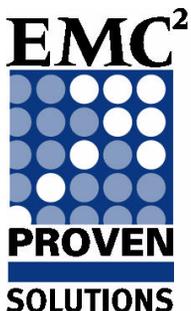




EMC Virtual Architecture for Microsoft SharePoint Server 2007

Enabled by EMC CLARiiON CX3-40, VMware ESX Server 3.5 and
Microsoft SQL Server 2005

Reference Architecture



EMC Global Solutions Operations

EMC Corporation
Corporate Headquarters
Hopkinton MA 01748-9103
1.508.435.1000
www.EMC.com

Copyright © 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part number: H4456

Contents

About this Solution

Purpose	4
The business challenge	4
The technology solution	5
Solution details	6
Environment profile	7
Conclusion	12

About this Solution

Purpose

The purpose of this solution is to demonstrate the functional, performance and scalability¹ capabilities of a virtualized Enterprise Microsoft SharePoint 2007 farm². In the solution, EMC's CLARiiON® CX3-40f array is used for storage and consolidation, while Microsoft SQL Server 2005 is used as the relational database management system that supports the various SharePoint 2007 content data types.

VMware High Availability (VMHA) provides high availability of all the applications running on the virtual machines (VMs), while VMware VMotion enables the live migration of the VMs across the physical ESX infrastructure servers with zero downtime. In addition, VMware Distributed Resource Scheduler (DRS) continuously monitors utilization and ensures particular VMs are allocated resources based on pre-defined rules. VMware DRS also enables live migration of VMs (when constrained by physical hardware resources) to different physical ESX Servers.

The business challenge

Organizations are creating huge volumes of unstructured content on a daily basis. This unstructured content varies from documents, e-mail, video files, and Web pages. The content is often in an unmanaged state, which can lead to inefficient information sharing and reduced employee productivity.

Portal sites are used to manage content and connect employees to business critical information, expertise, and applications. Microsoft's Office SharePoint Server is a world-class Enterprise portal platform that makes it easy to build and maintain portal sites for every aspect of the business.

-
1. Based on various user profiles.
 2. A farm is a collection of computer servers required to accomplish server needs far beyond the capability of one machine.

By using portals, organizations can streamline processes and transactions, increase employee productivity, and strengthen relationships with their customers and partners. This solution illustrates how organizations can:

- Manage diverse content and streamline business processes.
- Enable affordable uniform high availability across an entire virtualized environment.
- Ensure a more effective, streamlined method of handling growing data volumes.
- Minimize un-planned downtime and IT service disruption, as well as limit impact on server resources and networks.

The technology solution

This solution describes a virtualized Enterprise Microsoft SharePoint 2007 farm environment that can serve a particular user count³ (based on three user profiles). The environment consists of two clusters - an MSCS physical cluster and a VMHA cluster. The MSCS physical cluster is a two-node (active / passive) SQL cluster. The VMHA cluster is a three-node (active / active / active) cluster that contains the entire infrastructure required to operate a SharePoint 2007 farm, for example domain controllers, application servers, and Web front ends (WFE).

Microsoft SharePoint Server 2007 is used to link to Microsoft SQL Server 2005 to allow easy access to data such as documents and reports. EMC's CLARiiON CX3-40f is used to store and retrieve the required data (as well as provide storage for the VMs), while VMware is used to illustrate how it is possible to consolidate and virtualize large amounts of user data on the high-performance, highly available storage platform. The solution provides information on:

- Creating a well-performing storage design for a virtualized Microsoft SharePoint 2007 farm on an EMC[®] CLARiiON CX3-40f with a large and very active database.
- Documenting the process of installing and running the SharePoint 2007 farm on VMware machines.
- Documenting the performance of the SharePoint 2007 farm when running against a heavy user load, and documenting failover and failback procedures after hardware failures.

3. The environment supported up to 180,000 users with the limiting factor being server hardware (CPU).

Solution details

EMC CLARiiON CX3-40 — The CLARiiON CX3 UltraScale™ architecture, which is based on a high-performance, high availability design, enables the CX3 UltraScale series to address a broad range of application environments. The CX3 UltraScale series systems are built on the same redundant modular architecture and run the FLARE® storage operating environment.

Disk layout — The solution used the Microsoft maximum recommended sizing for SQL SharePoint 2007 content databases. 1 TB⁴ of data was split into 10 x 100 GB content databases and was spread across 30 x 146 GB disks using RAID 1/0 for redundancy and speed. All VMware VMs used the CX3-40f for operating systems and content drives, except for the SQL cluster, which had local operating system drives.

VMware virtualization technology — In essence, virtualization lets you transform hardware into software. The VMware ESX Server can transform or "virtualize" the hardware resources of an x64-based computer - including the CPU, RAM, hard disk, and network controller - to create a fully functional virtual machine that can run its own operating system and applications just like a physical computer.

VMware layout — All servers in the environment were VMs except for the SQL Server cluster. VMware VMotion was used to move the VMs from one physical server to another, and allowed for the scaling of the servers as required (based on usage and load).

Distributed Resource Scheduler (DRS) allowed for even load distribution of busy VMs across the ESX Servers, maximizing hardware utilization, while minimizing resource bottlenecks. VMHA was used to provide redundancy and failover of the SharePoint 2007 farm components. VMHA and Microsoft cluster services also allowed for full single-site redundancy.

Microsoft SQL Server 2005 — An active / passive MSCS SQL cluster was used for redundancy. The content and temporary (tempdb) databases were manually configured and created (datafile disk space was pre-allocated for optimum performance). Separate data and log LUNs were also used.

4. In total, 1 TB (equating to 4.6 million documents) of data was used in testing the environment.

Microsoft SharePoint 2007 — As crawling is a CPU-intensive activity the Index server was configured with four CPUs. To improve the crawl speed, the performance level for the Index server was set to "maximum" and incremental crawls were set to 15 minutes. In addition, every WFE server was a Query server for search purposes.

As the WFEs are CPU-intensive, 10 VMs were configured and dedicated to hosting the portal, each with four CPUs. The Excel calculation server was configured as the Central Administration, and there was one application server used for document conversions.

Environment profile

The environment contained one CX3-40f storage array. Ten SharePoint sites were populated with data in one site collection.

The SharePoint content database datafile LUN was a 200 GB volume / 100 GB data file, while the SharePoint content database log LUN was a 25 GB volume / 23 GB log file. User concurrency⁵ was 1% (0% think time⁶).

-
5. The percentage of users that are actively using the system. The Microsoft baseline is 1% concurrency.
 6. "0% think time" is the elimination of typical user decision time when browsing, searching, or modifying in SharePoint. For example, a single complete user request is completed from start to finish without user pause, therefore, creating a continuous workload on the system.

Hardware resources

Table 1 details the hardware resources used in the solution.

Table 1 Hardware resources

Equipment	Quantity	Configuration
Storage array	1	CX3-40f Fibre Channel (FC) array with FLARE 24 (version 3.24.40.5.016) <ul style="list-style-type: none"> • 15 x 300 GB FC, RAID 1, 15k rpm production disks (VMware VMFS OS drives) • 30 x 146 GB FC, 15k rpm disks (SQL)
SAN	2	Brocade Silkworm 4900 4 GB FC switches
ESX farm	3	Infrastructure: Dell R900, 16 CPU, 2.4 GHz, 32 GB RAM
Virtual Machines (VMs)	1	Index server (running WFE role dedicated for crawling): 4 CPU, 2.4 GHz, 10 GB RAM
	10	WFEs (running query): 4 CPU, 2.4 GHz, 4 GB RAM
	2	Application servers: 2 CPU, 2.4 GHz, 2 GB RAM
	2	Internal domain controllers: 2 CPU, 2.4 GHz, 2 GB RAM
	2	Public domain controllers: 2 CPU, 2.4 GHz, 2 GB RAM
	1	VMware Virtual Center Server: 2 CPU, 2.4 GHz, 4 GB RAM
HBAs for SAN	10	4 GB Emulex LPe-11002-E
Network switches	2	Cisco Catalyst 3750
SQL Server active node	1	Dell R900, 16 CPU, 2.4 GHz, 32 GB RAM
SQL Server passive node	1	Dell 2950, 8 CPU, 2.4 GHz, 16 GB RAM

Software resources

Table 2 details the software resources used in the solution.

Table 2 Software resources

Software	Configuration
Windows Server 2003 Datacenter Edition (64-bit)	2003 (SP2)
Windows 2003 (64-bit) resource kit	2003 DataCenter Edition x64 (SP2)
Microsoft SQL Server 2005 (SP2 64-bit Enterprise edition)	Used for the cluster (1 active / 1 passive)
Microsoft SharePoint Server 2007	SP1
VMware ESX 3.5	Used for the VMs
VMware Virtual Control Center	ESX 3.5 VirtualCenter 2.5
VMware High Availability (VMHA)	ESX 3.5 VirtualCenter 2.5
VMware VMotion	ESX 3.5 VirtualCenter 2.5

Farm content types

Table 3 lists the content types and sizes used in the SharePoint 2007 farm.

Table 3 Content types

Type	Size (KB)	Type	Size (KB)
.doc	251	.pptx	189
.docx	102	.jpg	93
.xls	820	.gif	75
.xlsx	20	.vsd	471
.ppt	2021	.mpp	235

Farm user load profiles

A Microsoft heavy user load profile was used to determine the maximum user count the SharePoint farm could sustain while ensuring average response times remained within acceptable user limits. As per Microsoft standard, a heavy user performs 60 requests per hour (RPH), that is, one request every 60 seconds.

[Table 4](#) details the acceptable user limits for SharePoint operations.

Table 4 Acceptable user response times

Type of operation	Example	Acceptable user response time
Common	Browse	< 3 seconds
Common	Search	< 3 seconds
Uncommon	Modify	< 5 seconds

Three user profiles were tested to help determine scalability. [Table 5](#) details the user profiles and the response times received.

Table 5 User profiles

User profile (browse / search / modify %)	User load profile (60 RPH)	Requests per sec	Concurrency	Max user capacity	Average user response time (secs) (browse / search / modify)
80 / 10 / 10	Heavy	25	1%	150,000	< 3 / < 3 / < 3
70 / 5 / 25	Heavy	29	1%	174,000	< 3 / < 3 / < 3
50 / 20 / 30	Heavy	30	1%	180,000	< 3 / < 3 / < 3

Physical architecture

Figure 1 illustrates the physical architecture of the solution.

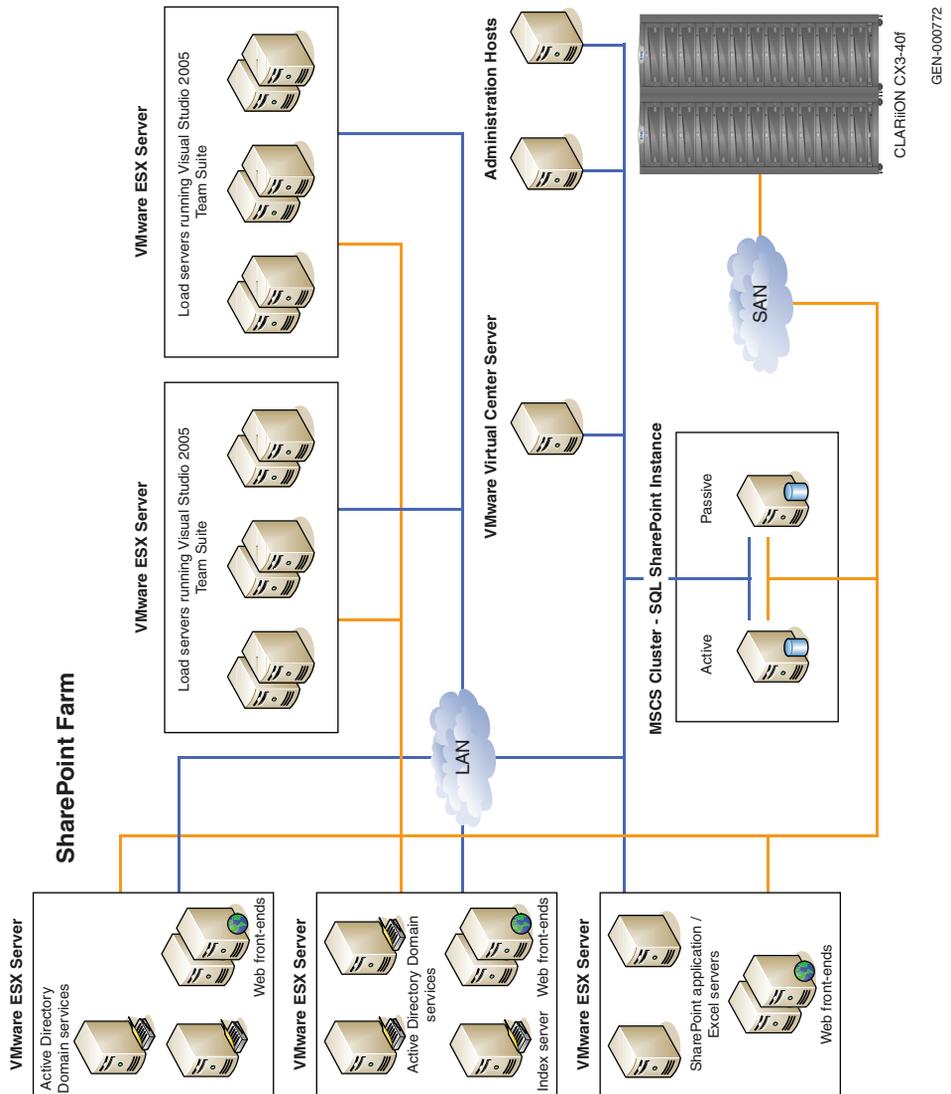


Figure 1 SharePoint 2007 farm environment

Conclusion

This reference architecture depicts a validated virtualized SharePoint 2007 farm that is enabled by SQL Server and VMware technology. The solution utilizes EMC's CLARiiON CX3-40f array for storage and consolidation of the content data types.

Sizing and configuring a SharePoint 2007 farm can be a complex activity as many requirements and aspects must be considered during the planning phase.

Time and attention should be invested to fully gather the current and future requirements of a customer's SharePoint 2007 farm. How a customer intends to utilize the infrastructure, coupled with the nature of the customer's business, will dictate where resources should be spent to eliminate future possible bottlenecks in the environment.

VMware's VMotion and DRS capabilities provide a very valuable proposition in offering greater flexibility by evenly spreading the load and optimizing the utilization of the physical hardware.

The solution provides the following benefits:

- Illustrates real-world expectations for realistic CLARiiON storage requirements and provisioning for a typical SharePoint 2007 farm.
- Illustrates how to design and scale a SharePoint 2007 farm to support a large number of users.
- Reduces the amount of physical servers: maintaining the same cumulative number of CPU cores and memory ensures similar performance levels. In addition, the solution promotes a more eco-friendly environment through the use of virtualization technology.
- Enhances the ability to correctly size and sell a storage and VMware solution based on a given SharePoint 2007 farm configuration.
- Reduces configuration timelines by documenting the key procedural areas.

In addition, it should be noted that the CLARiiON CX3-40f was not stressed during maximum user testing: SP utilization averaged at 20% and LUN utilization averaged at 30%.

EMC can help accelerate assessment, design, implementation, and management while lowering the implementation risks and cost of creating a virtualized SharePoint 2007 farm.

To learn more about this, and other solutions, contact an EMC representative or visit www.EMC.com/solutions/microsoft.