



Scale Test Report for Very Large Scale Document Repositories

This document is provided “as-is”. Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.

© 2011 Microsoft Corporation. All rights reserved.

Scale Test Report for Very Large Scale Document Repositories

Paul Andrew, Paul Learning, Barry Waldbaum, Frank Marasco
Microsoft Corporation
October 2011

Applies to: Microsoft® SharePoint® Server 2010, Microsoft FAST Search Server 2010 for SharePoint.

Summary

This white paper provides details about a test lab that was run at Microsoft to show large scale SharePoint Server 2010 content databases. It includes information about how two SharePoint Server content databases were populated with a total of 120 million documents over 30 terabytes (TB) of SQL Server® databases. It details how this content was indexed by FAST Search Server 2010 for SharePoint. It describes load testing that was performed on the completed SharePoint Server and FAST Search Server 2010 for SharePoint and shows results from that testing and conclusions about the results from the testing.

Contents

Introduction	5
Goals for the testing.....	5
Hardware Partners Involved	5
Definition of Tested Workload.....	6
Description of the document archive scale out architecture	7
The test transactions that were included	7
Test Transaction Definitions and Baseline Settings	8
Test Baseline Mix	9
Test Series	9
Test Load	11
Resource Capture During Tests.....	11
Test Farm Hardware Architecture Details	11
Virtual Servers.....	14
Disk Storage	15
Test Farm SharePoint Server and SQL Server Architecture	16
SharePoint Farm IIS Web Sites.....	17
SQL Server Databases	18
FAST Search Server 2010 for SharePoint Content Indexes	19
The Method, Project Timeline and Process for Building the Farm	20
Project Timeline	20
How the sample documents were created.....	20
Performance Characteristics for Large-scale Document Load.....	20
Input-Output Operations per Second (IOPS)	22
FAST Search Server 2010 for SharePoint Document Crawling	24
Results from Testing	25
Test Series A – Vary Users.....	25
Test Series B – Vary SQL Server RAM.....	27
Test Series C – Vary Transaction Mix	30
Test Series D – Vary Front-End Web Server RAM	33
Test Series E – Vary Number of Front-End Web Servers	36
Test Series F – Vary SQL Server CPUs.....	39
Service Pack 1 (SP1) and June Cumulative Update (CU) Test	42
SQL Server Content DB Backups	43

Conclusions 44

Recommendations 44

 Recommendations related to SQL Server 2008 R2 44

 Recommendations related to SharePoint Server 2010 44

 Recommendations related to FAST Search Server for SharePoint 2010 44

References 45

Introduction

Goals for the testing

This white paper describes the results of large scale SharePoint Server testing that was performed at Microsoft in June 2011. The goal of the testing was to publish requirements for scaling document archive repositories on SharePoint Server to a large storage capacity. The testing involved creating a large number of typical documents with an average size of 256 KB, loading them into a SharePoint farm, creating a FAST Search Server 2010 for SharePoint index on the documents and the running tests with Microsoft Visual Studio® 2010 Ultimate to simulate usage. With this testing we wanted to demonstrate both scale-up and scale-out techniques. Scale-up refers to using additional hardware capacity to increase resources and scale a single environment which for our purposes means a SharePoint content database. A SharePoint content database means all site collections, all metadata, and binary large objects (BLOBs) associated with those site collections that are accessed by SharePoint Server. Scale-out refers to having multiple environments, which for us means having multiple SharePoint content databases. Note that a content database is not just a SQL Server database but also includes various configuration data and any document BLOBs regardless of where those BLOBs are stored.

The workload that we tested for this report is primarily about document archive. This includes a large number of typical Microsoft Office documents that are stored for archival purposes. Storage in this scenario is typically for the long term with infrequent access.

Hardware Partners Involved

This test was made possible by support from several Microsoft hardware partners.

NEC Corporation of America

NEC provided an NEC Express5800/A1080a (GX) server containing 8 CPUs (processors) and 1 terabyte (TB) of total RAM. Each processor contained 8 cores for a total of 64 cores for the server. As detailed below, this server was used to run Microsoft Hyper-V with a number of virtual machines that made up the SharePoint Server and FAST Search Server 2010 for SharePoint farms.



Figure 1 - NEC Express Server 5800

Source:

Specifications for NEC Express 5800/A1080a server

- 8x Westmere CPU (E7-8870) each with 10 processor cores
- 1TB memory. Each Processor Memory Module has 1 CPU (10 cores) and 16 DIMMs.
- 2x dual port 8G FC HBA
- 5 HDDs

Intel

Intel provided a second NEC Express5800/A1080a server also containing 8 CPUs (processors) and 1 TB of RAM. Intel further upgraded this computer to Westmere EX CPUs each containing 10 cores for a total of 80 cores for the server. As detailed below, this server was used to run Microsoft SQL Server and FAST Search Server 2010 for SharePoint indexers directly on the computer without using Hyper-V.

EMC

EMC provided an EMC VNX 5700 SAN containing 300 TB of high performance disk.



EMC VNX5700 Unified Storage

Source: <http://www.emc.com/collateral/software/15-min-guide/h8527-vnx-virt-msapp-t10.pdf>

Specifications for EMC VNX 5700:

- 2 TB drives, 15 per 3U DAE, 5 units = total 75 drives, 150 TB raw storage
- 600 GB drives, 25 per 2U DAE, 10 units = total 250 drives, 150 TB raw storage
- 2x Storage Processors
- 2x Backup Battery Units

Definition of Tested Workload

This load test was designed to show large document archive capabilities of SharePoint Server 2010. The document archive workload is characterized by having a large number of documents that are added to (or ingested) slowly, infrequently accessed and almost never updated.

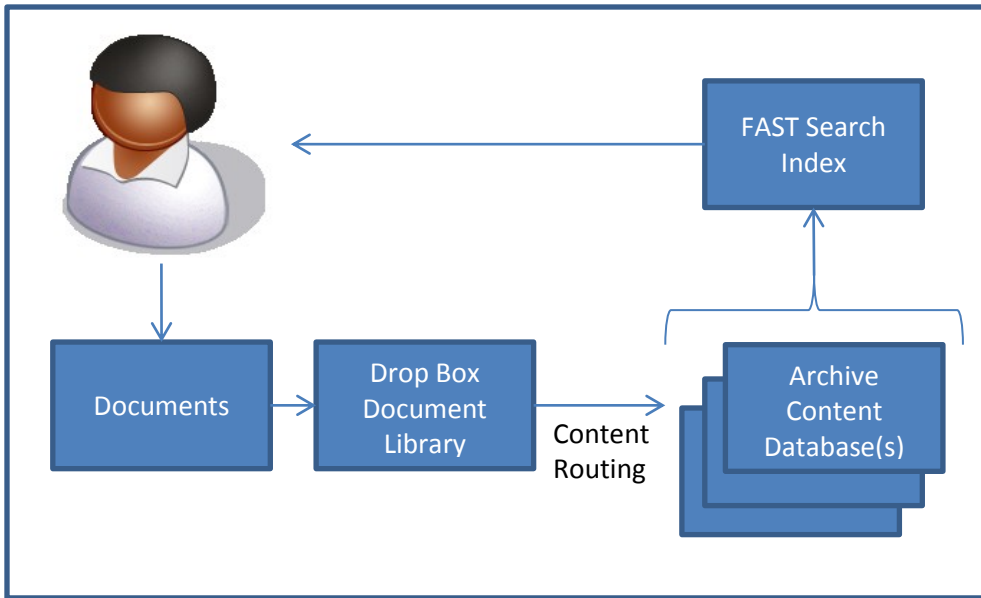


Figure 2 - Working with Large Document Archives

Description of the document archive scale out architecture

Content routing is recommended for a SharePoint farm with multiple content databases in order to send documents to the correct content database from an initial drop library. In the tests described in this report, content routing was not configured and we focused on scalability and performance of the installation.

While content routing is used to ingest documents into one of multiple SharePoint content databases, FAST Search Server 2010 for SharePoint can be used to optimally locate a document in one or more content databases. FAST Search Server 2010 for SharePoint builds an index with all documents from all content databases and searches can use metadata, refiners for selecting by date, author, or other properties and also by full text search.

The test transactions that were included

This white paper includes the results of a series of performance tests that were conducted on SharePoint Server 2010 and FAST Search Server 2010 for SharePoint in a document archive scenario. This section includes an explanation of the testing methodology that was used for tests that are discussed in this paper. Deviations from this methodology are noted where data is presented.

Workload

Important: It is important to note that the specific capacity and performance figures presented in this article are different from the figures in real-world environments. The figures that are presented are intended to provide a starting point for the design of an appropriately scaled environment. After you have completed your initial system design, test the configuration to determine whether your system will support the factors in your environment.

Testing workloads were designed in accordance with a large document archive storage scenario and are intended to help develop estimates of how different farm configurations are affected by a large-scale document repository scenario.

The test farm depicted in this scenario was designed to allow both scale out and scale up to accommodate additional capacity as required.

The ability to scale is as critical for small-scale implementations as it is for large-scale document archive scenarios. Scaling out allows you to add more servers to your farm (or farms), such as additional front end web servers or Application Servers. Scaling up allows you to increase the capacity of your existing servers by adding faster CPUs and/or memory to increase throughput and performance. Content routing should also be leveraged in archive scenarios to allow users to simply “drop” a file and have it dynamically routed to the proper document library and folder, if applicable, based on the metadata of the file.

Test Transaction Definitions and Baseline Settings

This section defines the test transactions and other baseline settings, and it provides an overview of the test process that was used for each scenario. Detailed information such as test results and specific parameters are given in each of the test results sections later in this white paper.

Baseline Item	Baseline Item Description	Baseline Setting (or Transaction Percent)
Document Upload	Upload a document to one of the Document Centers. One unique Folder and File were created in each Document Center each hour, 24 hours a day.	1%
Document Download (Open)	Download or open a document	30%
Browse	Access of a random Document Center Home page, Document Library List view page, or Folder list view page.	40%
Search	A random search query submitted to the FAST Search Center.	30%
Think Time	The time between transactions for each user. This is to represent the time a user spends reading or thinking between accesses to web pages.	10 seconds
Concurrent Users	The number of users connecting to the SharePoint farm from test agents to the SharePoint front-end Web servers. This does not represent a possible total user base, since in a typical environment a small proportion of total users will concurrently access the system.	10,000
Test Duration	The length of time the test is run	1 hour
Web Caching	Whether web caching is turned on for the front-end Web servers or not	On
FAST Content Indexing	Whether FAST Content indexing is operating during the test or not	Paused
Number WFEs	The number of front-end Web servers in the SharePoint farm that were used during the test	3 per content database
User Ramp	Each test was started off with 1,000 users and ramped to the target user load in 100 user increments. A 30 second ramp time was used and a 10 second step time.	100 users per 30 seconds

Test Agents	Visual Studio 2010 Ultimate was used to simulate the user transaction load. One test controller virtual machine and 19 test agent virtual machines were used to create this load.	19
-------------	---	----

Table 1 – Test Transactions and Baseline Settings

Test Baseline Mix

This section defines the test mixes that were taken advantage of and provides an overview of the test results for each test mix scenario.

The test mix used for each test varied, based on the particular test and load targets. All tests were conducted using Visual Studio 2010 Ultimate and were recorded, code-free scripts that were generated exclusively by Visual Studio. Specific data points for each test were populated, and then the test mix was run for different periods using different numbers of concurrent users to determine farm capacities and limits.

Notes

All tests conducted in the lab were run using 10 seconds of “think time.” Think time is a feature of the Microsoft Visual Studio 2010 Ultimate Test Controller that allows you to simulate the time that users pause between clicks on a page in a real-world environment.

The mix of operations that was used to measure performance for the purpose of this white paper is artificial. All results are only intended to illustrate performance characteristics in a controlled environment under a specific set of conditions. These test mixes are made up of an uncharacteristically high amount of list queries that consume a large amount of SQL Server resources compared to other operations. This was intended to provide a starting point for the design of an appropriately scaled environment. After you have completed your initial system design, test the configuration to determine whether your specific environmental variables and mix of operations will vary.

Test Series

There were six test series run which were labeled A through F. Each series involved running the baseline test with identical parameters and environment except for one parameter which was varied. The individual tests in each series were labeled after the test series followed by a number. This section outlines the individual test series that were run. Included in the list of tests is a note for which test was the same as the baseline. In other words, one test in each series did not vary the chosen parameter, but was in fact identical in all respects to the original baseline test.

Test Series A – Vary Users

This test series varies the number of users to see how the increased user load impacts the system resources in the SharePoint and FAST Search Server 2010 for SharePoint farm. Three tests were performed including 4,000 users, 10,000 users, and 15,000 users. The 15,000 user test required an increased test time to 2 hours to deal with the increased user ramp, and it also had increased front-end Web server (WFE) servers to 6 WFEs to handle the increased load.

Test	Number of users	Number of WFEs	Test Time
A.1	4,000	3	1 hour
A.2	10,000	3	1 hour (baseline)
A.3	15,000	6	2 hours

Test Series B – Vary SQL Server RAM

This test series varies the amount of RAM available to Microsoft SQL Server. Since the SQL Server computer had a large amount of physical RAM, we ran this test series to see how a server running SQL Server with less RAM might perform in comparison. Six tests were performed with the maximum SQL Server RAM set to: 16 GB, 32 GB, 64 GB, 128 GB, 256 GB, and 600 GB.

Test	SQL RAM
B.1	16 GB
B.2	32 GB
B.3	64 GB
B.4	128 GB
B.5	256 GB
B.6	600 GB – (baseline)

Test Series C – Vary Search Mix

This test series varies the proportion of searching done by the test users as compared to browsing and opening documents. The test workload applied to the farm is a mix of different user transactions, which follow the baseline by default of 30%, 40%, and 30% for Open, Browse and Search respectively. Tests in this series vary the proportion of search and hence also change the proportion of Open and Browse.

Test	Open	Browse	Search
C.1	30%	55%	15%
C.2	30%	40%	30% (baseline)
C.3	20%	40%	40%
C.4	20%	30%	50%
C.5	25%	25%	50%
C.6	5%	20%	75%

Test Series D – Vary WFE RAM

This test series varies the RAM allocated to the front-end Web servers. Also, four front-end Web servers were used for this test. The RAM on each of the 4 front-end Web servers was tested at 4 GB, 6 GB, 8 GB and 16 GB.

Test	WFE Memory
D.1	4 GB
D.2	6 GB
D.3	8 GB - (baseline)
D.4	16 GB

Test Series E – Vary Number WFEs

This test series varies the number of front-end Web servers in use. The different number of servers tested was 2, 3, 4, 5, and 6.

Test	Number of Web Front Ends
E.1	2
E.2	3 - (baseline)
E.3	4
E.4	5

Test Series F – SQL Server CPU Restrictions

This test series restricts the number of CPUs available to SQL Server. The different number of CPUs available to SQL Server tested was 2, 4, 8, 16 and 80 CPUs.

Test	CPUs Available to SQL Server
F.1	4
F.2	6
F.3	8
F.4	16
F.5	80 - (baseline)

Test Load

Tests were intended to stay below an optimal load point, or Green Zone, with a general mix of operations. To measure particular changes, tests were conducted at each point that a variable was altered. Test series were designed to exceed the optimal load point in order to find resource bottlenecks in the farm configuration. It is recommended that optimal load point results be used for provisioning production farms so that there is excess resource capacity to handle transitory, unexpected loads. For this project we defined the optimal load point as keeping resources below the following metrics:

- 75th percentile latency is less than 1 second
- Front-end Web server CPU is less than 85%
- SQL Server CPU is less than 50%
- Application server CPU is less than 50%
- FAST Search Server 2010 for SharePoint CPU is less than 50%
- Failure rate is less than 0.01

Resource Capture During Tests

During each test run, resource usage was captured by using Performance Monitor (Perfmon.exe) and Visual Studio 2010 Ultimate in order to determine the load on the test farm. The following details were captured and are shown in the reports section.

- CPU for each WFE, SharePoint application server, FAST Search Server 2010 for SharePoint Index, Fast Search Service Application (SSA), SQL Server computer
- RAM usage for each WFE, SharePoint application server, FAST Search Server 2010 for SharePoint Index, Fast SSA, SQL Server computer
- Page refresh time on all test elements
- Disk queues for each drive

Test Farm Hardware Architecture Details

The Document Center farm is the host for SharePoint Central Administration, Document Center 1, Document Center 2, Service Applications and the integrated FAST Search Center. The farm consists of three physical servers and 22 virtual servers.

Figure 3 has a diagram of the physical architecture.

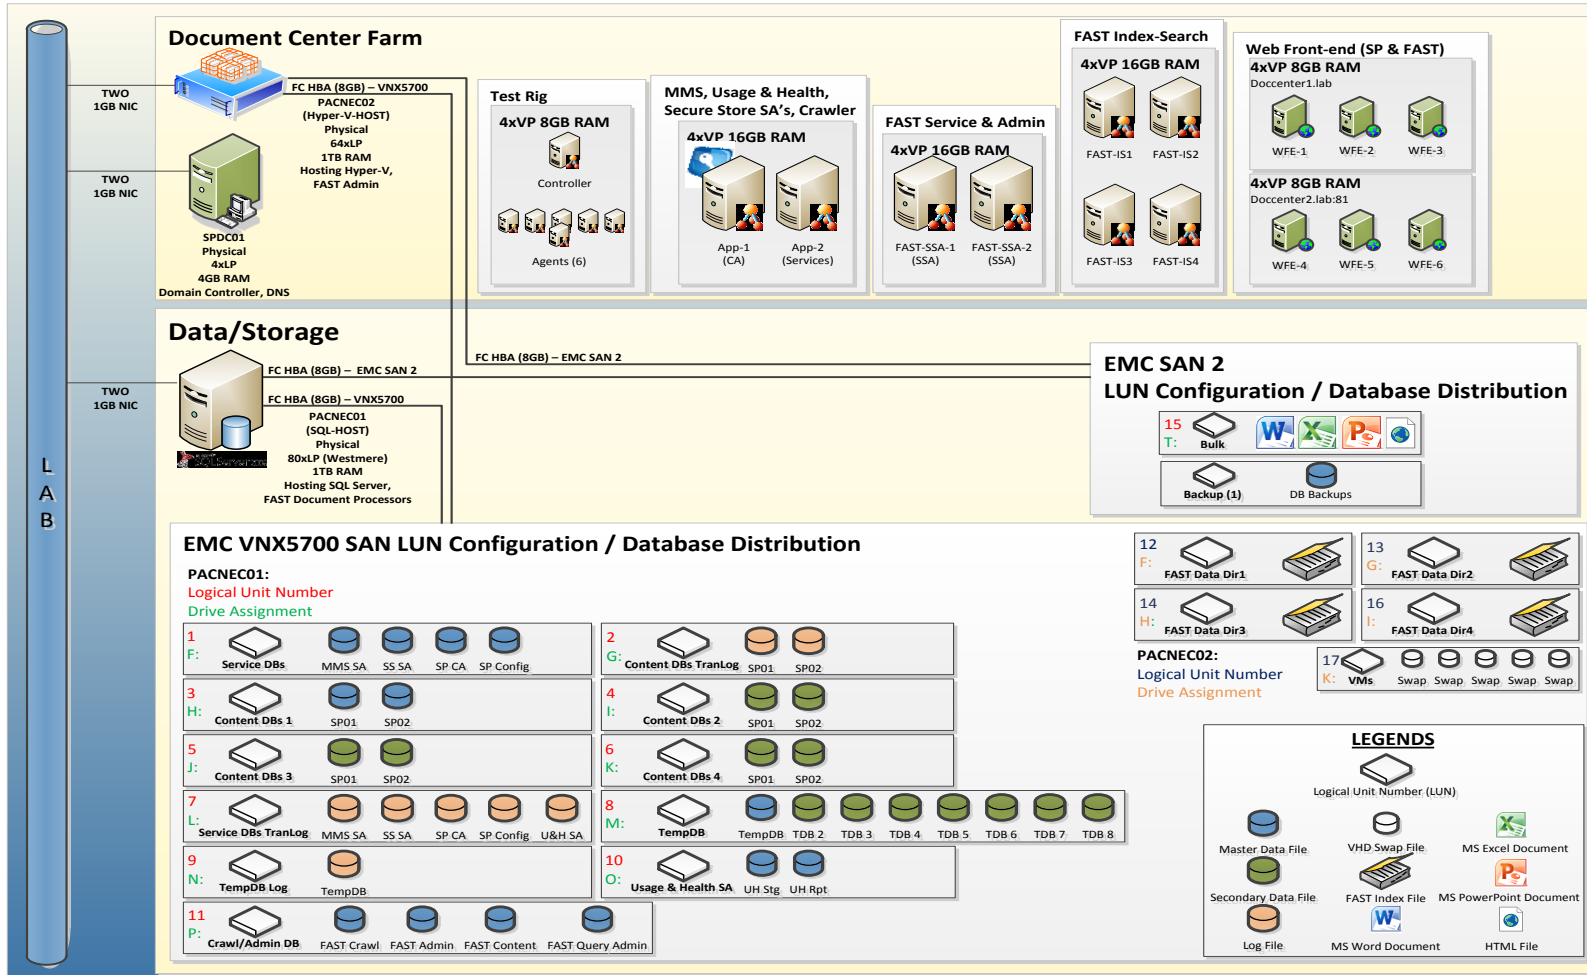


Figure 3 - Hardware Architecture Diagram

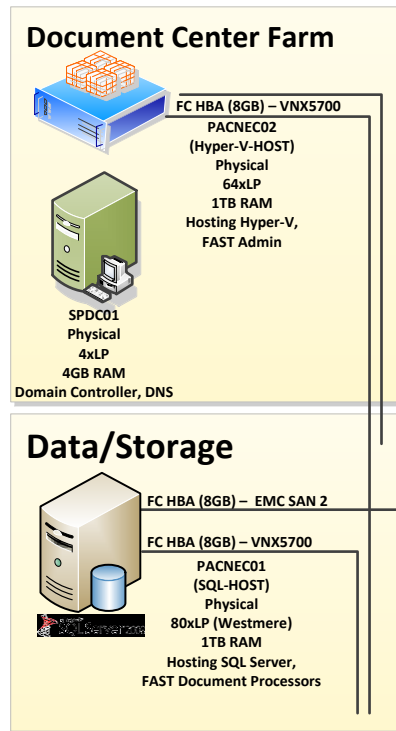


Figure 4 - Physical Servers

Hyper-threading was disabled on the physical servers because we did not need additional CPU cores and were limited to 4 logical CPUs in any one Hyper-V virtual machine. We did not want these servers to experience any performance degradation due to hyper-threading. There were three physical servers in the lab. All three physical servers plus the twenty two virtual servers were connected to a virtual LAN within the lab to isolate their network traffic from other unrelated lab machines. The LAN was hosted by a 1 GBPS Ethernet switch, and each of the NEC servers was connected to two 1 GBPS Ethernet ports.

- **SPDC01.** The Windows Domain Controller and Domain Naming System (DNS) Server for the virtual network used in the lab.
 - 4 physical processor cores running at 3.4 GHz
 - 4 GB of RAM
 - 33 GB RAID SCSI Local Disk Device
- **PACNEC01.** The SQL Server 2008 R2 hosting the master and secondary files for content DB's, Logs, and TempDB. Also ran 100 FAST Document Processors directly on this server.
 - NEC ExpressServer 5800 1080a
 - 8 Intel E7-8870 CPUs containing 80 physical processor cores running at 2.4 GHz
 - 1 TB of RAM
 - 800 GB of Direct Attached Disk
 - 2x Dual Port Fiber Channel Host Bus Adapter cards capable of 8 GB/s
 - 2x 1 GBPS Ethernet cards
- **PACNEC02.** The Hyper-V Host serving the SharePoint, FAST Search for SharePoint and Test Rig virtual machines within the farm.
 - NEC ExpressServer 5800 1080a
 - 8 Intel X7560 CPUs containing a total of 64 physical processor cores running at 2.27 GHz

- 1 TB of RAM
- 800 GB of Direct Attached Disk
- 2x Dual Port Fiber Channel Host Bus Adapter cards capable of 8 GB/s
- 2x 1 GBPS Ethernet cards

Virtual Servers

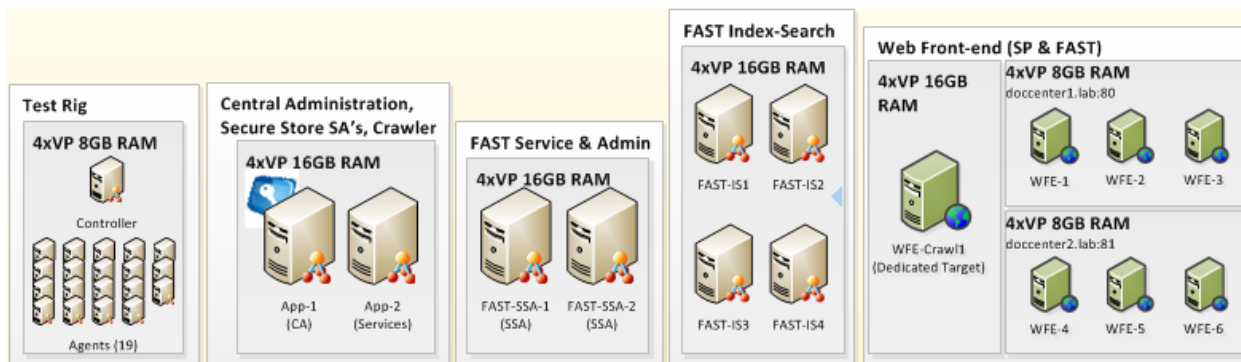


Figure 5 - Virtual Servers

These servers all ran on the Hyper-V instance on PACNECO2. All virtual servers booted from VHD files stored locally on the PACNECO2 server and all had configured access to the lab virtual LAN. Some of these virtual servers were provided direct disk access within the guest operating system to a LUN on the SAN. Direct disk access that was provided increased performance over using a VHD disk and was used for accessing the FAST Search indexes. Here is a list of the different types of virtual servers running in the lab and the details of their resources use and services provided.

Virtual Server Type	Description
Test Rigs (TestRig-1 through TestRig-20) <ul style="list-style-type: none"> • TestRig-1 is the Test Controller from Visual Studio 2010 Ultimate • TestRig-2 through TestRig-19 are the Test Agents from Visual Studio Agents 2010 that are controlled by TestRig-1 	The Test Controller and Test Agents from Visual Studio 2010 Ultimate for load testing the farm. These virtual servers were configured with 4 virtual processors and 8 GB memory. These servers used VHD for disk.
SP: Central Admin, Secure Store SA's, Crawler <ul style="list-style-type: none"> • APP-1 - SharePoint Central Administration Host and FAST Search Service Application Host. • APP-2 - SharePoint Service Applications and FAST Search Service Application Host. This application server ran the following SharePoint Shared Service Applications: <ul style="list-style-type: none"> ▪ Secure Store Service Application. ▪ FAST Search Service Application. 	These virtual machines host the SharePoint Central Administration and Service Applications used within the farm. These virtual servers were configured with 4 virtual processors and 16 GB memory. These servers used VHD for disk.
FAST Service and Administration <ul style="list-style-type: none"> • FAST-SSA-1 and FAST-SSA-2 – FAST Search Service Applications 1 and 2 respectively. 	These virtual machines host the FAST Search Service and Administration. They were each configured with 4 virtual processors, 16 GB memory and they used

	VHD for disk.
FAST Index-Search <ul style="list-style-type: none"> • FAST-IS-1, FAST-IS2, FAST-IS3, and FAST-IS4 - FAST Index, Search, Web Analyzer Nodes 1, 2, 3, and 4. 	These virtual machines host the FAST Index, and the Search and Web Analyzer Nodes used within the farm. They were configured with 4 virtual processors, 16 GB memory and they used VHD for their boot disk. They also each had direct access as disks to 3 TB of SAN LUNs for storage of the fast index.
Front-End Web server (SharePoint and FAST Search) <ul style="list-style-type: none"> • WFE-1, WFE-2, and WFE-3 - Front-end Web server #1, #2, and #3, part of a Windows load-balancing configuration hosting the first Document Center. These virtual servers were configured with 4 virtual processors and 8 GB memory. • WFE-4, WFE-5, and WFE-6 - Front-end Web server #4, #5, and #6, part of a Windows load-balancing configuration hosting the second Document Center. These virtual servers were configured with 4 virtual processors and 8 GB memory. 	These virtual machines host all of the front-end web servers and a dedicated FAST crawler host within the farm. Each content database contained one document center which was configured with 3 load-balanced SharePoint Server WFEs. This was done to facilitate the text mix for load testing across the two content databases. In a real farm each WFE would target multiple content databases. These servers used VHD for disk.

Disk Storage

The storage consists of EMC VNX5700 Unified Storage. The VNX5700 array was connected to each of the physical servers PACNEC01 and PACNEC02 with 8 GBPS Fiber Channel. Each of these physical servers contains two Fiber Channel host bus adapters so that it can connect to both of the Storage Processors on the primary SAN, which provides redundancy and allows the SAN to balance LUNs across the Storage Processors.

Storage Area Network - EMC VNX5700 Array

An EMC VNX5700 array (<http://www.emc.com/products/series/vnx-series.htm#/1>) was used for storage of the SQL Server databases and FAST Search Server 2010 for SharePoint search index. The VNX5700 as configured included 300 terabyte (TB) of raw disk. The array was populated with 250x 600GB 10,000 RPM SAS drives and 75x 2TB 7,200 RPM Near-line SAS drives (near-line SAS drives have SATA physical interfaces and SAS connectors whereas the regular SAS drives have SCSI physical interfaces). The drives were configured in a RAID-10 format for mirroring and striping. The configured RAID volume in the Storage Area Network (SAN) was split across 3 pools and LUNs are allocated from a specific pool as shown in Table 2.

Pool #	Description	Drive Type	User Capacity (GB)	Allocated (GB)
0	FAST	SAS	31,967	24,735
1	Content DB	SAS	34,631	34,081
2	Spare – not used	NL SAS	58,586	5,261

Table 2 - SAN Pools Allocated

The Logical Unit Numbers (LUNs) on the VNX 5700 were defined as shown in Table 3.

LUN #	Description	Size (GB)	Server	Disk Pool #	Drive Letter
0	SP Service DB	1,024	PACNEC01	0	F
1	PACNEC02 extra space	5,120	PACNEC02	0	
2	FAST Index 1	3,072	PACNEC02	0	F
3	FAST Index 2	3,072	PACNEC02	0	G
4	FAST Index 3	3,072	PACNEC02	0	H
5	FAST Index 4	3,072	PACNEC02	0	I
6	SP Content DB 1	7,500	PACNEC01	1	H
7	SP Content DB 2	6,850	PACNEC01	1	I
8	SP Content DB 3	6,850	PACNEC01	1	J
9	SP Content DB 4	6,850	PACNEC01	1	K
10	SP Content DB TransLog	2,048	PACNEC01	1	G
11	SP Service DB TransLog	512	PACNEC01	0	L
12	Temp DB	2,048	PACNEC01	1	M
13	Temp DB Log	2,048	PACNEC01	0	N
14	SP Usage Health DB	3,072	PACNEC01	0	O
15	FAST Crawl DB / Admin DB	1,024	PACNEC01	1	P
16	Spare – not used	5,120	PACNEC01	2	
17	Bulk Office Doc Content	3,072	PACNEC01	Additional	T
18	WMs Swap Files	1,024	PACNEC02	Additional	K
19	DB Backup 1	16,384	PACNEC01	Additional	R
20	DB Backup 2	16,384	PACNEC01	Additional	S

Table 3 - Logical Unit Numbers

Storage Area Network - Additional Disk Array

An additional lower performance disk array was used for backup purposes and to host the bulk Office document content that was loaded into the SharePoint Server 2010 farm. This array was not used during test runs.

Test Farm SharePoint Server and SQL Server Architecture

The logical architecture was defined to demonstrate the recommended limits of SharePoint Server 2010. The architecture consists of two Web applications, each containing a single site collection within a single unique content database. Each content database was loaded with 60 million documents of type **Microsoft Word (.docx)**, **Excel (.xlsx)**, **PowerPoint (.pptx)** and **Hyper-text Markup Language (.html)** pages, averaging **250 kilobytes (KB)** in size. Content database size was approximately **15 TB** each, for a total corpus of **30 TB**. The logical architecture for the large-scale lab is shown in Figure 6.

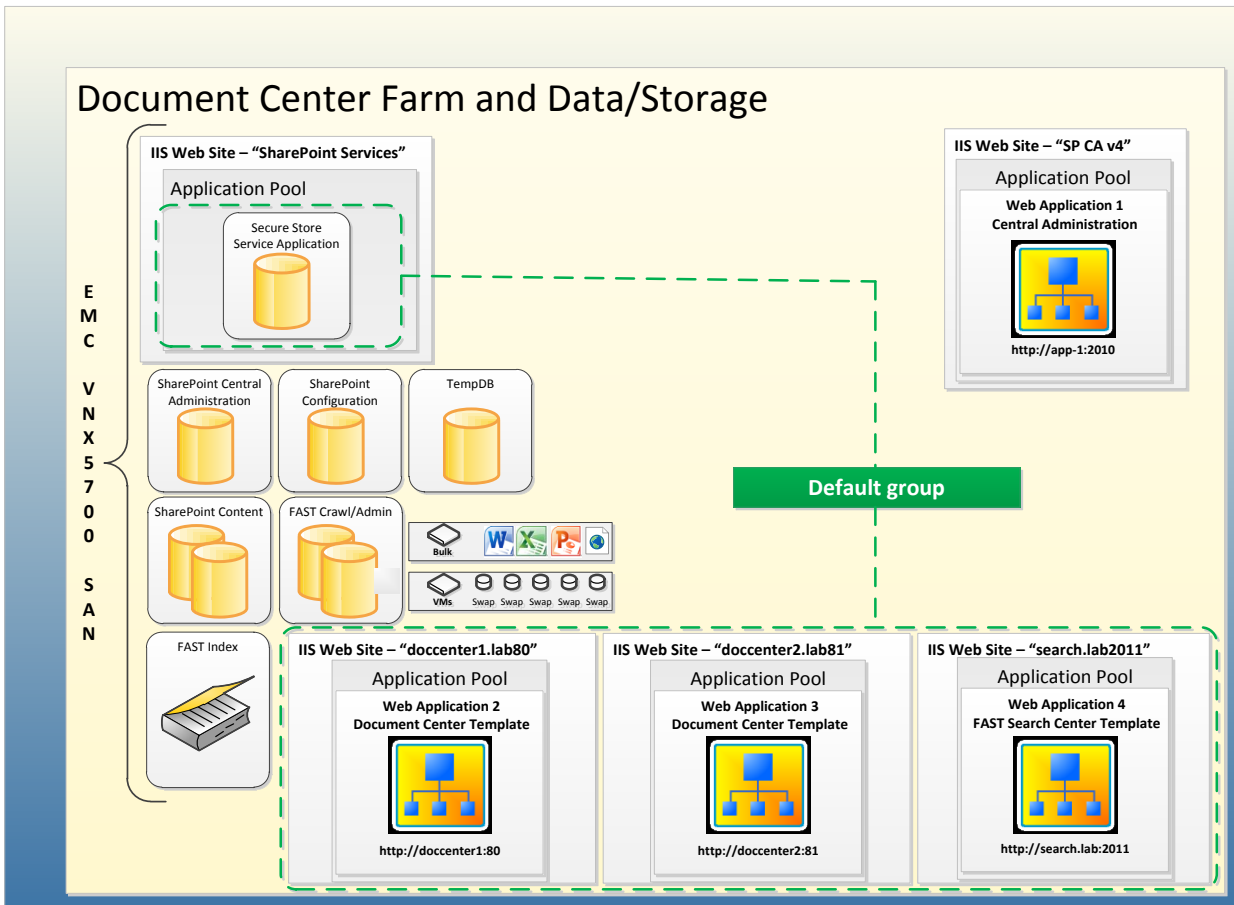


Figure 6 - Software

Architecture

The SharePoint Document Center farm is intended to be used in a document archival scenario and was designed to accommodate a large number of documents stored in several document libraries. Document libraries were limited to roughly one million documents each and a folder hierarchy limited the documents per container to approximately 2,000 items. This was done solely to accommodate the large document loading process and prevent the load time from decreasing after exceeding 1 million items in a document library.

SharePoint Farm IIS Web Sites

The two content site collections took advantage of the Document Center template. The Search Center site collection leveraged the FAST Search Center template. Each site collection was in a unique web application. Each Web application used a separate application pool.

IIS Web Site - SharePoint Services

The SharePoint Services IIS web site hosts the shared services used in SharePoint Server 2010. For the purposes of this lab, the Secure Store was utilized.

IIS Web Site - SharePoint Central Administration v4

The SharePoint Central Administration IIS Web Site hosts the Central Administration site and user interface for SharePoint Server 2010.

IIS Web Site - Document Center 1

The Document Center 1 IIS Web Site hosts the first Document Center archive.

IIS Web Site – Document Center 2

The Document Center 2 IIS Web Site hosts the second Document Center archive.

IIS Web Site – FAST Search Center

The Fast Search Center IIS Web Site hosts the search user interface for the farm.

At 70 million and above the crawl database started to become noticeably slower and some tuning work was required to take it from 100 million to 120 million.

SQL Server Databases

The following SQL Server databases are hosted on the EMC VNX 5700 Storage Area Network (SAN).

DB Name	Purpose	Size (MB)
SharePointAdminContent_<GUID>	SharePoint Central Administration Database	768
SharePoint_Config	SharePoint Configuration Database	1,574
System Databases – tempdb	SQL Server Temporary Database	16,384
ReportServer	A Microsoft SQL Server database that stores all report metadata including report definitions, report history and snapshots, and scheduling information.	10
ReportServerTempDB	A Microsoft SQL Server database that stores all of the temporary snapshots while reports are running.	3
SPContent01 (Document Center 1 content database)	SharePoint content databases	15,601,286
SPContent02 (Document Center 2 content database)	SharePoint content databases	15,975,266
FAST_Query_CrawlStoreDB_<GUID>	Crawler store for the FAST Search Query Search Service Application. This crawl store database is only used for user profiles (People Search).	15
FAST_Query_DB_<GUID>	Administration database for the FAST Search Query Search Service Application.	125
FAST_Query_PropertyStoreDB_<GUID>	Stores the metadata properties and security descriptors for the user profile items in the people search index. It is involved in property-based people search queries and returns standard document attributes for people search query results.	173

FASTContent_CrawlStoreDB_<GUID>	Crawler store for the FAST Search Content Search Service Application. This crawl store database is used for all crawled items except user profiles.	502,481
FASTContent_DB_<GUID>	Administration database for the FAST Search Content Search Service Application.	23
FASTSearchAdminDatabase	Administration database for the FAST Search Server 2010 for SharePoint farm. Stores and manages search setting groups, keywords, synonyms, document and site promotions and demotions, property extractor inclusions and exclusions, spell check exclusions, best bets, visual best bets, and search schema metadata.	4
WSS_Content_FAST_Search	FAST Search Center content database.	52
LoadTest2010	Load test results repository	4,099

Table 4 - SQL Server Databases

FAST Search Server 2010 for SharePoint Content Indexes

The FAST Search Server 2010 for SharePoint data directories are using a Hyper-V pass through drive directly to the SAN.

On the virtual server FAST-IS1, the data directory is using 745 GB of the 3 TB with no temp space being used (everything was cleaned up).

Table 5 shows the data storage in the FAST Search Server 2010 for SharePoint index file folders stored on the SAN.

Name	Purpose	Number Files	Size (GB)
data_fixml	Index Source used to Create Index	6 million	223
data_index	Actual Search Index used by Queries	3,729	490
sprel	SharePoint Relevancy Information. Used for boosting popular search results to top of list.	9	3
webanalyzer	Boosting search result order for more commonly linked documents.	135	12

Table 5 - Storage Used by 1 of 4 FAST Indexes

The Method, Project Timeline and Process for Building the Farm

Project Timeline

This is the approximate project timeline.

- Plan Farm Architecture 2 weeks
- Install Server and SAN Hardware 1 week
- Build Virtual Machines for Farm 1 week
- Creating Sample Content Items 2 weeks
- Load Items to SharePoint Server 3 weeks
- Develop Test Scripts 1 week
- FAST Search indexing of Content 2 weeks
- Load Testing 3 weeks
- Report Writing 2 weeks

How the sample documents were created

In order to provide a realistic document archive scenario, document uniqueness was critical. Two separate utilities were used; the first to create unique documents, and the second to read these files from disk and load them directly into targeted SharePoint Web Applications and document libraries.

Tool to Create Large Numbers of Documents

Documents were created using a command-line tool named **Bulk Loader**, which was written using the Microsoft .NET 4.0 Framework. This tool utilizes a dump file of Wikipedia content as input to allow the creation of up to 10 million unique documents to a disk location. Stock images are used to replace image references from the Wikipedia dumps. This tool is available as source code from <http://code.msdn.microsoft.com/Bulk-Loader-Create-Unique-eeb2d084>.

Tool to Load Documents into SharePoint

Documents were added to SharePoint Server by using a command-line tool named **LoadBulk2SP**, which was written using C# and the Microsoft .NET 3.5 Framework to be compatible with SharePoint Server. This tool takes the Bulk Loader tool disk output files as input and mimics the same folder and file structure directly into SharePoint Server using targeted web applications and document libraries specified in the application configuration. Using this tool, over **100** million 250 KB documents were loaded into SharePoint Server with a peak performance of **233** documents-per-second, and an overall average load time of **137** documents-per-second. This tool is available as source code on <http://code.msdn.microsoft.com/Load-Bulk-Content-to-3f379974>.

Performance Characteristics for Large-scale Document Load

Document loading was accomplished by using the **LoadBulk2SP** tool defined earlier in this document. This tool leverages the *SubFolderCollection.Add()* method to add new folders to specified document libraries, and the *SPFileCollection.Add()* method to add files directly into the SharePoint document library folders created. The folder and file structure created in SharePoint Server mimics the output hierarchy created by the Bulk Loader tool.

Document Library Content Database Sizes

Following are the details of each document library content database sizing, including SQL Server Filegroups, Primary and Secondary files used within the farm.

SQL Content File	FileGroup	LUN	Size (KB)	Size (MB)	Size (GB)	Size (TB)
SPCPrimary01.mdf	Primary	H:/	53,248	52.000	0.050	0.000

SPCData0102.mdf	SPCData01	I:/	3,942,098,048	3,849,697.312	3,759.470	3.671
SPCData0103.mdf	SPCData01	J:/	4,719,712,768	4,609,094.500	4,501.068	4.395
SPCData0104.mdf	SPCData01	K:/	3,723,746,048	3,636,470.750	3,551.240	3.468
SPCData0105.mdf	SPCData01	H:/	3,371,171,968	3,292,160.125	3,215.000	3.139
SPCData0106.mdf	SPCData01	O:/	4,194,394	4,096.087	4.000	0.003
Document Center 1	Totals:		15,760,968,474	15,391,570.775	15,030.820	14.678
SPCPPrimary02.mdf	SPCData02	H:/	52,224	51.00	0.049	0.000
SPCData0202.mdf	SPCData02	I:/	3,240,200,064	3,164,257.875	3,090.095	3.017
SPCData0203.mdf	SPCData02	J:/	3,144,130,944	3,070,440.375	2,998.476	2.928
SPCData0204.mdf	SPCData02	K:/	3,458,544,064	3,377,484.437	3,298.324	3.221
SPCData0205.mdf	SPCData02	H:/	3,805,828,608	3,716,629.500	3,629.521	3.544
SPCData0206.mdf	SPCData02	O:/	2,495,168,448	2,436,687.937	2,379.578	2.323
Document Center 2	Totals:		16,143,924,352	15,765,551.125	15,396.046	15.035
Corpus Total:			31,904,892,826	31,157,121.900	30,426.876	29.713

Table 6 - SQL Server Database Sizes

Document Library Hierarchies, Folders and Files

Following are the details of the document library hierarchies, total number of folders and documents generated for each Document Center using the LoadBulk2SP tool. The totals across both Document Centers are **60,234 Folders** and **120,092,033 Files**.

Document Center 1

The total number of folders and files contained in each document library in the content database are shown in Table 7. As stated previously, documents were limited to 1 million per document library strictly for the purposes of a large content load process. For SharePoint 2010 farm architecture results and advice related to large document library storage, please refer to an earlier test report, [Estimate performance and capacity requirements for large scale document repositories in SharePoint Server 2010](http://technet.microsoft.com/en-us/library/hh395916.aspx) (<http://technet.microsoft.com/en-us/library/hh395916.aspx>), which focused on scaling numbers of items in a document library.

Please also reference SharePoint Server 2010 boundaries for items in document libraries and items in content databases as detailed in [SharePoint Server 2010 capacity management: Software boundaries and limits](http://technet.microsoft.com/en-us/library/cc262787.aspx) (<http://technet.microsoft.com/en-us/library/cc262787.aspx>) on TechNet.

Document Center 1		
	Counts	
Document Library	Folders	Files
DC1 TOTALS:	30,447	60,662,595

Table 7 - Document Libraries in Document Center 1

Document Center 2

The total number of folders and files contained in each document library in the content database are shown in Table 8.

Document Center 2		
	Counts	
Document Library	Folders	Files
DC2 TOTALS:	29,787	59,429,438
DC1 TOTALS:	30,447	60,662,595
CORPUS TOTALS:	60,234	120,092,033

Table 8 - Document Libraries in Document Center 2

Following are statistic samples from the top five LoadBulk2SP tool runs, using four concurrent processes, each with 16 threads targeting different Document Centers, document libraries and input folders and files.

Run 26:		Time	Seconds	Folders	Files	Docs/Sec
5 Folders @ 2k files	Hours	0	0	315	639,980	233
	Minutes	45	2,700			
	Seconds	46	46			58264
	Total:		2,746			

Run 9:		Time	Seconds	Folders	Files	Docs/Sec
30 Folders @ 2k files	Hours	5	18,000	1,920	3,839,864	178
	Minutes	58	3,480			
	Seconds	46	46			
	Total:		21,526			

Run 10:		Time	Seconds	Folders	Files	Docs/Sec
30 Folders @ 2k files	Hours	6	21,600	1,920	3,839,881	162
	Minutes	33	1,980			
	Seconds	50	50			
	Total:		23,630			

Run 8:		Time	Seconds	Folders	Files	Docs/Sec
30 Folders @ 2k files	Hours	6	21,600	1,920	3,839,857	155
	Minutes	51	3,060			
	Seconds	30	30			
	Total:		24,690			

Run 7:		Time	Seconds	Folders	Files	Docs/Sec
30 Folders @ 2k files	Hours	6	21,600	1,920	3,839,868	154
	Minutes	55	3,300			
	Seconds	0	0			
	Total:		24,900			

Table 9 - Detailed performance results from LoadBulk2SP

Input-Output Operations per Second (IOPS)

SQLIO is a stress tool used to determine the I/O capacity of a given configuration. It was run on the system after performance tests had been completed. Therefore, several disks backed by SAN LUNs could not be included as they had too much existing data on them. The SQLIO test runs on each drive letter individually and then does a test on all drives at once. You can see the IOPS/GB in the right column which is calculated by dividing the IOPS by the drive capacity. For these drives all being tested at once, we achieved 105,730 IOPS.

IOPS as tested with SQLIO tool

LUN	LUN Description	Size (GB)	Reads IOPS (MAX)	Writes IOPS (MAX)	Total IOPS (MAX)	IOPS per GB
F:	SP Service DB	1024	2,736	23,778	26,514	25.89
G:	Content DBs TranLog	2048	3,361	30,021	33,383	16.30
L:	Service DBs TranLog	512	2,495	28,863	31,358	61.25
M:	TempDB	2048	2,455	21,778	24,233	11.83
N:	TempDB Log	2048	2,751	29,522	32,273	15.76
O:	Content DBs 5	3,072	2,745	28,767	31,511	10.26
P:	Crawl/Admin DBs	1024	2,603	22,808	25,411	24.81
	All at once	11776	16,665	89,065	105,730	8.98
TOTAL:		11,776	19,145	185,536	310,412	
AVERAGE:		1,682	2,735	26,505	38,801	22

Table 10 - IOPS test results for SAN from SQLIO tool

IOPS Achieved during Load Testing

Performance Monitor jobs were run consistently along with concurrent FAST Indexing, content loading, and Visual Studio load tests running. The following table reflects the maximum IOPS achieved by LUN and identifies each **LUN**, **Description**, **Total size**, **Max Reads**, **Max Writes**, **Totals IOPS**, and **IOPS per GB**.

Because these results were obtained during testing, they reflect the IOPS that the test environment was able to drive into the SAN. Because drives H:, I:, J:, and K: were able to be included, the total IOPS achieved was much higher than for the SQLIO testing.

LUN	LUN Description	Size (GB)	Reads IOPS (MAX)	Writes IOPS (MAX)	Total IOPS (MAX)	IOPS per GB
G:	Content DBs TranLog	2048	5,437	11,923	17,360	8.48
H:	Content DBs 1	6,850	5,203	18,546	23,749	3.47
I:	Content DBs 2	6,850	5,284	11,791	17,075	2.49
J:	Content DBs 3	7,500	5,636	11,544	17,180	2.29
K:	Content DBs 4	6,850	5,407	11,146	16,553	2.42

L:	Service DBs TranLog	512	5,285	10,801	16,086	31.42
M:	TempDB	2048	5,282	11,089	16,371	7.99
N:	TempDB Log	2048	5,640	11,790	17,429	8.51
O:	Content DBs 5	3072	5,400	11,818	17,218	5.60
P:	Crawl/Admin DBs	1024	5,249	11,217	16,467	16.08
TOTAL:		31,365	53,824	121,667	175,491	
AVERAGE:		3,136	5,382	12,167	17,549	5.60

Table 11 - IOPS as measured from Perfmon Logs

FAST Search Server 2010 for SharePoint Document Crawling

Crawling SharePoint sites for search is done by using the SharePoint crawler configured to feed to the FAST Content Distributors. The content Search Service Application (SSA) was configured to run on two servers, APP-1 and APP-2, and the query SSA was run on the servers FAST-1 and FAST-2.

100 FAST indexing document processors were run on the SQL Server machine. We took this screenshot from task manager on the computer showing the activity while both document processor work and a 10,000 user load test were running with SQL Server also located on this computer.

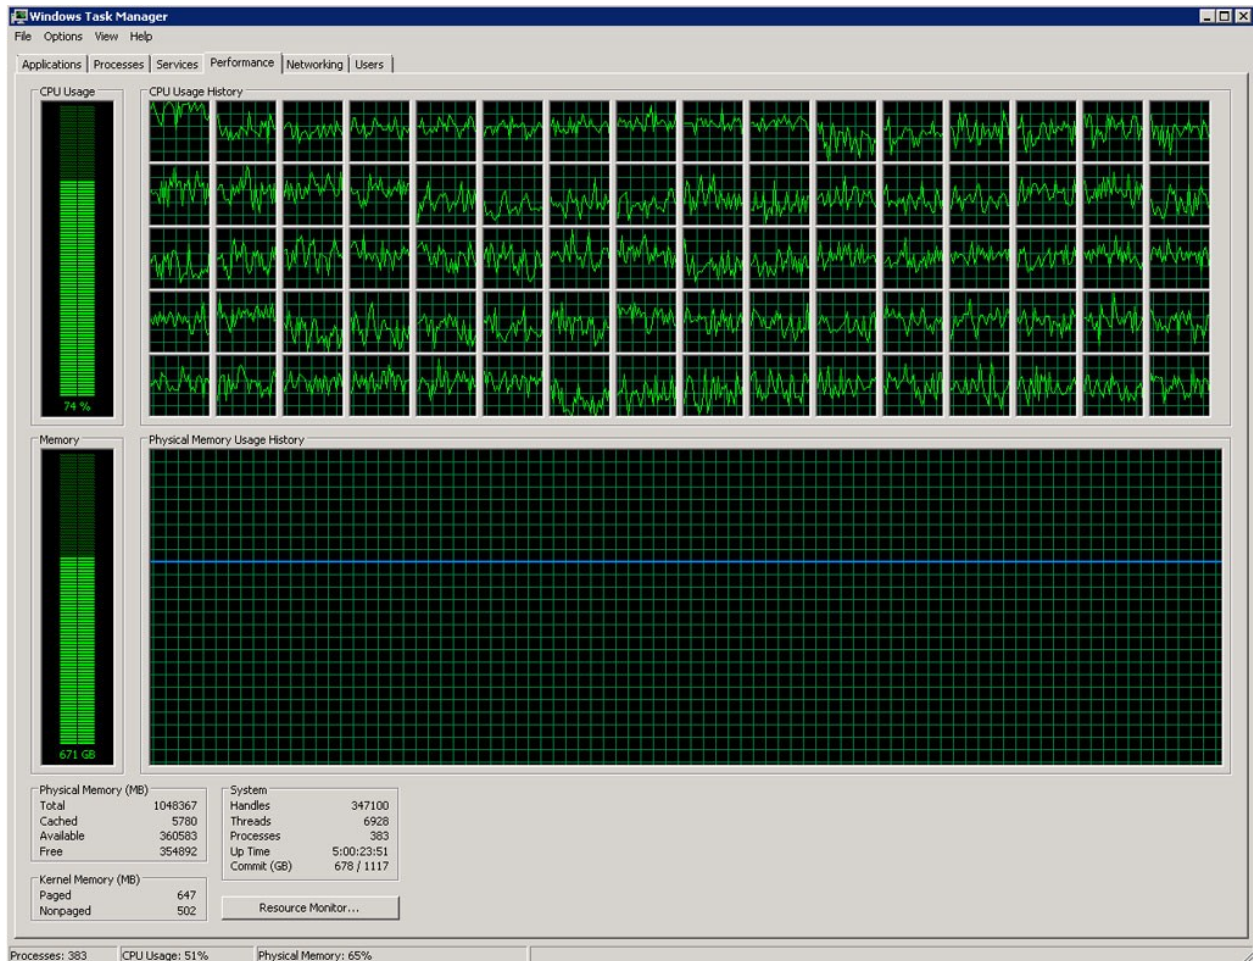


Figure 7 - Task Manager on PACNEC01 during FAST Indexing and Load Test

Results from Testing

In order to generate a significant load during testing, the following software was used: Visual Studio 2010 Ultimate, Visual Studio 2010 Load Control, and Microsoft Visual Studio Agents 2010 ¹. A Test Rig is required in order to simulate a number of users as well as produce a significant load. A Test Rig is made up of a Test Controller machine and one or more Test Agent machines. The test controller manages and coordinates with agent machines, and the agents are used to generate load against SharePoint Server. The Test controller is also responsible for collecting performance monitor data from the machines that are under test and from the agent machines.

This section identifies the results of the performance test runs.

Test Series A – Vary Users

In this test series, we vary the number of users loaded onto the test farm. Figure 8 shows the requests per second that the Visual Studio 2010 Ultimate Test Controller was able to process through the SharePoint farm during the tests for each of the user load sizes. You can see that as additional user load is applied, the requests go up due to the larger user number, but when it gets to 15,000, we are heavily loading the farm so it does not increase as much as the load applied.

Because the 15,000 user test took additional time to ramp up, we ran this test for 2 hours instead of the baseline of 1 hour. Due to the load, we also found that 3 front-end Web servers were not sufficient. We ran this test with 6 front-end Web servers.




Figure 8 - Average RPS for series A

In Figure 9 you can see that test transaction response time goes up along with the page refresh time for the large 15,000 user test. This shows that there is a bottleneck in the system for this large user load. We experienced high IOPS load on the H: drive which contains the primary data file for the content database during this test. Additional investigation could have been done in this area to remove this bottleneck.

¹ [Visual Studio Agents 2010](#)




Figure 9 - Times and WFE's Used for Series A

In Figure 10 you can see the increasing CPU use as we move from 4,000 to 10,000 user load, and then you can see the reduced CPU use just for the front-end Web servers (WFEs) as we double the number of them from 3 to 6. At the bottom, you can see the APP-1 server has fairly constant CPU use, and the large PACNEC01 SQL Server computer does not get to 3% of total CPU use.

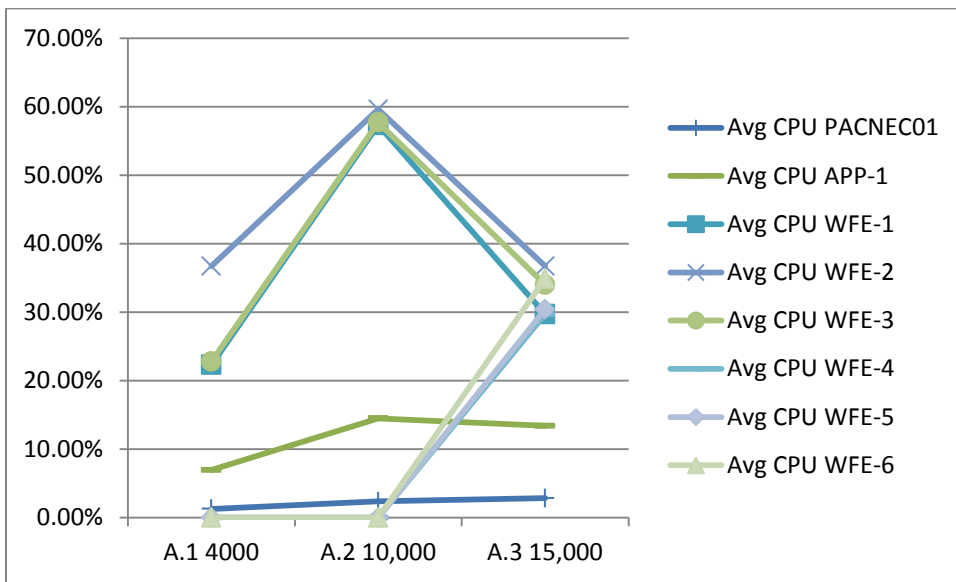


Figure 10 - Average CPU Use for Series A

Table 12 shows a summary of data captured during the three tests in test series A. Data items that show “NA” were not captured.

Test	A.1	A.2	A.3
Users	4,000	10,000	15,000
WFEs	3	3	6
Duration	1 hour	1 hour	2 hours
Avg RPS	96.3	203	220

Avg Page Time	0.31 sec	0.71 sec	19.2 sec
Avg Response Time	0.26 sec	0.58 sec	13.2 sec
Avg CPU WFE-1	22.3%	57.3%	29.7%
Available RAM WFE-1	5,828	5,786	13,311
Avg CPU WFE-2	36.7%	59.6%	36.7%
Available RAM WFE-2	5,651	5,552	13,323
Avg CPU WFE-3	22.8%	57.7%	34%
Available RAM WFE-3	5,961	5,769	13,337
Avg CPU PACNEC01	1.29%	2.37%	2.86%
Available RAM PACNEC01	401,301	400,059	876,154
Avg CPU APP-1	6.96%	14.5%	13.4%
Available RAM APP-1	13,745	13,804	13,311
Avg CPU APP-2	0.73%	1.09%	0.27%
Available RAM APP-2	14,815	14,992	13,919
Avg CPU WFE-4	NA	NA	29.7%
Available RAM WFE-4	NA	NA	13,397
Avg CPU WFE-5	NA	NA	30.4%
Available RAM WFE-5	NA	NA	13,567
Avg CPU WFE-6	NA	NA	34.9%
Available RAM WFE-6	NA	NA	13,446
Avg Disk Write Queue Length, PACNEC01 H: SPCContent DB1	0.0 (with peak of 0.01)	0.0 (with peak of 0.02)	0.3 (with peak of 24.1)

Table 12 - Detailed Results from Series A Testing

Test Series B – Vary SQL Server RAM

In this test series we vary the amount of RAM available to SQL Server. You can see in Figure 11 that the requests-per-second was not impacted by the RAM allocated to SQL Server.




Figure 11 – Average Requests per Second for series B

In Figure 12 you can see that all tests had page and transaction response times under 1 second.




Figure 12 - Page and transaction response times for series B

Figure 13 shows the CPU use for the front-end Web servers (WFE), the App Server, and the SQL Database Server. You can see that the 3 WFEs were constantly busy for all tests, the App Server is mostly idle, and the database server does not get above 3% CPU.




Figure 13 - Average CPU Use for series B

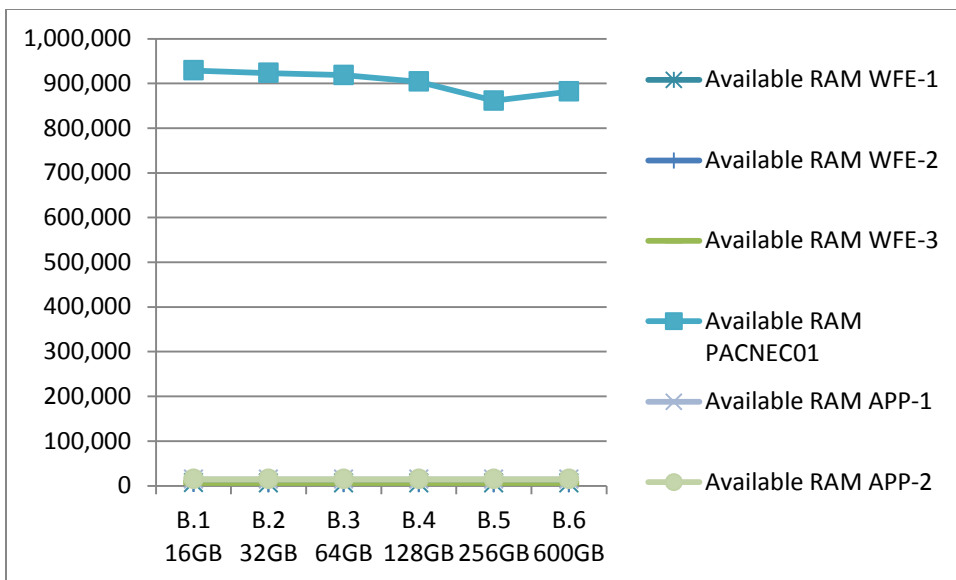


Figure 14 - Available RAM for series B

Table 13 shows summary of data captured during the three tests in test series B.

Test	B.1	B.2	B.3	B.4	B.5	B.6
SQL RAM	16 GB	32 GB	64 GB	128 GB	256 GB	600 GB
Avg RPS	203	203	203	204	203	202
Avg Page Time	0.66	0.40	0.38	0.42	0.58	0.89
Avg Response Time	0.56	0.33	0.31	0.37	0.46	0.72
Avg CPU WFE-1	57.1%	58.4%	58.8%	60.6%	60%	59%
Available	6,239	6,063	6,094	5,908	5,978	5,848

RAM WFE-1						
Avg CPU	55.6%	60.1%	57.1%	59.6%	60.3%	58.1%
WFE-2						
Available RAM	6,184	6,079	6,141	6,119	5,956	5,828
WFE-2						
Avg CPU	59.4%	56%	56.9%	58.4%	61.4%	59.8%
WFE-3						
Available RAM	6,144	6,128	6,159	6,048	5,926	5,841
WFE-3						
Avg CPU	2.84%	2.11%	2.36%	2.25%	2.38%	2.29%
PACNEC01						
Available RAM	928,946	923,332	918,526	904,074	861,217	881,729
PACNEC01						
Avg CPU	14.3%	12.6%	13.3%	12.5%	13.4%	13.8%
APP-1						
Available RAM	14,163	14,099	14,106	14,125	14,221	14,268
APP-1						
Avg CPU	1.29%	1.14%	1.2%	1.2%	1.03%	0.96%
APP-2						
Available RAM	15,013	14,884	14,907	14,888	14,913	14,900
APP-2						

Table 13 - Detailed Results from Series B

Test Series C – Vary Transaction Mix

In this test series, we vary the proportion of search transactions done in the workload mix.




Figure 15 - Average RPS for series C

In Figure 16 you can see that test C.5 had significantly longer page response times, which indicates that the SharePoint Server 2010 and FAST Search Server 2010 for SharePoint farm was overloaded during this test.

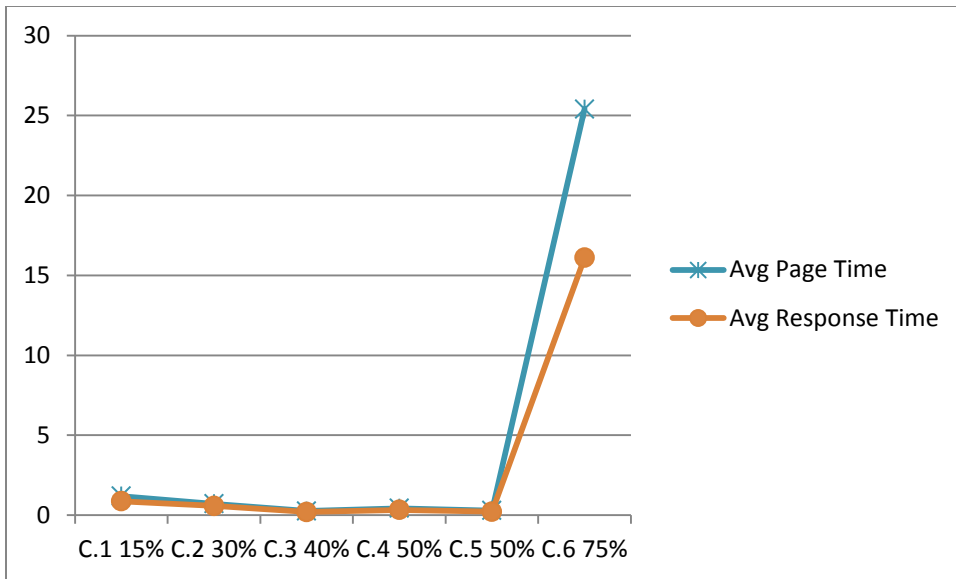


Figure 16 - Page and transaction response times for series C

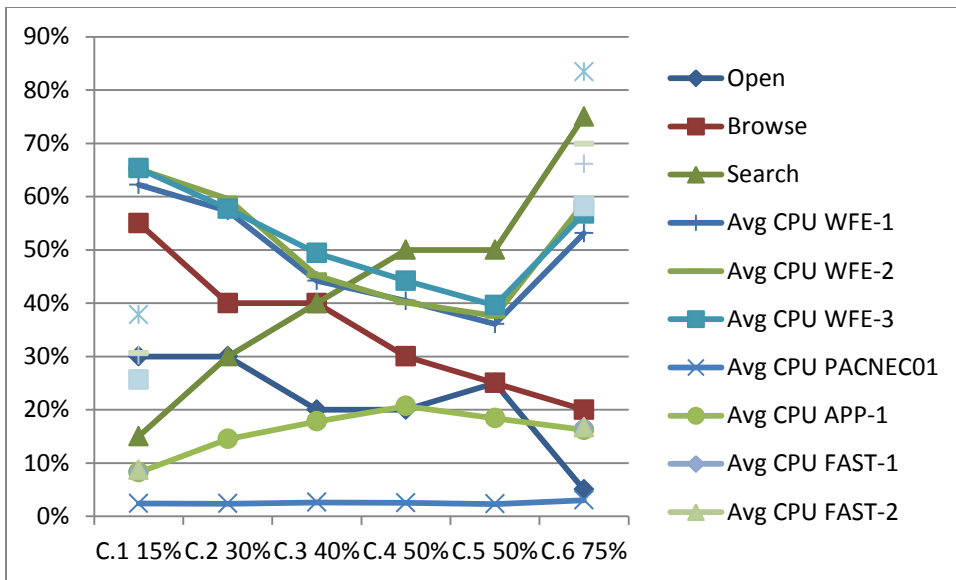


Figure 17 - Average CPU Time for series C




Figure 18 - Average RAM for series C

Table 14 shows a summary of data captured during the three tests in test series C.

Test	C.4	C.2 (baseline)	C.1	C.2	C.3	C.5
Open	30%	30%	20%	20%	25%	5%
Browse	55%	40%	40%	30%	25%	20%
Search	15%	30%	40%	50%	50%	75%
Avg RPS	235	203	190	175	168	141
Avg Page Time (S)	1.19	0.71	0.26	0.43	0.29	25.4
Avg Response Time (S)	0.87	0.58	0.20	0.33	0.22	16.1
Avg CPU WFE-1	62.2%	57.30%	44.2%	40.4%	36.1%	53.1%
Available RAM WFE-1	14,091	5,786	6,281	6,162	6,069	13,766
Avg CPU WFE-2	65.2%	59.60%	45.2%	40.1%	37.6%	58.8%
Available RAM WFE-2	13,944	5,552	6,271	6,123	6,044	13,726
Avg CPU WFE-3	65.3%	57.70%	49.4%	44.2%	39.6%	56.8%
Available RAM WFE-3	13,693	5,769	6,285	6,170	6,076	13,716
Avg CPU PACNEC01	2.4%	2.37%	2.6%	2.51%	2.32%	3.03%
Available RAM PACNEC01	899,613	400,059	814,485	812,027	808,842	875,890
Avg CPU	8.27%	14.50%	17.8%	20.7%	18.4%	16.2%

APP-1						
Available RAM APP-1	13,687	13,804	14,002	13,991	13,984	13,413
Avg CPU APP-2	0.28%	NA	0.88%	0.8%	0.79%	0.14%
APP-2						
Available RAM APP-2	13,916	NA	14,839	14,837	14,833	13,910
Avg CPU FAST-1	8.39%	NA	NA	NA	NA	16.6%
Available RAM FAST-1	13,998	NA	NA	NA	NA	13,686
Avg CPU FAST-2	8.67%	NA	NA	NA	NA	16.7%
Available RAM FAST-2	14,135	NA	NA	NA	NA	13,837
Avg CPU FAST-IS1	37.8%	NA	NA	NA	NA	83.4%
Available RAM FAST-IS1	2,309	NA	NA	NA	NA	2,298
Avg CPU FAST-IS2	30.2%	NA	NA	NA	NA	66.1%
Available RAM FAST-IS2	5,162	NA	NA	NA	NA	5,157
Avg CPU FAST-IS3	30.6%	NA	NA	NA	NA	69.9%
Available RAM FAST-IS3	5,072	NA	NA	NA	NA	5,066
Avg CPU FAST-IS4	25.6%	NA	NA	NA	NA	58.2%
Available RAM FAST IS-4	5,243	NA	NA	NA	NA	5,234

Table 14 - Detailed Results from Series C Testing

Test Series D – Vary Front-End Web Server RAM

In this test series, we vary the amount of RAM on each front-end Web server virtual machine.




Figure 19 - Average RPS

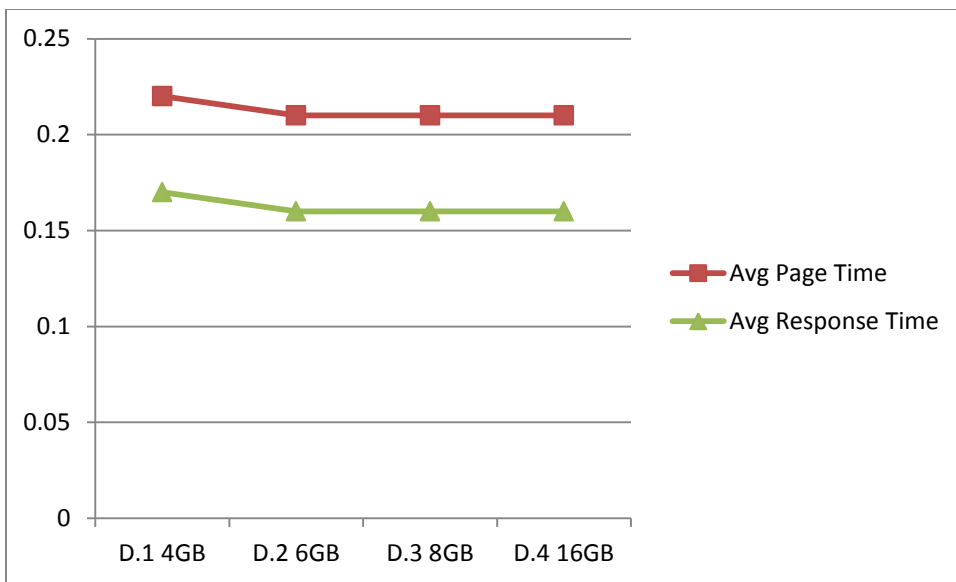


Figure 20 - Page and transaction response time




Figure 21 - Average CPU times

In Figure 22 you can see that the available RAM on each front-end Web server in all cases is the RAM allocated to the virtual machine less 2 GB. This shows that for the 10,000 user load and this test transaction mix, the front-end Web servers require a minimum of 2 GB of RAM plus any reserve.




Figure 22 - Available RAM for series D

Table 15 shows a summary of data captured during the three tests in test series D.

Test	D.1	D.2	D.3	D.4
WFE RAM	4 GB	6 GB	8 GB	16 GB
Avg RPS	189	188	188	188
Avg Page Time (S)	0.22	0.21	0.21	0.21
Avg Response	0.17	0.16	0.16	0.16

Time (S)				
Avg CPU WFE-1	40.5%	37.9%	39.6%	37.3%
Available RAM WFE-1	2,414	4,366	6,363	14,133
Avg CPU WFE-2	42.3%	40%	40.3%	39.5%
Available RAM WFE-2	2,469	4,356	6,415	14,158
Avg CPU WFE-3	42.6%	42.4%	42.2%	43.3%
Available RAM WFE-3	2,466	4,392	6,350	14,176
Avg CPU PACNEC01	2.04%	1.93%	2.03%	2.14%
Available RAM PACNEC01	706,403	708,725	711,751	706,281
Avg CPU APP-1	11.8%	13.1%	12.9%	12.3%
Available RAM APP-1	13,862	13,866	13,878	13,841
Avg CPU APP-2	0.84%	0.87%	0.81%	0.87%
Available RAM APP-2	14,646	14,650	14,655	14,636
Avg CPU WFE-4	42.3%	43.6%	41.9%	45%
Available RAM WFE-4	2,425	4,342	6,382	14,192

Table 15 - Detailed Results from Series D Testing

Test Series E – Vary Number of Front-End Web Servers

In this test series we vary the number of front-end Web servers in the farm. Notice that Figure 23 shows the average RPS slightly lower with 2 and 3 front-end Web servers as the system does not quite keep up with the applied user load. But see that with 4, 5 or 6 front-end Web servers, requests-per-second is constant as the system is handling the full load from the test agents.

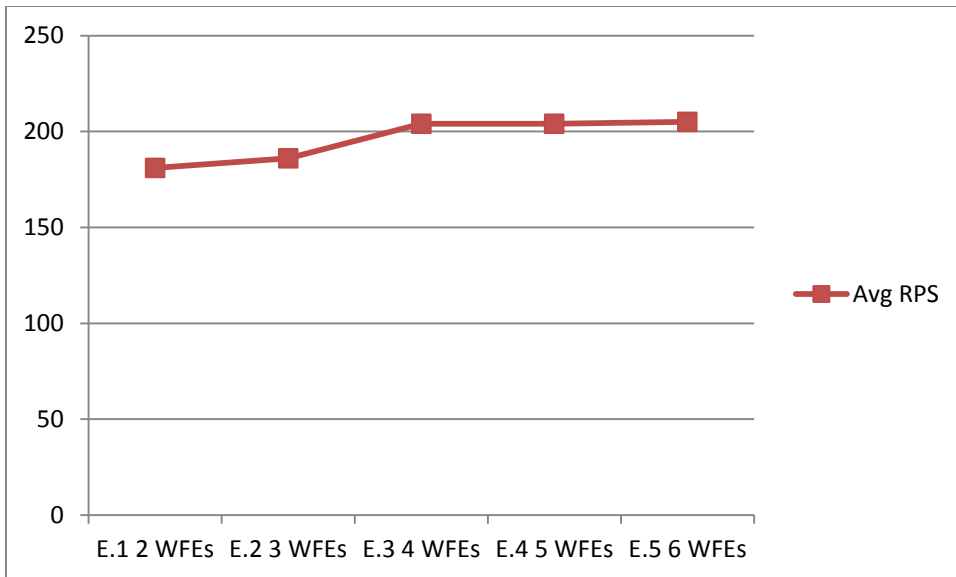


Figure 23 - Average RPS for Series E

A similar pattern is shown in Figure 24 where you can see the response times high for 2 and 3 WFEs and then very low for the higher numbers of front-end Web servers.




Figure 24 - Page and transaction time for Series E

In Figure 25 you can see the CPU time is lower when more front-end Web servers are available. Using 6 front-end Web servers clearly reduces the average CPU utilization across the front-end Web servers, but only 4 front-end Web servers are required for the 10,000 user load. Notice that you cannot tell from this chart which configurations are handling the load and which ones are not. See that for 3 front-end Web servers which we identified as not completely handling the load, the front-end Web server CPU is just over 50%.




Figure 25 - Average CPU for Series E

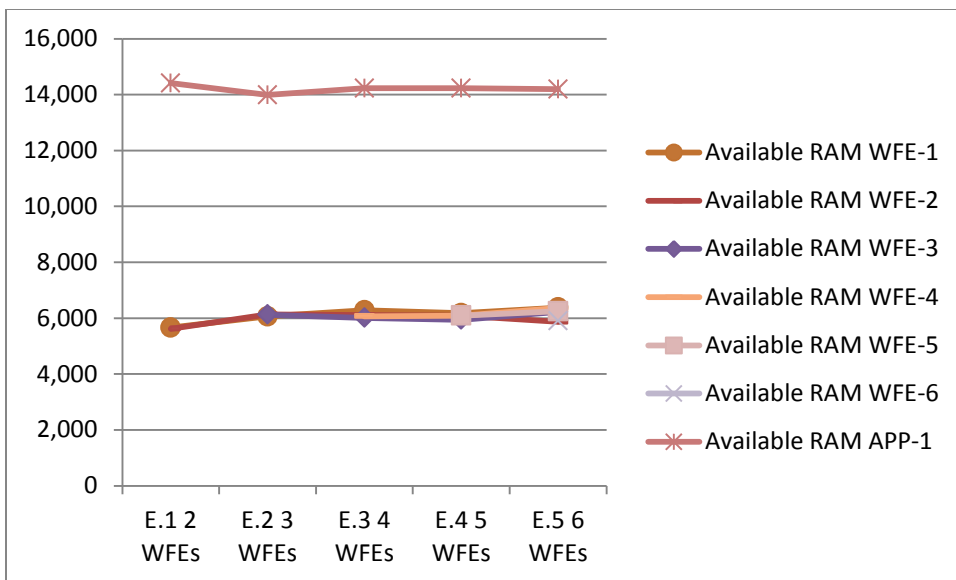


Figure 26 - Available RAM for Series E

Table 16 shows a summary of data captured during the three tests in test series E.

Test	E.1	E.2	E.3	E.4	E.5
WFE Servers	2	3	4	5	6
Avg RPS	181	186	204	204	205
Avg Page Time (S)	8.02	0.73	0.23	0.20	0.22
Avg Response Time (S)	6.34	0.56	0.19	0.17	0.18
Avg CPU WFE-1	77.4	53.8	45.7	39.2	32.2
Available	5,659	6,063	6,280	6,177	6,376

RAM WFE-1					
Avg CPU WFE-2	76.2%	53.8%	45.9%	38.2%	28.8%
Available RAM WFE-2	5,623	6,132	6,105	6,089	5,869
Avg CPU WFE-3	NA	52.5%	43.9%	37.7%	31.2%
Available RAM WFE-3	NA	6,124	6,008	5,940	6,227
Avg CPU WFE-4	NA	NA	44.5%	34.8%	34.7%
Available RAM WFE-4	NA	NA	6,068	6,083	6,359
Avg CPU WFE-5	NA	NA	NA	35.1%	32%
Available RAM WFE-5	NA	NA	NA	6,090	6,245
Avg CPU WFE-6	NA	NA	NA	NA	33.9%
Available RAM WFE-6	NA	NA	NA	NA	5,893
Avg CPU PACNEC01	2.13%	1.93%	2.54%	2.48%	2.5%
Available RAM PACNEC01	899,970	815,502	397,803	397,960	397,557
Avg CPU APP-1	9.77%	11.7%	15%	14.7%	13.6%
Available RAM APP-1	14,412	13,990	14,230	14,227	14,191
Avg CPU APP-2	1.06%	0.92%	1%	1%	1.04%
Available RAM APP-2	14,928	14,841	14,874	14,879	14,869

Table 16 - Detailed Results from Series E Testing

Test Series F – Vary SQL Server CPUs

In this test series we vary the number of CPUs available to SQL Server.

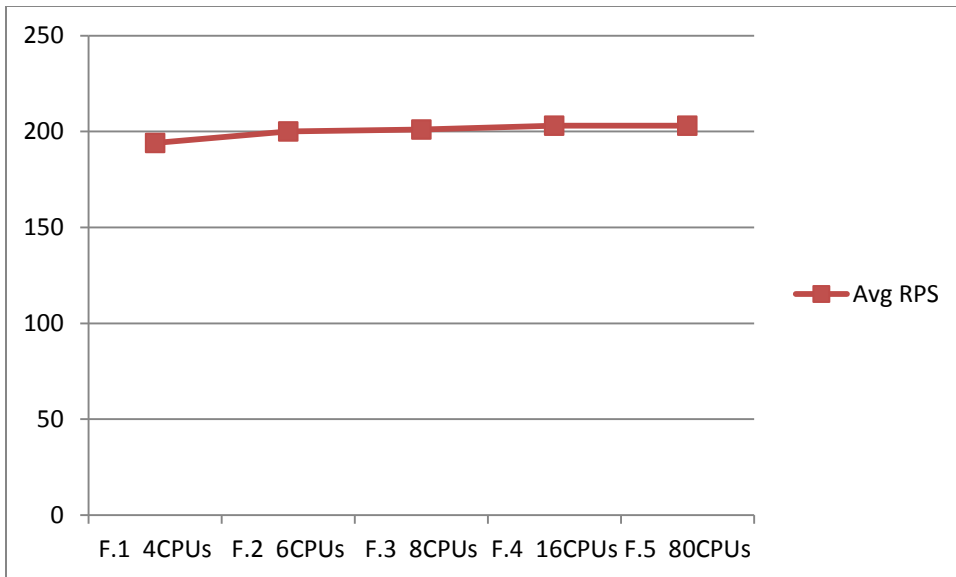


Figure 27 - Average RPS for series F

You can see in Figure 28 that despite minimal CPU use on the SQL Server computer, the page and transaction response times go up when SQL Server has fewer CPUs available to work with.

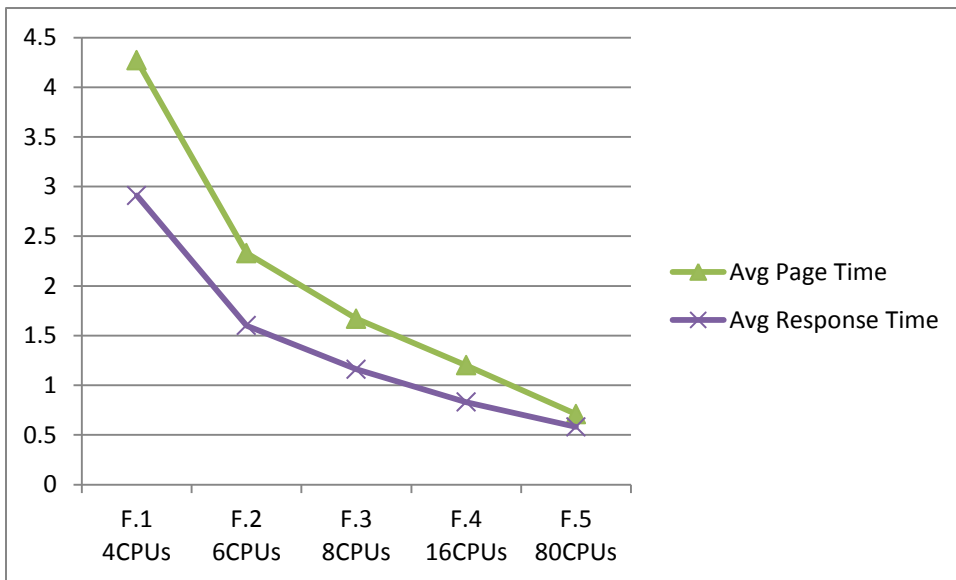


Figure 28 - Page and transaction time for series F

In Figure 29 the SQL Server average CPU for the whole computer does not go over 3%. The three front-end Web servers are all around 55% throughout the tests.

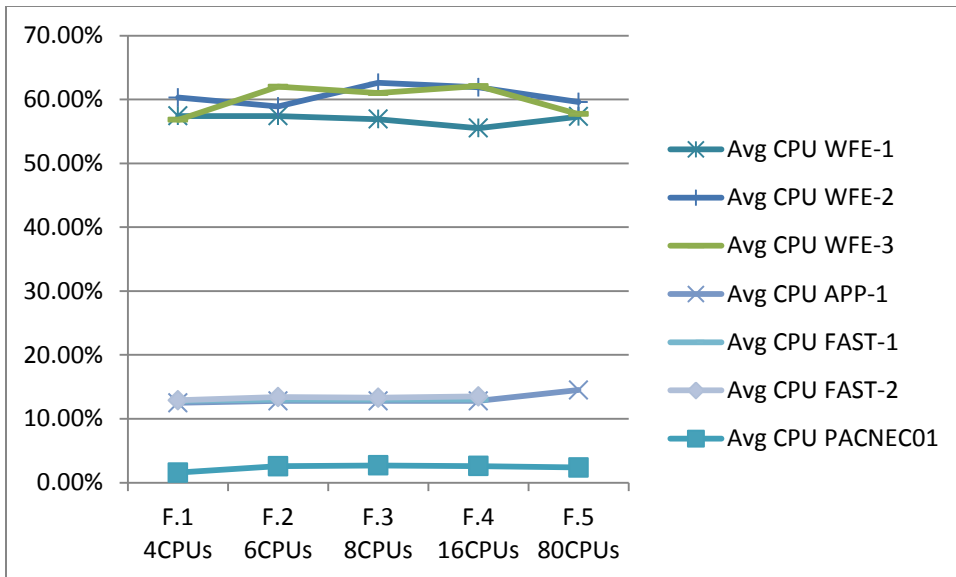


Figure 29 - Average CPU for series F

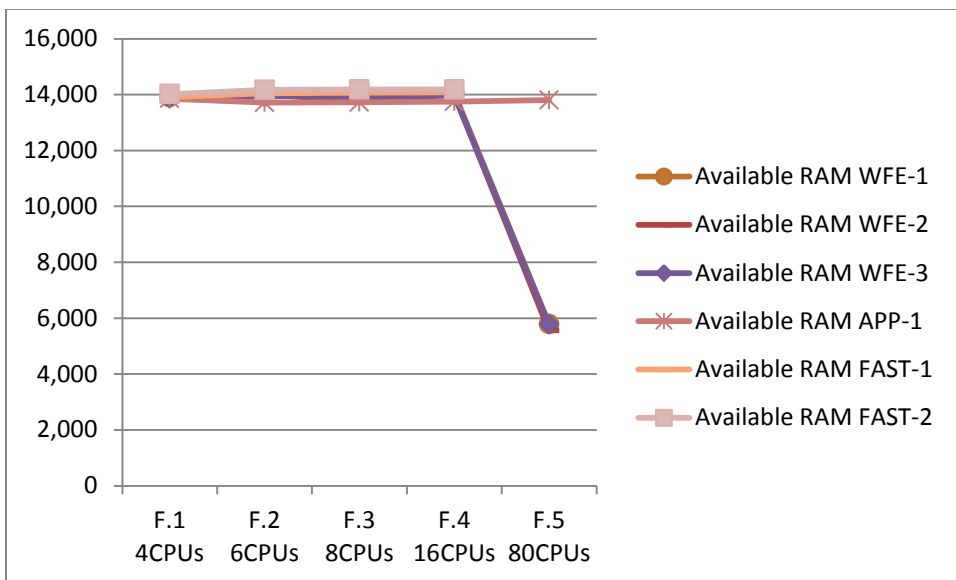


Figure 30 - Available RAM for series F

Table 17 shows a summary of data captured during the three tests in test series F.

Test	F.1	F.2	F.3	F.4	F.5
SQL CPUs	4	6	8	16	80
Avg RPS	194	200	201	203	203
Avg Page Time (S)	4.27	2.33	1.67	1.2	0.71
Avg Response Time (S)	2.91	1.6	1.16	0.83	0.58
Avg CPU WFE-1	57.4%	57.4%	56.9%	55.5%	57.30%
Available	13,901	13,939	13,979	14,045	5,786

RAM WFE-1					
Avg CPU	60.3%	58.9%	62.6%	61.9%	59.60%
WFE-2					
Available	13,920	14,017	13,758	14,004	5,552
RAM WFE-2					
Avg CPU	56.8%	62%	61%	62.1%	57.70%
WFE-3					
Available	13,859	13,942	13,950	13,971	5,769
RAM WFE-3					
Avg CPU	1.56%	2.57%	2.69%	2.6%	2.37%
PACNEC01					
Available	865,892	884,642	901,247	889,479	400,059
RAM					
PACNEC01					
Avg CPU	12.5%	12.8%	12.8%	12.8%	14.50%
APP-1					
Available	13,856	13,713	13,725	13,745	13,804
RAM APP-1					
Avg CPU	0.22%	0.25%	0.26%	0.25%	NA
APP-2					
Available	14,290	14,041	14,013	13,984	NA
RAM APP-2					
Avg CPU	12.8%	13%	13%	13%	NA
FAST-1					
Available	13,913	14,051	14,067	14,085	NA
RAM FAST-1					
Avg CPU	12.9%	13.4%	13.3%	13.5%	NA
FAST-2					
Available	14,017	14,170	14,183	14,184	NA
RAM FAST-2					

Table 17 - Detailed Results from Series F Testing

Service Pack 1 (SP1) and June Cumulative Update (CU) Test

After the SharePoint Server 2010 farm was fully populated with 120 million items, we applied SharePoint Server 2010 SP1 and FAST Search Server 2010 for SharePoint SP1 to see how long the process would take on a large populated farm.

SharePoint Server 2010

SharePoint Server 2010 Service Pack 1 (SP1) and the June Cumulative Update were applied in the lab to determine a base upgrade time for a large-scale Document Center farm scenario. The following chart reflects the servers in the farm that required the SP1 and June CU upgrades, the start and end time of each install, total time of installs, start and end time of PSCONFIG upgrade command, total time of PSCONFIG upgrade command, and total time of upgrade by server name, and total installation times.

Server Name	SP1 Start	SP1 End	Diff (h:mm:ss)	June CU Start	June CU End	Diff (h:mm:ss)	PSCONFIG Start	PSCONFIG End	Diff (h:mm:ss)
APP-1	7/12/2011 4:00:00	7/12/2011 4:15:51	0:15:51	7/29/2011 10:45:00	7/29/2011 11:00:05	0:15:05	7/29/2011 13:25:50	7/29/2011 13:30:15	0:04:25

APP-2	7/12/2011 4:26:07	7/12/2011 4:39:31	0:13:24	7/29/2011 11:02:30	7/29/2011 11:17:23	0:14:53	7/29/2011 7/29/2011 13:33:15 13:35:11	0:01:56
WFE-1	7/12/2011 4:41:05	7/12/2011 4:49:16	0:08:11	7/29/2011 11:23:00	7/29/2011 11:31:07	0:08:07	7/29/2011 7/29/2011 13:36:35 13:38:11	0:01:36
WFE-2	7/12/2011 4:50:24	7/12/2011 4:57:47	0:07:23	7/29/2011 11:32:45	7/29/2011 11:40:46	0:08:01	7/29/2011 7/29/2011 13:39:20 13:40:54	0:01:34
WFE-3	7/12/2011 4:59:00	7/12/2011 5:06:39	0:07:39	7/29/2011 11:42:00	7/29/2011 11:49:47	0:07:47	7/29/2011 7/29/2011 13:42:40 13:44:14	0:01:34
WFE-4	7/12/2011 5:10:06	7/12/2011 5:17:30	0:07:24	7/29/2011 11:51:00	7/29/2011 11:58:49	0:07:49	7/29/2011 7/29/2011 13:46:05 13:47:41	0:01:36
WFE-5	7/12/2011 5:18:49	7/12/2011 5:27:07	0:08:18	7/29/2011 11:59:45	7/29/2011 12:08:19	0:08:34	7/29/2011 7/29/2011 13:49:00 13:50:36	0:01:36
WFE-6	7/12/2011 5:28:25	7/12/2011 5:35:40	0:07:15	7/29/2011 12:09:30	7/29/2011 12:17:10	0:07:40	7/29/2011 7/29/2011 13:52:00 13:53:35	0:01:35
WFE-CRAWL1	7/12/2011 5:37:20	7/12/2011 5:44:35	0:07:15	7/29/2011 12:18:10	7/29/2011 12:25:51	0:07:41	7/29/2011 7/29/2011 13:54:35 13:56:19	0:01:44
FAST-SSA-1	7/12/2011 5:49:00	7/12/2011 5:57:45	0:08:45	7/29/2011 12:39:40	7/29/2011 12:48:24	0:08:44	7/29/2011 7/29/2011 13:57:30 13:59:07	0:01:37
FAST-SSA-2	7/12/2011 5:59:08	7/12/2011 6:08:29	0:09:21	7/29/2011 12:51:30	7/29/2011 13:00:11	0:08:41	7/29/2011 7/29/2011 14:00:00 14:01:58	0:01:58
Total Time:			1:40:46			1:43:02		0:21:11
Grand Total:			3:44:59					

Table 18- Times to apply SP1 and June Cumulative Updates

FAST Search Server for SharePoint 2010

The FAST Search Server for SharePoint 2010 SP1 upgrade took approximately 15 minutes per node to upgrade.

SQL Server Content DB Backups

Document Center 1

A SQL Server database backup was executed on the content database for Document Center 1 (SPContent01). A backup (B/U) was performed pre-SP1, June Cumulative Update (CU), and post-SP1. Backup time and size details follow.

Database Name	B/U Start	B/U End	Diff (h:mm:ss)	Size (TB)	Notes
SPContent01	7/10/2011 9:56:00	7/10/2011 23:37:00	13:41:00	14.40	Pre-SP1
SPContent01	7/29/2011 14:22:10	7/30/2011 4:28:00	14:05:50	14.40	Post- SP1 / June CU

Table 19 - Time to Run Backups

Conclusions

The SharePoint Server 2010 farm was successfully tested at 15,000 concurrent users with two SharePoint content databases totaling 120 million documents. We were not able to support the load of 15,000 concurrent users with three front-end Web servers as was specified in the baseline environment and needed six front-end Web servers for this load.

Recommendations

A summary list of the recommendations follows. A forthcoming large scale document library best practices document is planned for more detail on each of these recommendations. In each section the hardware notes are not intended to be a comprehensive list, but rather indicate the minimum hardware that was found to be required for the 15,000 concurrent user load test against a 120 million document SharePoint Server 2010 farm.

Recommendations related to SQL Server 2008 R2

- Hardware notes for load:
 - 64 GB RAM on SQL Server
 - 16 CPU cores on SQL Server
- Provide adequate 2 IO Capability per Second per GB stored in the SharePoint content database
- Set SQL Server 2008 R2 property Maximum Degree of Parallelism (MAXDOP)=1; the default is 0
- Use multiple LUNs (drive letters) on SAN each with a SQL Server data file and one virtual CPU allocated for each LUN used. We used 5 data files all on separate LUNs

Recommendations related to SharePoint Server 2010

- Hardware notes for load:
 - 8 GB RAM on each front-end Web server
 - 6 front-end Web servers
- Add the Disable Loopback Check Registry Key at
`\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Lsa\DisableLoopbackCheck=1`
- Reduce table index fragmentation issues manually during bulk document import by running ALTER INDEX on the affected table indexes.
- Prefer SPFileCollection.ADD for bulk import of documents over creating duplicate documents with SPFolder.CopyTo.

Recommendations related to FAST Search Server for SharePoint 2010

- Hardware notes for load:
 - 4 rows of FAST Search Server 2010 for SharePoint index Servers
- Registry Updates for SharePoint Server 2010 document crawler

On nodes running the FAST Content SSA crawler (*APP-1* and *APP-2*), the following registry values were updated to improve the crawler performance in the hive:

HKLM\SOFTWARE\Microsoft\Office Server\14.0\Search\Global\Gathering Manager

1. FilterProcessMemoryQuota

Default value of 100 megabytes (MB) was changed to 200 MB

2. DedicatedFilterProcessMemoryQuota

Default value of 100 megabytes (MB) was changed to 200 MB

3. FolderHighPriority

Default value of 50 was changed to 500

- Monitor the FAST Search Server 2010 for SharePoint Index Crawl

The crawler should be monitored at least three times per day. 100 million items took us about 2 weeks to crawl. Each time monitoring the crawl the following four checks were done:

1. rc -r | select-string "# doc"

Checks how busy the document processors are

2. Monitoring crawl queue size

Use reporting or SQL Server Management Studio to see MSCrawlURL

3. Indexerinfo -a doccount

Make sure all indexers are reporting to see how many are indexed in 1000 milliseconds. We saw this run from 40 to 120 depending on the type of documents being indexed at the time.

4. Indexerinfo -a status

Monitor the health of the indexers and partition layout

References

- [SharePoint Server 2010 capacity management: Software boundaries and limits](http://technet.microsoft.com/en-us/library/cc262787.aspx) (http://technet.microsoft.com/en-us/library/cc262787.aspx)
- [Estimate performance and capacity requirements for large scale document repositories in SharePoint Server 2010](http://technet.microsoft.com/en-us/library/hh395916.aspx) (http://technet.microsoft.com/en-us/library/hh395916.aspx)
- [Storage and SQL Server capacity planning and configuration \(SharePoint Server 2010\)](http://technet.microsoft.com/en-us/library/cc298801.aspx) (http://technet.microsoft.com/en-us/library/cc298801.aspx)
- [SharePoint Performance and Capacity Planning Resource Center on TechNet](http://technet.microsoft.com/en-us/office/sharepointserver/bb736741) (http://technet.microsoft.com/en-us/office/sharepointserver/bb736741)
- [Best practices for virtualization \(SharePoint Server 2010\)](http://technet.microsoft.com/en-us/library/hh295699.aspx) (http://technet.microsoft.com/en-us/library/hh295699.aspx)
- [Best practices for SQL Server 2008 in a SharePoint Server 2010 farm](http://technet.microsoft.com/en-us/library/hh292622.aspx) (http://technet.microsoft.com/en-us/library/hh292622.aspx)
- [Best practices for capacity management for SharePoint Server 2010](http://technet.microsoft.com/en-us/library/hh403882.aspx) (http://technet.microsoft.com/en-us/library/hh403882.aspx)
- [Performance and Capacity Recommendations for FAST Search Server 2010 for SharePoint](http://technet.microsoft.com/en-us/library/gg702613.aspx) (http://technet.microsoft.com/en-us/library/gg702613.aspx)
- [Bulk Loader tool](http://code.msdn.microsoft.com/Bulk-Loader-Create-Unique-eeb2d084) (http://code.msdn.microsoft.com/Bulk-Loader-Create-Unique-eeb2d084)
- [LoadBulk2SP tool](http://code.msdn.microsoft.com/Load-Bulk-Content-to-3f379974) (http://code.msdn.microsoft.com/Load-Bulk-Content-to-3f379974)
- [SharePoint Performance Testing Scripts](http://code.msdn.microsoft.com/SharePoint-Testing-c621ae38) (http://code.msdn.microsoft.com/SharePoint-Testing-c621ae38)

