

Microsoft® SharePoint® Server 2010

Large Farm with Compellent Storage - Performance Study

Ravikanth Chaganti and Quocdat Nguyen

Dell™ | SharePoint Solutions

January 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, PowerConnect, Compellent, and PowerEdge are trademarks of Dell Inc. Big-IP is the trademark of F5 Networks. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, SharePoint, SQL Server, Visual Studio, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

January 2012 | Rev 1.0

Contents

Executive Summary	5
Introduction	6
SharePoint 2010 Farm Topologies	6
Large Server Farm Topology	7
SharePoint Farm Configuration	11
Farm Architecture and Configuration of Farm Roles	12
Configuration of Web Front-End and Application Servers	14
Hardware Load Balancers	14
HTTP Request Throttling	15
Search Service Application Configuration	15
Network Configuration	16
Configuration of Database Servers	16
SQL Server Memory Configurations	17
Performance Study of a Large Farm	17
Dell SharePoint Load Generation Framework	17
Content Population Tool	18
Visual Studio Load Testing Framework	18
Load Testing Workload Test Mix	19
Test Methodology	20
Performance Results and Analysis	21
Summary	25
References	25

Tables

Table 1. Hi-Level Overview of Farm Configuration	7
Table 2. Dell Compellent Hardware Capabilities	9
Table 3. Hi-Level Farm Server Hardware Configuration	12
Table 4. OS and SharePoint Software Editions Overview	13
Table 5. Content Database and LUN Size Used in This Study	17
Table 6. SCP Collaboration Profile Mix	19
Table 7. Dell Light Collaboration Test Mix	20
Table 8. Table 8 Dataset Categorization	21
Table 9. High Level Overview of the Farm Performance	21
Table 10. Storage Performance Metrics	24

Table 11.	Detailed Storage Performance Metrics per Content Database	24
-----------	---	----

Figures

Figure 1.	Farm Physical Architecture and Cabling Details	12
Figure 2.	Farm Logical Architecture	13
Figure 3.	BIG IP Load Balancer Cabling Details	15
Figure 4.	Search Architecture.....	16
Figure 5.	LoadGen Tool Framework	18
Figure 6.	VSTS Test Rig.....	19
Figure 7.	Farm Response Time.....	22
Figure 8.	Processor Utilization of Farm Servers	22
Figure 9.	Database Server Network Utilization	23

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 workload, 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 large SharePoint farm, built using Dell PowerEdge™ blade servers and Dell Compellent™ Storage Center storage arrays with Fibre Channel Controllers Series 40, performed under load testing. The key findings from this study are:

- The recommended farm architectures were 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 87ms or 0.087 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 with an average disk response time of 2ms and ~2400 average disk transfers per second.
- Using hardware load balancing solution enabled us to scale the farm beyond 4 Web Front-Ends (WFE) and achieve the maximum desired concurrent user load.

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.

A companion paper, *SharePoint Server 2010: An Introduction*, is available from www.dell.com/sharepoint. It offers an overview of SharePoint Server 2010, and provides common concepts and definitions that form a basis for understanding the reference architectures presented in this paper.

Another companion paper, *SharePoint 2010: Designing and Implementing a Large Farm*, is available from www.dell.com/SharePoint. This companion paper provides the reference architecture and infrastructure best practices for implementing a SharePoint 2010 large farm. These reference architectures formed the basis of the performance study described in this paper.

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 and future growth. A SharePoint implementation can be used in several ways, including custom-developed applications. This brings in the complexity factor while sizing the servers and storage for a SharePoint implementation. The six pillars¹ of SharePoint can be used to drive clarity around how SharePoint is used. This performance study paper provides performance capacity details of a SharePoint 2010 large farm configured with Dell PowerEdge blade servers and Dell Compellent Storage Center Series 40 storage arrays in the context of SharePoint collaboration² workload.

SharePoint 2010 Farm Topologies

A SharePoint server farm is a set of servers, which collectively provides the services needed 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. These optional components may include some of the service applications such as managed metadata service, Excel services, and so on. Some constraints and best practices determine which components should be located on each server in the farm. Also, 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 varies based on several factors such as number of concurrent users, service application in the farm, the expected uptime Service Level Agreement (SLA), and so on. These factors dictate how many servers are needed in the SharePoint farm and how the overall farm architecture looks. Based on these factors, SharePoint 2010 farm implementations are classified into small farm³, medium⁴ farm and a large farm deployments.

¹ SharePoint capabilities - <http://sharepoint.microsoft.com/en-us/product/capabilities/Pages/default.aspx>

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

³ SharePoint 2010 - Designing and implementing a small farm
http://www.dell.com/downloads/global/services/dell_small_sharepoint_farm.pdf

⁴ SharePoint 2010 - Designing and implementing a medium farm
http://www.dell.com/downloads/global/services/dell_medium_sharepoint_farm.pdf

Large Server Farm Topology

A typical SharePoint large⁵ server farm consists of three tiers:

- Web front-end
- Application
- Database.

Dedicated servers are 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 large 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. This performance study paper used SharePoint 2010 large farm architecture to understand how several components of a farm perform at incrementing user loads.

Within the scope of this paper, single farm was used to study the performance characteristics of SharePoint 2010 on Dell servers and storage. Figures 1 depicts the reference architecture of the farm used in this performance study.

Table 1. Hi-Level Overview of Farm Configuration

	Large Farm Configuration
Blade Chassis	One (1) Dell PowerEdge M1000e with Ethernet Pass-through modules
Web Front-end Servers	Six Dell PowerEdge M710HD Servers
Application Servers	Two Dell PowerEdge M710 Servers
Database Servers	Two Dell PowerEdge M910 Servers
Storage Arrays	A Dual-Controller Dell Compellent Series 40 storage array with two 2U - 24 drive bay - SAS enclosures

The following section describes the servers chosen for each of the farm roles and provides a technical overview of the servers used in this performance study.

Dell PowerEdge M1000e Blade Enclosures

The PowerEdge M1000e modular blade enclosure is the foundation for Dell's blade server architecture, providing one of the most energy-efficient, extremely reliable, 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

⁵ SharePoint 2010 - Designing and implementing a Large farm
http://www.dell.com/downloads/global/services/dell_large_sharepoint_farm.pdf

architecture. The M1000e is optimized for use with all Dell PowerEdge Blades including the M710, M710HD and M910 blades servers. Features include:

- **Energy Efficiency**
M1000e is built on Dell's energy smart technology, which can help you to increase capacity and lower operating costs while delivering better performance/watt.
- **Effortless Scalability**
With scale on-demand switch design and additional I/O slots and switch options, the M1000e provides a flexibility to meet the increasing demand for I/O consumption. Plus, Dell's FlexIO modular switch technology offers a 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 for IT administrators to manage and monitor multiple enclosures and blades from a single console.
- **Multi-Chassis Management**
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.

Dell PowerEdge M910

The PowerEdge M910 is a four-socket, full-height blade server with support of up to 512GB of physical RAM (32 x 16GB DDR3 DIMMs) and the latest six-, eight-, and 10-core Intel Xeon processors. This server supports a maximum of two internal SAS disk drives, and hence the maximum internal storage capacity is 1.8TB when using 2 x 900GB, 10K RPM SAS drives in a RAID 0 configuration. Similar to the PowerEdge M710, the M910 also supports four 1GB network ports without any additional I/O expansion cards.

Within the scope of this performance study paper, the PowerEdge M910 server was used at the database tier of the SharePoint farm. The enormous processing power and physical memory capacity makes this server the best choice for a database server.

In this study, two LOMs were used in a network team to connect the database server to the farm network. For connecting the database server to the fiber channel storage network, two QLogic QME2472s were used and configured in MPIO.

Dell PowerEdge M710HD

The PowerEdge M710HD is a two-socket, half-height blade server with support of up to 192GB of physical RAM (32 x 16GB DDR3 DIMMs) and the latest quad- and six-core Intel Xeon processors. This server supports a maximum of two internal SAS disk drives and hence the maximum internal storage capacity is 1.2TB when using 2 x 600GB, 10K RPM SAS drives in a RAID 0 configuration. This server provides four 1GB network ports⁶ without any additional I/O expansion cards.

⁶ When using an Ethernet Switch module and not a pass-through module.

Within the scope of this performance study, M710HD has been used at the web front-end (WFE) is the tier of the farm configuration.

Dell PowerEdge M710

The PowerEdge M7107 is a two-socket, full-height blade server with support for up to 288GB of physical RAM and the latest quad-core and six-core Intel Xeon processors. The M710 supports a maximum internal disk storage capacity of 3.6TB when using 4 x 900GB and 10K RPM SAS drives in a RAID 0 configuration. This server supports up to four 1GB network connections without using additional expansion cards.

Within the scope of this performance study paper, the PowerEdge M710 servers were used at the application tier of the farm. These additional HDDs were used to contain the SharePoint farm's index queries.

In this study, two 1GB network ports - on both M710 and M710HD - in a load balancing team were used for farm connectivity. A load balancing network team was used to connect the web front-end and application servers to the farm network used in this performance study.

Dell Compellent Storage Center Storage Arrays

Dell Compellent 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 combining SSD, FC and SAS drive technologies as well as iSCSI, FCoE and FC front-end 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

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
Storage Capacity	739TB of raw storage capacity with FC drives or 1,920TB with SAS drives per system.
Hard Drives	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.

3.5" Drive Performance and Capacities	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 and 2T
2.5" Drive Performance and Capacities	15,000 RPM SAS drives available in 146GB 10,000 RPM SAS drives available in 450GB and 600GB 7,200 RPM SAS drives available in 1TB SAS SSD available in 200GB
Expansion Capabilities	Each system supports up to 1,232 FC drives or 960 SAS drives
Host Connectivity	Maximum of 22 front-end ports per controller. A switch is recommended for all configurations and required for clustered controller configurations.
RAID Levels	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 array without RAID group limitation.

In this performance study, we used two 2U - 24 drive bay SAS enclosures containing 28 spindles to host the content database for the large SharePoint farm. These are 6G SAS EBOB enclosures which can house up to 48 - 2.5" 6Gb SAS drives and has an aggregated SAS connection speed of 24 Gb (6Gb speed per channel, in 4 channel SAS cable). One 2U SAS EBOB enclosure contains the 24 spindles of a 15K 146GB SAS drive; the second SAS EBOB enclosure holds the other four spindles. Those available drive bays provide room for future need as of storage demand. In addition there is almost 3TB of raw storage created as of the tier-1 storage (RAID10 and RAID5-9 storage profile) to host the studied large SharePoint farm configuration using the 26 active spindles and 2 hot spare spindles. As mentioned above we included only tier-1 storage for the performance study. Users can add in near-line SAS-drive enclosures as of a tier-3 storage to leverage Compellent's Storage Center feature such as data progression, as well as to improve the total cost of ownership.

Compellent Dynamic Capacity™ (Thin Provisioning) was used to create volumes for the SharePoint content database. Dell Compellent Thin Provisioning delivers the highest enterprise storage utilization possible by eliminating pre-allocated but unused capacity. The software, Dynamic Capacity, completely separates allocation from utilization, enabling users to provision any size volume upfront yet only consume disk space when data is written. Thin Write technology assesses the incoming payload and designates capacity for each write on demand, leaving unused disk space in the storage pool for other servers and applications. Users can now plan for growth up front without the need to worry about unused capacity. Not only is capacity not used until the data is written to the volume, users also have the ability to expand volumes on the fly without downtime.

The other 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 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, hence reducing the total cost of ownership.

As part of a Fibre channel (FC) Storage Area Networking (SAN) configuration shown in Figure 1, we used Brocade 5300 switches in a full redundant configuration. The Brocade 5300 is a high-performance 8 Gbit/second Fibre Channel switch designed for the needs of enterprise environment. In addition, the Brocade 5300 switch is an 80-port auto-sensing 1, 2, 4, and 8 Gbit/sec switch which is suitable to interface with both back-end Compellent storage array and the front-end SharePoint farm servers. Dell Compellent storage Series-40 controllers are equipped with 8 Gbit/sec FC Host Bus Adapters (HBA) as an interface to the SAN configuration. However, database servers are supplied with 4 Gbit/sec QLogic QME2472 HBAs because these would provide sufficient bandwidth for the database server. Therefore, Brocade 5300 is an appropriate choice not only to provide back-ward compatibility but also to support future growth for the SAN environment.

Figure 1 shows the cabling details along with the complete farm physical architecture.

SharePoint Farm Configuration

The experimental farm configuration in this study used one M1000e blade chassis and farm servers were spread across both chassis to provide blade chassis-level redundancy. The configuration accommodates the entire SharePoint farm while leaving enough room for other workloads or future farm growth. As mentioned, this configuration, as shown in Figure 1, used Dell PowerEdge M710HD servers at the Web front-end tier, Dell PowerEdge M710 servers at the application tier, and Dell PowerEdge M910 servers at the DB tier.

Figure 1. Farm Physical Architecture and Cabling Details

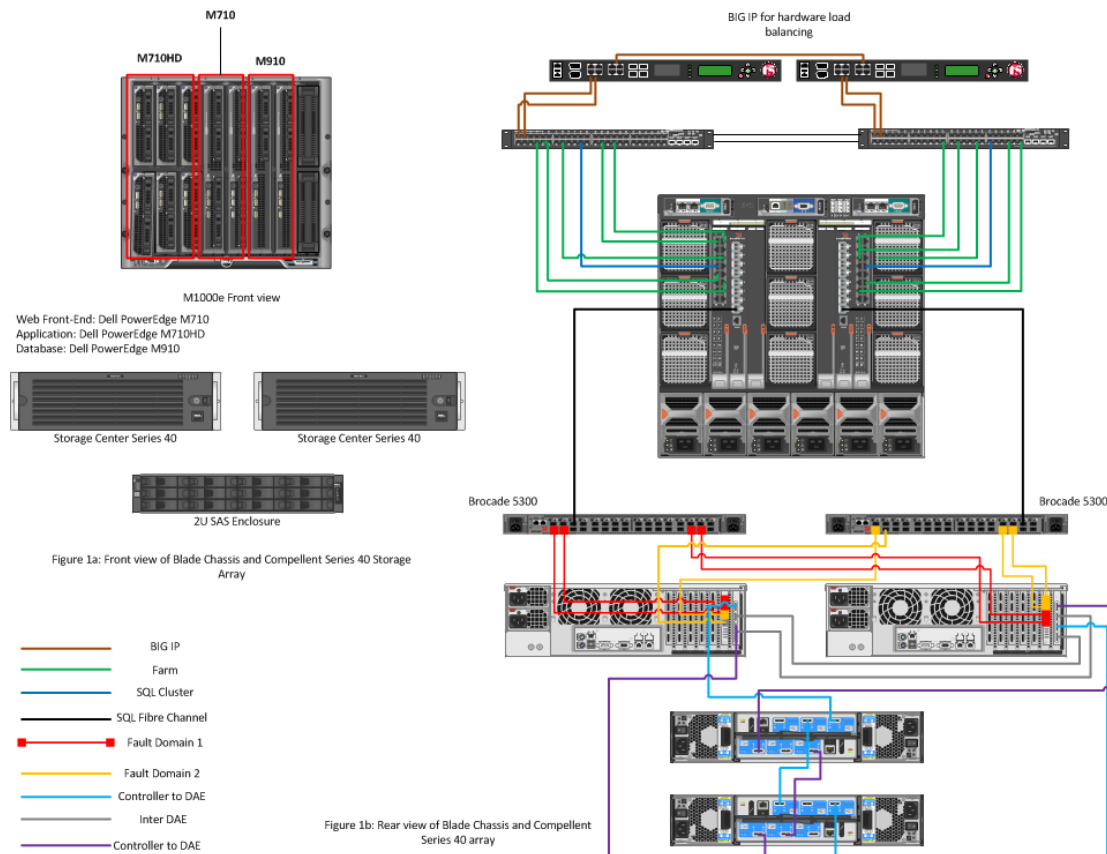


Table 3. Hi-Level Farm Server Hardware Configuration

Server Role	Web Front-end	Application Server	Database Server
Server Model	M710HD	M710	M910
Processor	Two Sockets - X5670, 6 cores, 2.93Ghz	Two Sockets - X5550, 4 cores, 2.66Ghz	Four Sockets - E7540, 6 cores, 2.0Ghz
Memory	24GB	24GB	96GB
Internal Storage	146GB RAID 1 for OS	146GB-RAID1 for OS and 146GB RAID for Index Query	146GB RAID 1 for OS
Network Controller	2 * 1Gb BCM5709S NIC teaming for Farm connections	2 * 1Gb BCM5709S NIC teaming for farm connections	4 * 1Gb BCM5709S NIC teaming for farm connections and one NIC for cluster private network.
Fiber Channel HBA	-	-	2 * Qlogic QME2472 FC4 HBA
External Storage (Compellent)	-	-	28 * 146GB, 15K, 6Gbps SAS Drives


Farm Architecture and Configuration of Farm Roles

The experimental farm configuration, as shown in Figure 1, included six WFEs, two application servers, and two database servers in a failover cluster. This farm was configured to use Microsoft Windows® authentication and hence all the requests during the load test were authenticated requests.

In general, any SharePoint farm with content database size more than or equal to 2 TBs is considered a large farm. However, pre SharePoint 2010 SP1, a limit of 200GB per content database was imposed in a general usage scenario, such as collaboration. This performance study used only SharePoint 2010 and not SharePoint 2010 SP1. Hence, the content database limitation was considered and 11 web applications were created, each with a content database size of maximum 250GB.

In this performance study, each Web application hosted four site collections under which several document libraries and other SharePoint list items were created. Figure 2 illustrates the logical architecture of the farm configuration.

Figure 2. Farm Logical Architecture



The following table shows the operating system and software used as a part of this performance study.

Table 4. OS and SharePoint Software Editions Overview

	Web Front-Ends	Application Servers	Database Servers
Operating System	Windows Server® 2008 R2 Enterprise Edition SP1	Windows Server 2008 R2 Enterprise Edition SP1	Windows Server 2008 R2 Enterprise Edition SP1
SharePoint Server	SharePoint 2010 Server Standard Edition	SharePoint 2010 Server Standard Edition	NA
Database Server	NA	NA	SQL Server® 2008 R2 x64 Enterprise Edition

Note

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

Configuration of Web Front-End and Application Servers

The SharePoint 2010 farm design included six Web front-end servers. The software matrix for these Web front-end servers is as shown in Table 4. SharePoint 2010 Standard Edition used in the performance study included only the out-of-the-box features of SharePoint and was a collaboration workload only. As a part of the collaboration workload, only the search service application was deployed. No other service applications, such as Excel services or Visio Services, were deployed.

On the PowerEdge M710 and M710HD, simultaneous multi-threading or logical processor support was enabled for increased performance. This option is disabled by default in the system BIOS and must be enabled manually.

Hardware Load Balancers

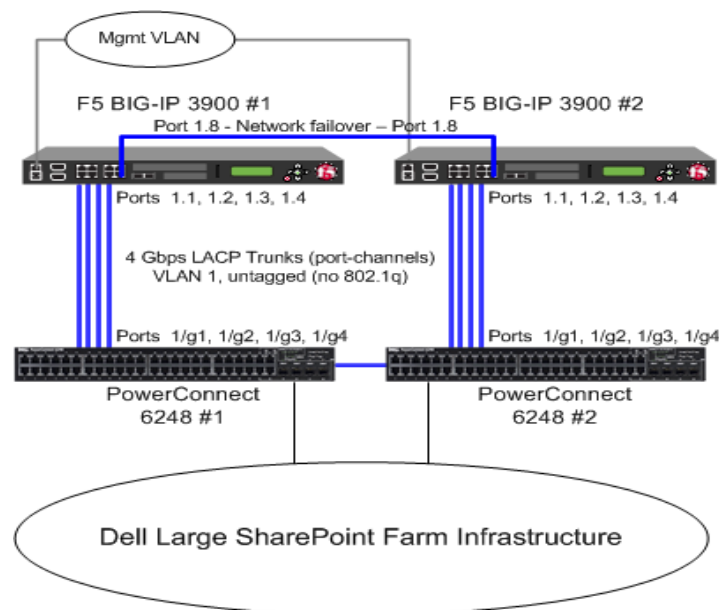
Within the scope of this study? Project?, the farm configuration used F5 Networks® BIG-IP® Local Traffic Manager™ (LTM) hardware load balancers to enable load balancing across Web front-end nodes. The native software network load balancing (NLB) clusters can become unstable when there are more than 4 or 5 Web front-ends. Hence, two F5 BIG-IP hardware load-balancing switches were used for Web front-end load balancing.

As a part of the farm architecture, two F5 BIG-IP 3900 series switches were used. These load balancer systems feature high-performance SSL acceleration hardware and software compressions⁷ as well as advanced connection management to remove processing intensive tasks from application servers. A BIG-IP 3900 switch system features a Quad core CPU, 8GB of memory, and supports up to 4Gbps of traffic throughput. Using hardware load balancers instead of software NLB enabled the farm capacity to go beyond four web front-ends and achieve higher concurrent user load without compromising the farm performance.

No custom load balancer profiles were defined for this study. The farm configuration used the out-of-box acceleration functionality and BIG-IP application templates, thereby simplifying the administrative tasks and shortening the required set up time. The following figure shows how the load balancers were connected to the SharePoint farm infrastructure.

⁷ Hardware datasheet: <http://www.f5.com/pdf/products/big-ip-platforms-ds.pdf>

Figure 3. BIG IP Load Balancer Cabling Details



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 503⁸ 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 come up with 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.

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. The index does, however, still propagate and is stored locally on the query servers. Two

⁸ Throttling starts alert- Events 8032 8062 - <http://technet.microsoft.com/en-us/library/ee513044.aspx>

application servers hosting the crawler role were used in this performance study. The query role was hosted on two application servers to provide better availability and improved search performance.

The farm configuration implemented the search service application, as illustrated in Figure 4.

Figure 4. Search Architecture



Note

In the above figure, ‘m’ in the index partition name represents a mirror. For example, index 1 represents index partition 1, and index 1m represents the mirror of index partition 1.

Network Configuration

On the PowerEdge M1000e blade chassis, Dell Ethernet pass-through modules were used for network connectivity. For both the Web front-end servers and applications servers, teamed network connections were used. These teamed connections (shown in Figure 1) were configured to be in the smart load balancing (SLB) mode, which supports both load balancing and failover.

Configuration of Database Servers

As shown in Figure 1, this performance study deployed PowerEdge M910 blade servers at the database server tier. Two database servers were deployed in a fail-over cluster to enable redundancy at the database tier of the SharePoint farm. A SharePoint farm’s performance depends largely on the performance of the database server and the database backend. The PowerEdge M910 blade servers are the best choice for hosting the SQL database. The PowerEdge M910 supports only 2 internal drives; therefore, the SQL instance had to be hosted on an external storage arrays. In both the farm configurations, Dell Compellent storage arrays were used to add performance and storage capacity. 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 volumes were provisioned in the farm to provide storage for the SharePoint content and other databases.

Table 5. Content Database and LUN Size Used in This Study

Database	Number of LUNs	LUN Size	Total Size
11 x SharePoint Content Databases	11	250GB	2.5TB
11 x SharePoint Content Logs	11	100GB	1TB
2 * Temporary databases	2	100GB	200GB
Search database (Crawl, Property, and Admin)	1	200GB	200GB
WSS Usage database	1	200GB	200GB
Other SharePoint Databases (Config and AdminContent)	1	100GB	100GB

SQL Server Memory Configurations

By default, SQL Server uses all available physical memory⁹. This is because SQL Server dynamically grows and shrinks the size of its buffer pool depending on the physical memory reported by the operating system. However, this behavior is 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 Dell PowerEdge M910 server used at the database tier, out of 96GB of physical memory, 77GB was allocated to SQL server.

Also, as a part of this study, Processor node interleaving feature in BIOS has been enabled to disable Non-Uniform Memory Access (NUMA). The **Node Interleaving** setting can be found under the **Memory Settings** section in the Dell PowerEdge system BIOS.

Performance Study of a Large 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 to provide guidance, developed 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.

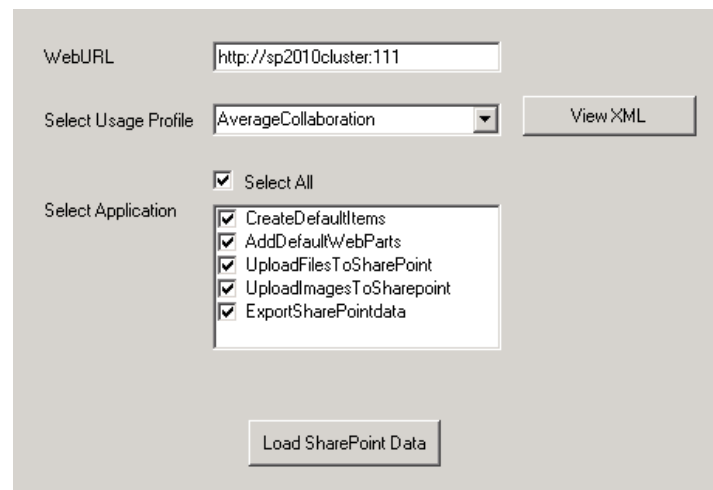
The Dell SharePoint load-generation framework has two components - a content population tool and Visual Studio Team Suite (VSTS) Web test framework.

⁹ SQL Server memory options - <http://msdn.microsoft.com/en-us/library/ms178067.aspx>

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 5. LoadGen Tool Framework



The screenshot shows the LoadGen Tool Framework interface. It includes a 'WebURL' text box with the value 'http://sp2010cluster:111'. Below it is a 'Select Usage Profile' dropdown menu set to 'AverageCollaboration', with a 'View XML' button to its right. Under the 'Select Application' section, there is a 'Select All' checkbox which is checked, and a list of application tasks: 'CreateDefaultItems', 'AddDefaultWebParts', 'UploadFilesToSharePoint', 'UploadImagesToSharepoint', and 'ExportSharePointdata', all of which are also checked. 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, 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 Planner¹⁰ and used other Microsoft recommendations. Although SharePoint capacity planner was intended for MOSS 2007, there are several aspects of these recommendations¹¹ 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 the Visual Studio® test framework.


Visual Studio Load Testing Framework

¹⁰ SharePoint capacity planner - <http://www.microsoft.com/downloads/details.aspx?FamilyID=dbee0227-d4f7-48f8-85f0-e71493b2fd87&displaylang=en>

¹¹ Microsoft SharePoint 2010 performance and capacity management - <http://technet.microsoft.com/en-us/library/cc262971.aspx>

Dell's SharePoint load generation framework uses Visual Studio 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 2008 helps to rapidly create and parameterize use cases. SharePoint load testing is performed using a test rig - shown in Figure 6 - of several physical test agents. The results are captured in to a SQL database on the test controller. Figure 6 represents only a portion of the actual farm and test rig. The actual test rig used for the study included 45 test agents and the farm as depicted in Figure 1.

Figure 6. VSTS Test Rig



Load Testing Workload Test Mix

As mentioned earlier, the load test usage profiles were based on the SharePoint Capacity Planner and other Microsoft recommendations for SharePoint 2010. System Center SharePoint capacity planner defines several usage profiles for both collaboration and publishing workloads. These usage profiles are categorized into 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, light collaboration usage profile was used. Table 6 shows the light collaboration test mix as suggested by SharePoint Capacity Planner (SCP).

Table 6. SCP Collaboration Profile Mix

SCP Usage Profiles	Light Collaboration
Home Page Access (%)	30
List Page Access (%)	20

Document/Picture Download (%)	15
Document/Picture Upload (%)	8
Search (%)	15
List Item insertion / deletion (%)	12
Total requests/hour/connected user	20

As shown in Table 6, SCP defines only a high-level test mix for each usage profile. Table 7 shows a more granular translation of this SCP light collaboration usage profile. Several use cases were mapped into each of the categories described by SCP, and the number of use cases per hour per connected user has been assigned.

Table 7. Dell Light Collaboration Test Mix

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

Dell's test mix, shown in Table 7, 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 light collaboration profile, this translates into more requests than 20 per hour as the usage profile uses 20 tests per hour for each connected user. In addition, 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 and are specific to the workload mix defined in Table 7.

Test Methodology

The intent of the experiments conducted as a part of this performance study was to understand the capacity of a large SharePoint farm as shown in Figure 1 with the configuration described in Table 3. Several load test iterations were conducted with incremental user load. For example, an initial user load of 500 virtual users was used and then 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 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 8 shows a distribution of file content sizes in each Web application used in this performance study.

Table 8. Table 1 Dataset Categorization

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 7. 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. 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. Based on the results, the farm configuration was modified to reach the final farm configuration shown in Figure 1. 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. 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 7.


Table 9. High Level Overview of the Farm Performance

	SharePoint Large Farm
--	-----------------------

Maximum concurrent ¹² user load supported	10000
Requests per second ¹³ at Max concurrent user load	426/sec

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 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 stayed almost flat during the load tests duration.

Figure 7. Farm Response Time




The average processor usage on the Web front-end servers in the farm configurations was minimal even at the maximum user load. The following figures show the processor usage metrics for several user load iterations in the SharePoint farm.

Figure 8. Processor Utilization of Farm Servers

¹² Concurrency refers to number of simultaneous requests to the farm servers

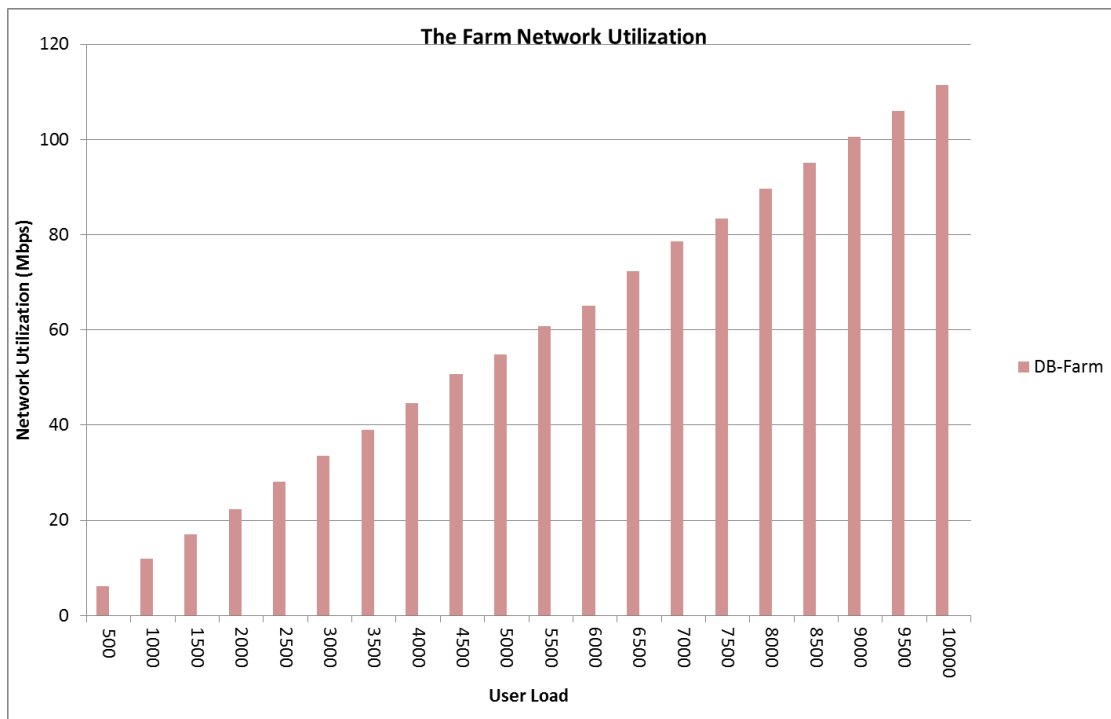
¹³ This number indicates the average requests per second generated during the load test duration. This is a Visual Studio reported metric.



From the above processor usage metrics, it may seem that fewer than six front-ends may be sufficient to support the desired maximum concurrent user load of 10000 users.

However, in a heavily loaded scenario, even though the average processor usage is below 10 percent, the SharePoint farm configuration may not support a larger number of users than what is shown in Table 9. 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.

Figure 9. Database Server Network Utilization



The overall network usage on web front-end and application servers was very minimal and in the range of 30 - 40 Mbps. Each of these web front-end servers have two embedded NICs (or LOMs) configured in a load balancing network team, resulting in 2Gbps overall available throughput. Considering the available throughput, the network usage on WFEs and Application servers was not a significant load. 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 Table 10, this performance study used 11 SharePoint Web applications, each with a separate content database. The total SharePoint content size was approximately 2TB. The other SharePoint databases, such as Search, Usage Data, and SQL TempDB, were also stored on the Compellent™ storage. The following table shows the I/O read-write statistics and overall IOPS achieved at the maximum supported concurrent user load.

Table 10. Storage Performance Metrics

IO Performance metric	SharePoint Large Farm
Avg. Disk Transfers/Second	~2,400
Avg. Disk Writes/second	1,736
Avg. Disk Reads/Second	623
Avg. Disk seconds/Transfer	0.0022 or 2ms
Avg. Disk Queue Length (_Total)	0.4
Avg. Disk Bytes/Transfer	55,341 (56KB)

The above storage metrics indicate that the workload had an I/O mix of ~30 percent reads and ~70 percent writes with an approximate I/O size of 56KB. The following tables provide the detailed storage usage statistics for each farm configuration used in this performance study.

Table 11. Detailed Storage Performance Metrics per Content Database

Database Name	Avg. Disk Queue Length	Disk Transfers/ Second	Disk Seconds/ Transfer
Content_DB1	0.085	30.7	0.0028
Content_Logs1	0.12	117	0.001
Content_DB2	0.08	29.5	0.0027
Content_Logs2	0.0057	3.81	0.0015
Content_DB3	0.078	29.4	0.0026
Content_Logs3	0.0059	3.96	0.0015
Content_DB4	0.08	29.6	0.0027
Content_Logs4	0.0064	4.25	0.0015
Content_DB5	0.075	28.7	0.0026
Content_Logs5	0.0059	4.04	0.0015
Content_DB6	0.079	28.4	0.0028
Content_Logs6	0.0056	3.83	0.0015
Content_DB7	0.076	28.2	0.0027
Content_Logs7	0.0056	3.77	0.0015

Content_DB8	0.082	29.2	0.0028
Content_Logs8	0.0062	4.13	0.0015
Content_DB9	0.074	28.8	0.0026
Content_Logs9	0.006	3.96	0.0015
Content_DB10	0.068	26.7	0.0025
Content_Logs10	0.0056	3.76	0.0015
Content_DB11	0.039	11.1	0.0035
Content_Logs11	0.0072	4.13	0.0017
TempDB 1	6.2	1,455	0.0043
TempDB 2	2.6	441	0.0059

With approximately ~2400 IOPS and a 28 disk backend, this study showed that the Dell Compellent Storage Center Series 40 array was capable of handling a collaboration workload of up to 10,000 concurrent users.

Summary

A SharePoint 2010 farm consists of multiple servers, each of which is provisioned with different SharePoint components. A large SharePoint farm in general is the best choice for large enterprises with relatively high concurrent user load.

These farms employ three-tier architecture, and the reference architecture used in this performance study enables high availability at all tiers of the farm, while providing 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 is intended to describe the performance capacity of a large SharePoint 2010 farm using Dell servers and storage. 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 well below one second. The Dell Compellent Storage Center Series 40 array provided highly optimal performance for the SharePoint 2010 deployment used in this performance study.

References

Dell SharePoint solutions: <http://dell.com/sharepoint>

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

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

F5 | Dell Solutions: <http://www.f5.com/dell>

F5 | Microsoft Technical Forum: <http://devcentral.f5.com/microsoft>

Capacity and sizing overview for SharePoint 2010: <http://technet.microsoft.com/en-us/library/ff758647.aspx>