPAR on Dell PowerEdge servers, see <http://www.broadcom.com/collateral/faq/2CS57712-K-FAQ102-R.pdf>

The following figure depicts the logical network architecture on the ESXi hosts used in this performance study. As shown in the the following figure, three of the four available network partitions were used to create VMware vSwitches and add both vmkernel and virtual machine port groups as required.

Figure 5. VMware ESXi Host Networking



As shown in Figure 5, a combination of network partitions from each port of the NDC were used to create the vSwitches. The following table shows a mapping of network partitions to vSwitch configuration on the ESX hosts.

**Table 7. Network partitions to vSwitch mapping**

vSwitch Name	Network Partitions used
vSwitch#0	NIC 0
	NIC 1
vSwitch#1	NIC 2
	NIC 3
vSwitch#2	NIC 4
	NIC 5

The three vSwitches, as shown in Figure 5, were created based on the network traffic segregation provided by VLANs. The following table shows the mapping of VLANs to virtual switches and the purpose of these VLANs in the SharePoint architecture:

**Table 8. VLAN to virtual switch mapping - LAN**

VLAN name	Virtual switch name	Purpose	Virtual machine mapping
VM_Mgmt_Net	vSwitch#0	Used only for the host and VM Management and AD / DNS connectivity	WFE1, WFE2, WFE3, WFE4, APP1, APP2, BIGIP1, and BIGIP2
Host_Mgmt_Net			
VM_WFE-APP_Net	vSwitch#1	Used for WFE to APP communication and WFE/APP to SQL communication	WFE1, WFE2, APP1, and APP2
VM_WFE_APP-SQL-Net			
VM_Farm_Public_Net	vSwitch#1	Used only by WFE for public farm connectivity	WFE1 and WFE2
vMotion_Net	vSwitch#2	Used by ESXi hosts for performing vMotion	NA

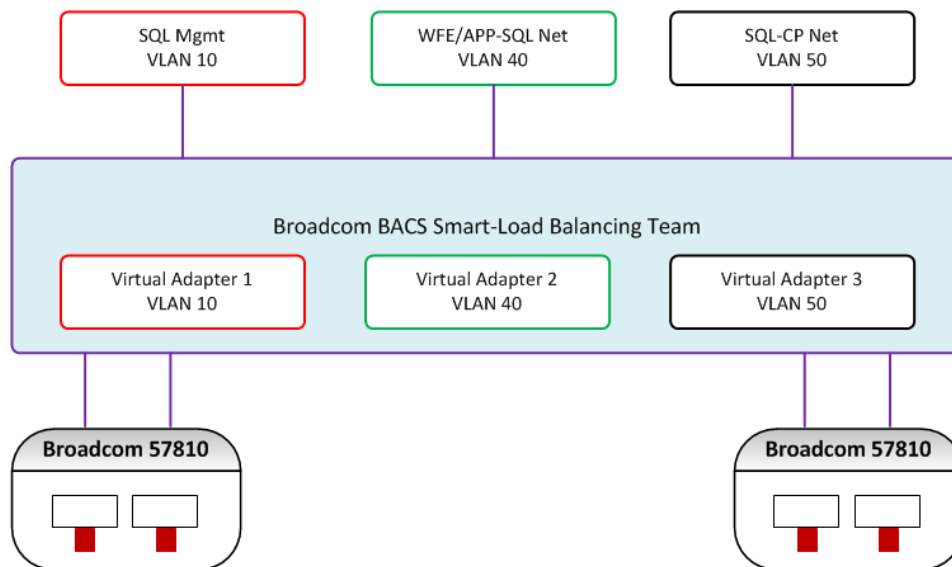


The VLAN segregation for the different network types helps isolate the traffic among components of the SharePoint server farm. Also, as an application best practice, the WFE/APP to SQL communication and the WFE to APP communication traffic has been isolated as discussed earlier.

## Microsoft SQL Server (M820) Networking

As shown in Figure 1, Dell PowerEdge M820 blade servers were used as the database servers in the SharePoint farm architecture. In this reference architecture, the database role was not virtualized to utilize the physical server capacity available on the M820 blade servers. On these database hosts, Similar to the ESXi hosts, Broadcom 57810-K 10GbE blade NDC were used for LAN connectivity. The 2x 10GbE NDC per database host provided four 10GbE network connections. The following diagram shows the high-level network design on the database hosts.

Figure 6. High-Level Network Design on the Database Hosts



As shown in Figure 6, the four 10GbE network ports were teamed using the BACS utility and teamed as a Smart Load Balancing (SLB) team to provide load balancing as well as failover. Three virtual adapters were created with tagged VLANs to isolate various types of traffic from the database hosts.

Table 9. SQL network connections overview

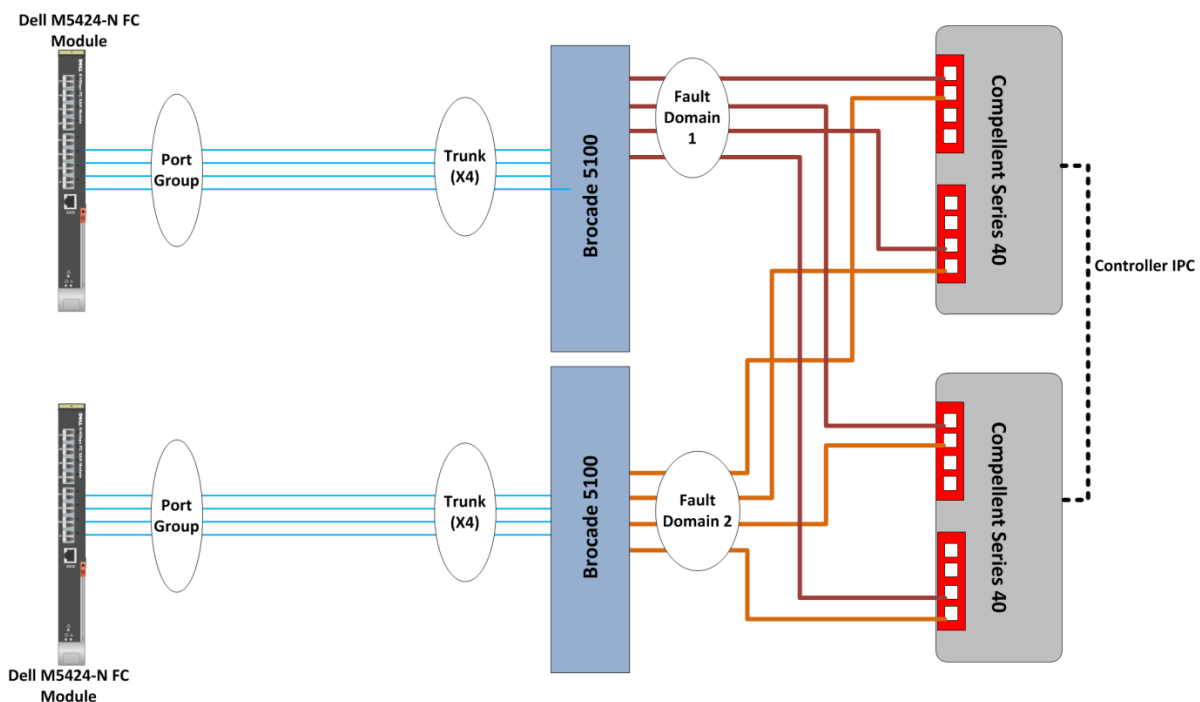
Virtual Adapter Name	Purpose
SQL_Mgmt	Used only management and AD/DNS connectivity
WFE/APP-SQL_Net	Used for SQL to WFE/APP VM traffic
SQL_CP_Net	Used for SQL cluster private communication

## Storage architecture

As shown in Figure 1 and described in Table 1, Dell Compellent Series 40 Storage Center Fibre Channel (FC) storage arrays were used in this performance study. For hosting the VM storage and SharePoint databases, two 2.5'' 15K SAS 300GB disk enclosures each with 24 drives were used. As described in Table 3, the QLogic QME2572 blade mezzanine controllers are used to connect the blade servers to the storage fabric. These modules were connected to the in-chassis Dell M5424-N 8/4 Gbps Fibre Channel IO modules in fabrics B1 and B2. The M5424-N FC switches, as described earlier, provide the access gateway functionality to connect to the top-of-rack Brocade 5100 FC switches.

## Storage Frontend Architecture

Figure 7. High-Level Storage Frontend Architecture



The high-level storage frontend architecture used in this performance study is shown in Figure 7. As described earlier, the Dell FC SAN Module is configured to operate as a port aggregator or access gateway for aggregating 16 internal ports to 2, 4, or 8 external ports. Dell FC SAN module port aggregator operates in access gateway mode for providing N\_Port ID virtualization (NPIV) functionality. The following are the FC ports that the FC SAN Module uses:

- F\_Port - internal fabric port that connects a blade server (HBA)
- N\_Port - external node port that connects to a switch

The FC SAN Module uses internal F-port to external N-port mappings as configured to direct traffic from the blade server (HBAs) to the fabric. Using Brocade Port Trunking, four external N\_Ports were combined to form a single logical port. This provides 32 Gbps of bandwidth from the blade chassis to the top-of-rack FC fabric. Fibre channel IO modules in the Dell Compellent storage arrays were connected to the Brocade 5100 FC switches.

In this performance study, one trunk group was configured on each 5100 with two FC ports per trunk group. The M5424-N SAN module automatically forms a trunk with the corresponding external N-ports.

**Table 10. Trunk and N\_Port mapping**

Blade Module	Trunk Group	N_Port used
M5424-N in Fabric B1	Trunk Group 1	N_17 and N_18
	Trunk Group 2	N_19 and N_20
M5424-N in Fabric B2	Trunk Group 1	N_17 and N_18
	Trunk Group 2	N_19 and N_20

In each Dell FC SAN module, the blades in odd numbered slots are mapped to one trunk group and blades in even numbered slot are mapped to another trunk group.

Blade Module	Blade Server Slots	F_Ports Used	Trunk Group
M5424-N in Fabric B1	Slot 1, Slot 5, Slot 9, Slot 13	F_1, F_5, F_9, F_13	Trunk Group 1
	Slot 2, Slot 4, Slot 10, Slot 12	F_2, F_4, F_10, F_12	Trunk Group 2
M5424-N in Fabric B2	Slot 1, Slot 5, Slot 9, Slot 13	F_1, F_5, F_9, F_13	Trunk Group 1
	Slot 2, Slot 4, Slot 10, Slot 12	F_2, F_4, F_10, F_12	Trunk Group 2

With trunking, any link within a trunk group can go offline or become disabled, but the trunk remains functional with the remaining links and no re-configuration is required. If the trunk fails for some reason, the Dell SAN module will remap the internal F-ports of failed N-port trunk to the pre - configured stand-by N-port trunk group. Note that in this case, the FC ports actually failover from one FC fabric to another.

As illustrated in Figure 7, the trunk groups were connected to the Brocade 5100 switches. Trunk group with odd numbered of blades from fabric B1 and trunk group with even numbered blades from fabric B2 are connected to the top-of-rack Brocade 5100 switch. Similarly, trunk group with even numbered blades from fabric B1 and trunk group with odd numbered blades from fabric B2 are connected to the bottom Brocade 5100 switch. This configuration optimally distributes traffic between the two SAN fabrics.

The fault domains - as shown in Figure 7- combine ports from each controller to provide redundancy and port failover using the virtual ports feature of Compellent storage. On both ESXi hosts and the

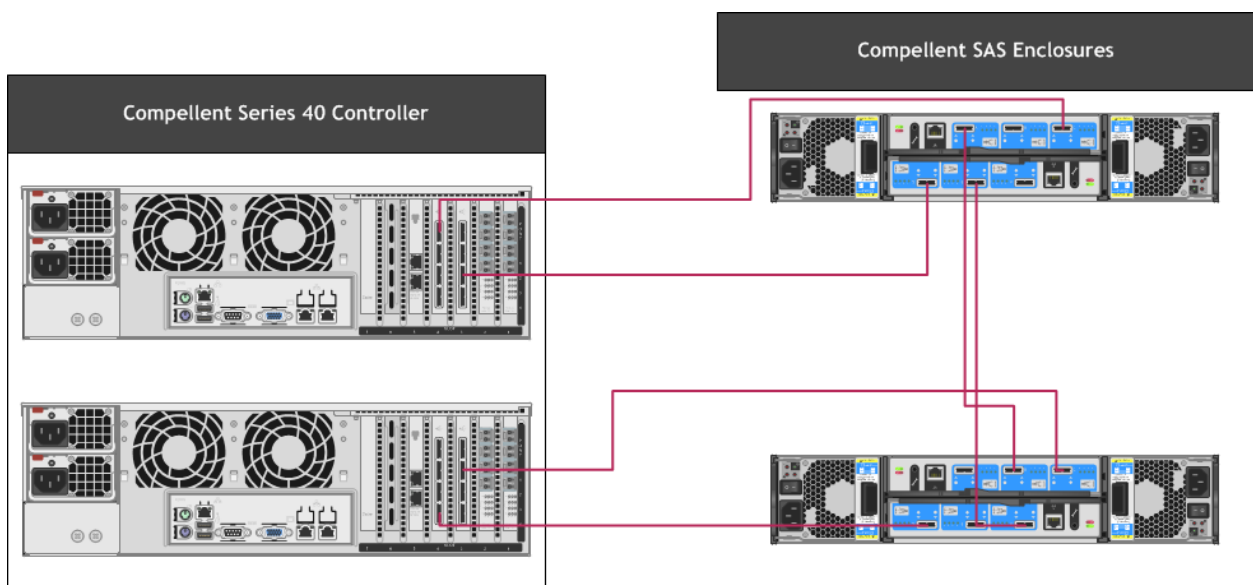
database hosts, multipathing was used to enable IO load balancing and path failover for the hosts. In this reference architecture, VMware Native Multipath Plug-In (NMP) is used to provide multi-pathing. Path Selection Plug-Ins (PSPs) run with the VMware NMP and is responsible for choosing a physical path for I/O requests. On the database hosts, native Windows MPIO feature was used to configure multipathing.

## Storage Backend Architecture

As mentioned earlier, within this reference architecture, two 15K SAS 2.5'' 300GB drive enclosures were used to host the VM store and SharePoint content databases. These enclosures were connected to the Dell Compellent Series 40 Storage Center controllers using 6Gbps SAS connectivity.

The following diagram depicts the high-level architecture of the storage backend architecture.

Figure 8. High-Level Storage Backend Architecture



## SharePoint Farm Configuration

The SharePoint 2010 farm configuration, as shown in Figure 1, included two Web Frontend (WFE), two Application Servers (APP), and two database servers in a failover cluster. The application servers in the farm were used to host the search query and crawler roles. The farm servers at all tiers used teamed network connections to provide load balancing and failover capabilities. The following table lists the operating system and software editions used in the farm configurations.

Table 11. Operating systems and software editions

	WFE and application servers	Database servers
Operating system	Windows Server 2008 R2 SP1 Enterprise Edition	
SharePoint server	SharePoint Server 2010 Standard Edition SP1	NA
Database server	NA	SQL Server 2008 R2 x64 Standard Edition SP1

**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.

For the purpose of this performance study, ten SharePoint 2010 web applications were used with different site collection templates. This included several site collections with both collaboration (team site) and publishing templates. Also, each of the web applications was created with an initial content DB size of 250GB.

Figure 9. SharePoint Farm Logical Architecture

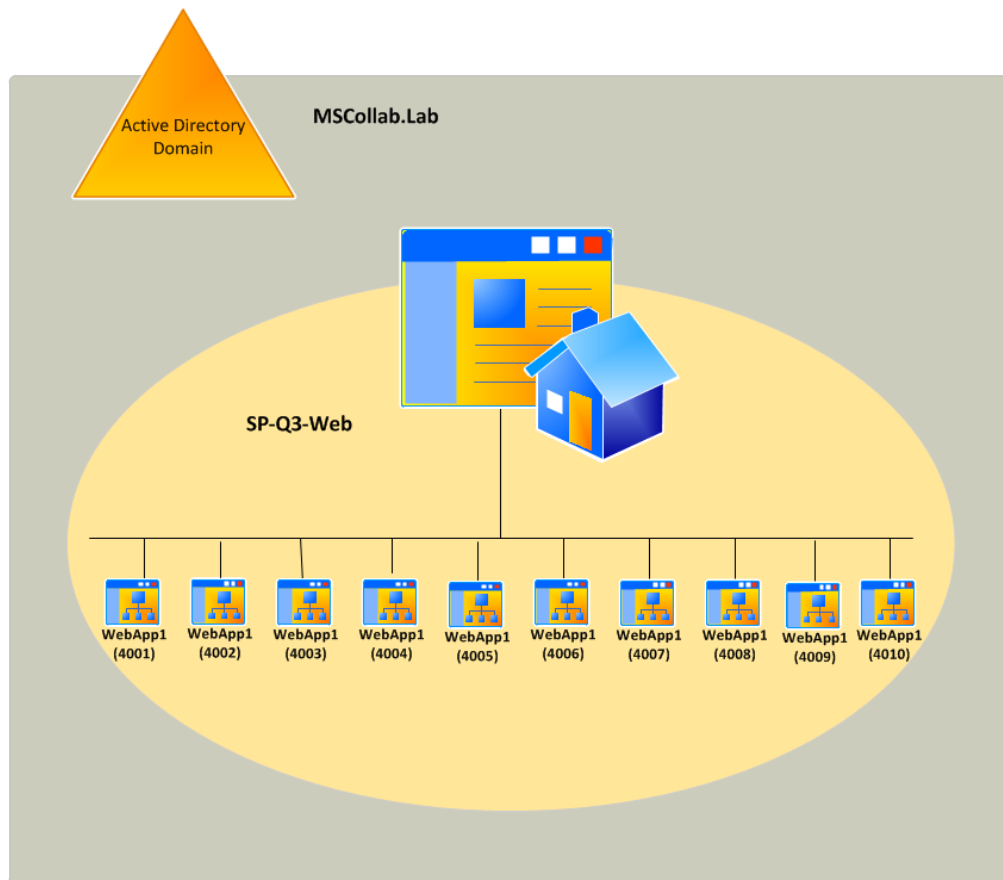


Figure 9 illustrates the logical architecture of the farm configuration. In Figure 9, the SharePoint web applications are represented as WebApp1, WebApp2, and so on. The port numbers - 4001, 4002, and so on - identify the SharePoint web application port numbers used in this performance study.

### Configuration of Web Front End (WFE) Servers

This SharePoint 2010 farm design included two WFE servers. The software matrix for these WFE servers is shown in Table 11. This performance study used SharePoint 2010 Standard Edition SP1 and included only out-of-the-box features used for collaboration workloads. As a part of the SharePoint mixed (collaboration and publishing) workload, only the standard search service application was deployed and no other service applications, such as Excel and Visio, were deployed.

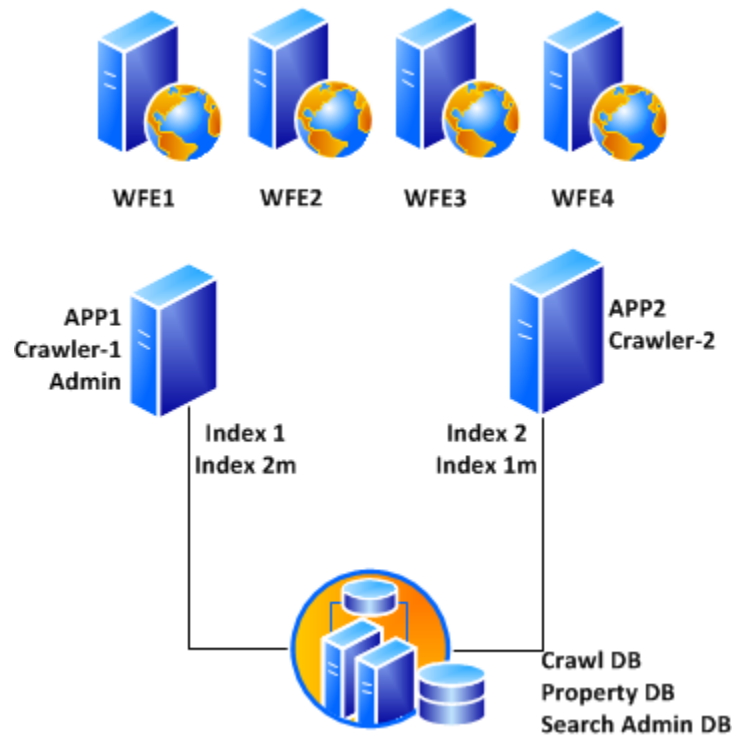
As mentioned earlier, BIG-IP LTM virtual edition load balancers were used in the reference architecture to balance the load on the WFE servers. In this reference architecture, the application server virtual machines were used for the search crawler and query roles only. However, these APP virtual machines can be added to the BIG-IP LTM virtual server pool when additional capacity at the WFE tier is desired.

### Search service application configuration

SharePoint 2010 changed the search architecture and introduced high availability at the application tier; for instance, the search crawler. The new search service application architecture in SharePoint 2010 includes greater redundancy. This 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 and query roles were used in the SharePoint 2010 farm.

The farm configuration used the search service application configuration shown in Figure 10. The search crawler and query roles were hosted only on the application servers for improved search performance. Both the crawlers were associated with the same crawl database, and indexed the same content source.

Figure 10. High-Level Search Service Architecture



The search architecture shown in Figure 10 is a logical representation of how the SharePoint Search service was configured. This configuration provides complete redundancy for both search crawler and query roles.

In Figure 10, the index partitions are represented as index 1 and index 2. The two primary index partitions were mirrored, for redundancy, and mirror, which is a copy and is referred to as index1m and index2m in Figure 10, of the index partition was placed on each application server. On both of the application servers, VMware raw-device mapping was used to connect the volumes created on Compellent storage arrays. Four volumes - two per application server - were used to store each index partition.

### HTTP Request Throttling

SharePoint 2010 offers resource throttling features that are 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, Request in Queue, and Request Wait Time are monitored. After three unsuccessful checks, the server enters a throttling period and remains in this state until a successful check is completed. Requests that were generated prior to the server's entering throttling mode are completed. Any new

HTTP GET and Search Robot requests generates a 503<sup>10</sup> error message and is logged in the event viewer.

The throttle settings can be modified to increase the overall load supported by the farm servers. However, this itself requires a complete study to find accurate throttle setting recommendations for any given user load or requests per second. The default HTTP throttle monitor settings 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.

## Configuration of the Database Server

As shown in Figure 10, the SharePoint 2010 farm deployed two Dell PowerEdge M820 physical servers as the database servers. Two database servers were deployed in a highly available SQL failover 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 storage backend. In this reference architecture, the storage resources on the database server were limited, so the SQL instance was connected directly to the external Dell Compellent FC storage arrays.

These arrays were configured with a single tier and used a RAID10 and RAID5-9 storage profile with no replay configuration.

The following table shows how the storage was used to host the SharePoint content and other databases.

Table 12. SharePoint database LUN details

Database	Number of LUNs	LUN Size	Total Size
10 x SharePoint Content Databases	10	250GB	2.5TB
10 x SharePoint Content Logs	10	100GB	1TB
Temp DB	1	100GB	200GB
Search DB (Crawl, Property, and Admin)	1	200GB	200GB
WSS Usage DB	1	200GB	200GB
Other SharePoint Databases (Config and AdminContent)	1	100GB	100GB

---

<sup>10</sup> Throttling starts alert- Events 8032 8062 - <http://technet.microsoft.com/en-us/library/ee513044.aspx>



The disk volumes described in Table 12 were configured as NTFS mount points on clustered SQL servers. Additional configuration is required when using NTFS mount points with the SQL failover cluster to ensure the disk dependencies failover with the SQL service. For example, the disk dependencies must be added in the SQL Server service dependencies in the Failover Cluster MMC.

Figure 11. SQL Service Disk Dependency Configuration

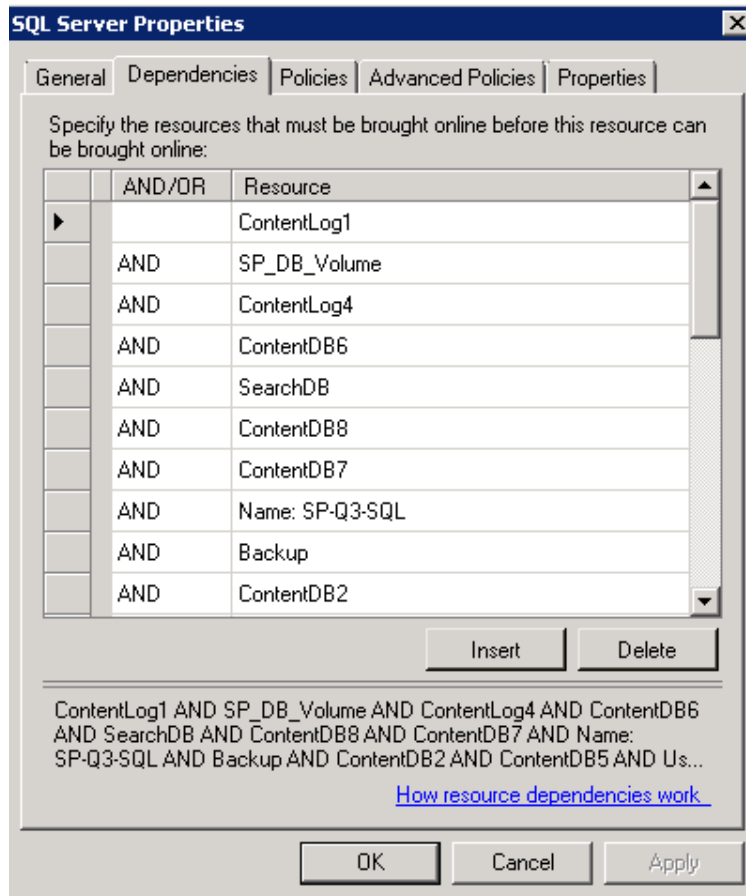


Figure 11 only shows a part of the configuration of cluster disk dependencies used in this performance study. All the shared disk volumes must be added to the SQL server services.

**NOTE:** Step-by-step instructions for configuring a cluster disk dependencies are outside the scope of this paper. For more information, see <http://support.microsoft.com/kb/280297>

### Memory configuration

By default, SQL Server service uses all available memory<sup>11</sup> as 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 physical memory

<sup>11</sup> SQL Server memory options - <http://msdn.microsoft.com/en-us/library/ms178067.aspx>

available. For example, out of 96GB of memory on the database hosts, 78GB was allocated to SQL Server.

## Performance Study of a SharePoint Farm

Microsoft SharePoint 2010 is a versatile platform that is used in a large 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 a multitude of ways of using SharePoint, which makes it almost impossible to accurately size servers and storage for a SharePoint farm. In addition, there is no standard benchmark for sizing SharePoint workloads at this time. It is very important to provide proper guidance to customers when it comes to recommending infrastructure elements of a SharePoint implementation. This led to the development of the Dell SharePoint Load Generation framework used to perform load testing of a SharePoint farm.

### Dell SharePoint Load Generation Framework

An internally developed load generation framework was used to understand the performance characteristics of the 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 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 12. LoadGen Content Population

The screenshot shows a web-based configuration interface for the LoadGen Content Population tool. It features several input fields and a list of application tasks. The 'WebURL' field is set to 'http://sp2010cluster:111'. The 'Select Usage Profile' dropdown is set to 'AverageCollaboration'. A 'View XML' button is located to the right of the dropdown. Below this, there is a 'Select Application' section with a 'Select All' checkbox checked. A list of application tasks is displayed, all of which are checked: 'CreateDefaultItems', 'AddDefaultWebParts', 'UploadFilesToSharePoint', 'UploadImagesToSharepoint', and 'ExportSharePointdata'. A 'Load SharePoint Data' button is located at the bottom of the interface.

The content population tool was developed to:


- Create SharePoint web applications
- Create site collections
- Add web parts to home pages
- Create document libraries
- Create SharePoint list items
- Upload documents/images and so on.

This utility is capable of populating hundreds of gigabytes of SharePoint content in few hours. The size of SharePoint content Database and other aspects such as number of site collections, and so on, vary based on the usage profile selection. A usage profile is a collection of use cases closely mapped to real world SharePoint usage. To some extent, these usage profiles were mapped in to SharePoint Capacity Planner<sup>12</sup> and other Microsoft recommendations. Although SharePoint capacity planner was intended for MOSS 2007, there are several aspects of these recommendations<sup>13</sup> 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.

## VSTS Load Testing Framework

Dell's SharePoint load generation framework uses VS 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 Visual Studio 2010 helps in rapid creation of use cases and parameterize those use cases. SharePoint load testing is performed using a test rig - shown in Figure 13 - of several test agents and the results are captured in to a SQL database on the test controller. Figure 13 represents only a portion of the actual farm and test rig. The actual test rig used for the study included 43 test agents and the farm as depicted in Figure 1.

Figure 13. Load Test Deployment



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

<sup>13</sup> Microsoft SharePoint 2010 performance and capacity management - <http://technet.microsoft.com/en-us/library/cc262971.aspx>

## Load Testing Workload Test Mix

As mentioned earlier, the load test usage profiles were based on the SharePoint Capacity Planner and other Microsoft recommendations<sup>14</sup> for SharePoint 2010. System Center SharePoint capacity planner (SCP) defines several usage profiles for both collaboration and publishing workloads.

Within the scope of this performance study paper, collaboration and publishing usage profiles were used. The following tables - Table 13 and Table 14 - provide an overview of the test profiles used in this load test.

---

<sup>14</sup> SharePoint Server 2010 Capacity Management - <http://technet.microsoft.com/en-us/library/cc261700.aspx>

Table 13. Dell's Collaboration test mix

Collaboration Test Mix	Number of tests/hr/user
Read Site Home Page	6
Read Survey	2
Read Lists	2
Read Document Library	1
Read Wiki Page	1
Read Picture Library	1
Create Wiki Page	1
Upload Document	1
Search Site	3
Respond to Survey	1
Edit Wiki Page	1
Total tests/hour/connected user	20

Table 14. Dell's Publishing test mix

Publishing Test Mix	Number of tests/hr/user
Read Site Home Page	9
Read Picture Library	1
Read Survey	2
Search Site	3
Read Document Library	1
File Upload	1

Publishing Test Mix	Number of tests/hr/user
Create Wiki Page	2
Edit Wiki Page	1
<b>Total tests/hour/connected user</b>	<b>20</b>

Dell’s test mix, shown in Table 13 and Table 14 is not a one-to-one mapping in to the System Center Capacity Planner (SCP) for MOSS 2007 and MS recommendations. For example, SCP defines total requests per hour per connected user. However, within Dell’s test mix for these usage profiles, this translates in to more requests than 20 per hour as the usage profile deploys 20 tests per hour for each connected user and one test could mean more than one request. Hence, the results published in this paper may or may not map directly in to SharePoint capacity planner recommendations directly and are specific to the workload mix defined in Table 13 and Table 14.

## Test Methodology

The intent of this performance study was to understand the capacity of a SharePoint farm as shown in Figure 1 with the configuration described in Table 3 and Table 4. Several load test iterations were conducted with incrementing user load. For example, an initial user load of 500 virtual users was used and the same had been incremented by 500 users until the farm resources reached an optimal level of usage. The overall goal of the load test was to make sure that the processor usage is below 60 percent and the average farm response time is sub 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. The following table shows a distribution of file content sizes in each Web application used in this performance study.

Table 15. Data set used in the performance study

Average File size	Number of files
1KB to 10KB	224122
10KB to 100KB	47235
100KB to 1MB	138262
1MB to 16MB	31517
16MB to 128MB	617
Greater than 128MB	12

The aggregated SharePoint content database size was around 2TB. During the load test duration, this content DB 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 13 and Table 14. A full content crawl was performed once at the beginning of the load tests. There were no subsequent crawls after load test 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 3 and Table 4. The following sections of this paper described the performance data and how several components within the farm performed at incremental user loads.

## Performance Results and Analysis

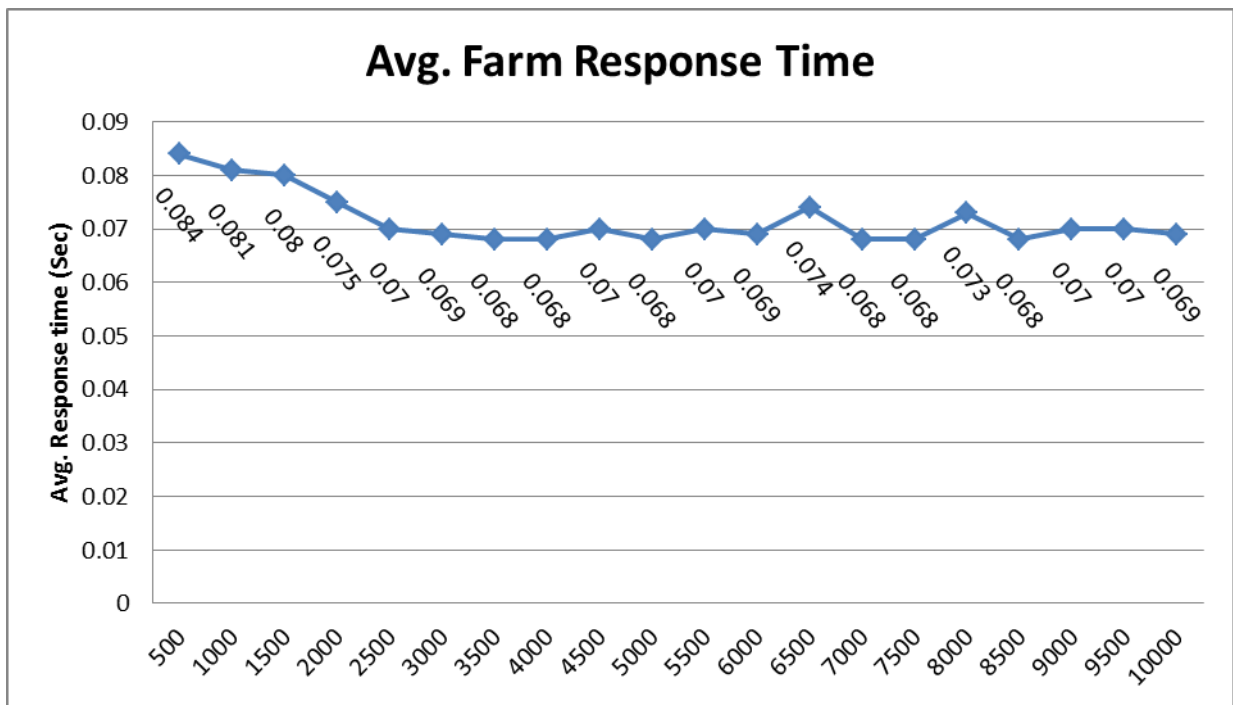
As a part of this performance study, several performance metrics were collected and analyzed. This section describes how various components of the SharePoint farm performed under incremental user load. As mentioned earlier, this study included a mixed (collaboration and publishing) SharePoint workload. So, all the results shown here are relative to the workload used and may differ with any other implementation outside of the test mix shown in Table 13 and Table 14.

Table 16. Validated farm capacity

	SharePoint Large Farm
Maximum concurrent <sup>15</sup> user load validated	10000
Requests per second <sup>16</sup> at Max concurrent user load	325/sec
Avg. Response time at the maximum validated user load	0.069 seconds

The above metrics indicate that the farm configuration used for this performance study could support faster, sub one-second farm response times even at the maximum validated concurrent user load. The following charts show average farm response time for various user load iterations. As seen below, it is clear that the average farm response time had always been less than a second and almost stayed flat during the load tests duration.

Figure 14. Average Farm Response Time



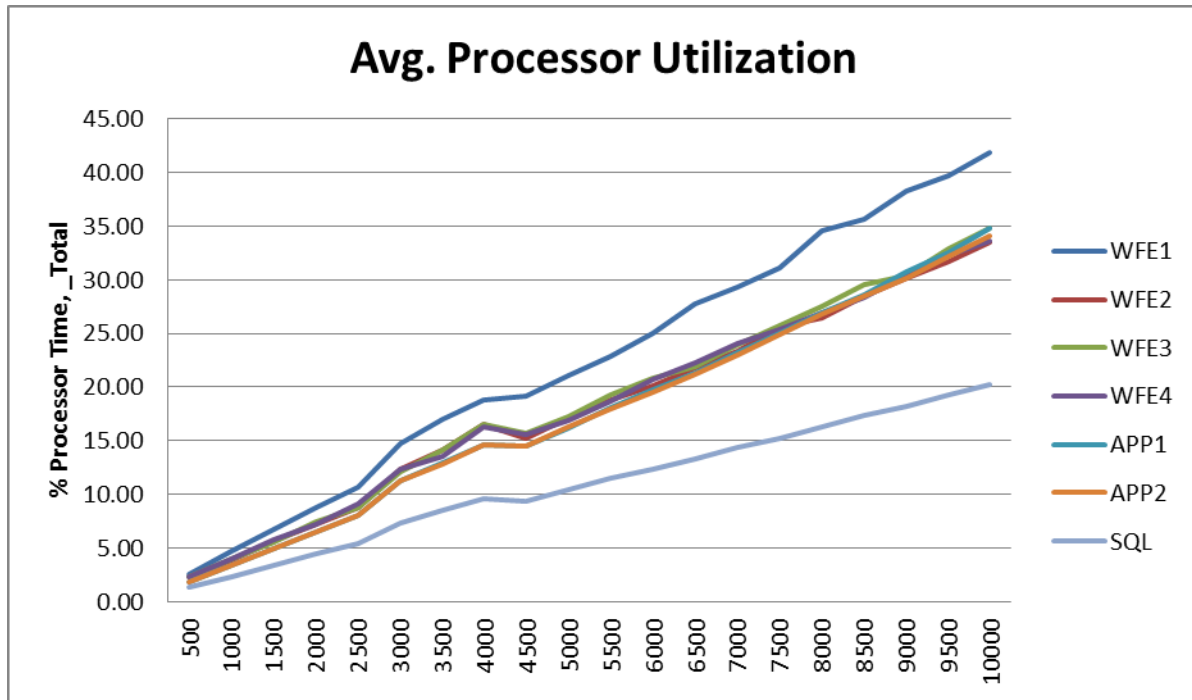
<sup>15</sup> Concurrency refers to number of simultaneous requests to the farm servers

<sup>16</sup> This number indicates the avg. requests per second generate during the load test duration and this is a Visual Studio reported metric



The average processor utilization on the web frontend servers in the farm configuration was minimal even at the maximum validated user load. The following charts show the processor utilization metrics for several user load iterations in the SharePoint farm.

Figure 15. Processor Utilization on Farm Servers



From the above processor utilization metrics, it may seem that fewer than four frontends may be sufficient to support the desired maximum concurrent user load of 10000 concurrent users in a mixed usage profile shown in Table 13 and Table 14.

However, in a heavily loaded scenario, even though the average processor usage is below 45-50 percent, the SharePoint farm configuration may not support a larger number of users than what is shown in Table 16. This 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 is outside the scope of this paper and may require an in-depth study in itself.

The overall network usage on Web frontend and Application servers was very minimal and in the range of 280 Mbps. Considering the available throughput, the network usage on WFEs and APP servers was not a significant load at all. The Database Server network utilization was also found to be very minimal.

As shown in Figure 1, a Dell Compellent Storage Center Series 40 array was used for the SharePoint storage backend. As shown in Figure 9, this performance study used 10 SharePoint Web applications each with a separate content database. The total SharePoint content size was approximately 2 Tera Bytes(TB). Also, all the other SharePoint databases such as Search, Usage Data, and SQL TempDB were also stored on the Dell Compellent Series 40 Storage Center arrays. The following table shows the I/O read-write statistics and overall IOPS achieved at the maximum validated concurrent user load.

Table 17. Storage performance metrics

IO Performance metric	SharePoint Large Farm
Avg. Disk Transfers/Second	~1454
Avg. Disk Writes/second	~1250
Avg. Disk Reads/Second	~204
Avg. Disk seconds/Transfer	0.002 or 2ms
Avg. Disk Queue Length (_Total)	2.17
Avg. Disk Bytes/Transfer	110863 (108KB)

The above storage metrics indicate that the workload had an I/O mix of ~12 percent reads and ~88 percent writes with an approximate I/O size of 108KB. The Dell Compellent Series 40 Storage Center with 2x 24 disk 15K SAS backend is - theoretically - capable of providing approximately 7500 IOPS. In this reference architecture, the IOPS generated - approximately 1454 IOPS - at the maximum validated user load is well below the theoretical maximum of the underlying storage. This showed that the Dell Compellent Series 40 Storage Center array was capable of handling the mixed SharePoint workload of up to 10000 concurrent users validated in this performance study.

## Summary

A SharePoint Server 2010 farm consists of multiple servers, each of which is provisioned with different SharePoint components. A large SharePoint farm in general is a best choice for large enterprises with relatively high concurrent user load. These farms employ a three-tier architecture. 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 Server 2010 can be used in many different ways and each implementation necessitates an in-depth study of requirements, such as expected user load, requests per second and future growth. The Dell SharePoint Load Generation Framework helps in generating SharePoint farm performance data during load test. These data can be used to provide recommendations for best performance. This performance study paper was intended to demonstrate the performance capacity of a partially virtualized SharePoint 2010 farm built using Dell servers, Dell Compellent storage, and vSphere 5.0 for the WFE and APP roles.

This study showed that the configuration, as illustrated above, could support more than 100,000 users with a minimum concurrency of 10 percent and an average farm response time was well below one second. The Dell Compellent Series 40 Storage Center array provided highly optimal performance for the SharePoint 2010 deployment used in this performance study. The Dell PowerEdge M1000e blade chassis provided an energy efficient and scalable infrastructure for this SharePoint farm implementation. The blade IO options - Dell PowerConnect M8024-K - along with the top-of-rack Dell Force10 S4810 network switches provided an end to end non-blocking and non-oversubscribed 10GbE

Implementing Microsoft SharePoint Server 2010 on Dell 12<sup>th</sup>-Generation Blade Servers and VMware vSphere 5.0

network bandwidth for all the SharePoint 2010 farm roles. Overall, the reference architecture and the infrastructure components of this performance study were found capable of achieving a 10000 concurrent user load in a mixed SharePoint 2010 usage profile environment.

## References

Dell solutions for SharePoint: [Http://dell.com/sharepoint](http://dell.com/sharepoint)

Dell Data Center Systems Management Solutions: <http://content.dell.com/us/en/enterprise/data-center-system-management>

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

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

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>