


# EMC Virtual Architecture for Microsoft Office SharePoint Server 2007

Enabled by EMC CLARiiON CX3-40,  
Windows Server 2008, and Hyper-V

## Reference Architecture



**EMC Global Solutions**

**Microsoft** | Virtualization

Copyright © 2008, 2009 EMC Corporation. All rights reserved.

Published March, 2009

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

Benchmark results are highly dependent upon workload, specific application requirements, and system design and implementation. Relative system performance will vary as a result of these and other factors. Therefore, this workload should not be used as a substitute for a specific customer application benchmark when critical capacity planning and/or product evaluation decisions are contemplated.

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly. EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. 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.

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This white paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Microsoft, Excel, Hyper-V, SharePoint, SQL Server, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

© 2009 Microsoft Corporation. All rights reserved.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Part number: H5926.3

# 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 scalability (based on various user profiles) capabilities of a virtualized Enterprise Microsoft® Office SharePoint® Server 2007 enterprise farm<sup>1</sup>. In the solution, EMC's CLARiiON® CX3-40c array is used for storage and consolidation, while Microsoft SQL Server® 2005 is used as the relational database management system that supports the various Office SharePoint Server 2007 content data types.

The Hyper-V™ (a feature in the Windows Server® 2008 operating system) virtual machine migration, in conjunction with its Failover Clustering, provides high availability of all the applications running on the virtual machines. In addition, the Microsoft Performance Resource Optimization (PRO) via Microsoft System Center Virtual Machine Manager, and System Center Operations Manager, continuously monitors host utilization and ensures that all host server hardware is evenly utilized for optimized farm performance.

## The business challenge

On a daily basis, organizations are creating huge volumes of unstructured content within multiple organically grown, unmanaged, incompatible systems. This unstructured content varies, including documents, e-mail, video files, and Web pages. The content is often unorganized, which can lead to inefficient information sharing and reduced employee productivity.

The infrastructure required to support such disparate environments tends to cause physical server and storage sprawl, leading to additional data center demands. This results in increased energy consumption, physical footprint, and maintenance costs, contributing to a higher total cost of ownership (TCO).

- 
1. A farm is a collection of computer servers required to accomplish server needs far beyond the capability of one machine.

In addition, multiple isolated storage deployments contribute to inconsistent levels of service and increased complexity of backup and recovery procedures. Management of these systems is a complicated and costly task, which in turn, increases the risk of data loss and extends recovery times. This solution illustrates how organizations can:

- Manage diverse content and streamline business processes by using one common information infrastructure.
- Enable affordable high availability across an entire virtualized environment, by minimizing downtime and service level disruption.
- Ensure a more effective method of handling data growth.
- Optimize data center efficiency by minimizing physical footprint, connectivity requirements, and energy consumption, through virtualization and storage consolidation.

## The technology solution

This solution describes a virtualized Microsoft Office SharePoint Server 2007 enterprise farm environment that can serve a particular user count (based on three user profiles). The environment consists of two Microsoft failover clusters:

- SQL Server cluster
- Hyper-V cluster

The SQL Server cluster consists of an active and passive node. The Hyper-V cluster consists of a three-node active cluster that contains the entire infrastructure required to operate an Office SharePoint Server 2007 farm, for example, domain controllers, application servers, and Web front ends (WFE). Microsoft Office SharePoint Server 2007 uses Microsoft SQL Server 2005 as its datastore. EMC's CLARiiON CX3-40c provides flexible, highly available storage, while ensuring performance to the entire SharePoint farm. Hyper-V enables server consolidation through virtualization. This solution provides information<sup>2</sup> on:

- Creating a well-performing storage design for a virtualized Microsoft Office SharePoint Server 2007 farm on an EMC® CLARiiON CX3-40c with a large and very active database.
- Designing and deploying an Office SharePoint Server 2007 farm on Hyper-V machines.

---

2. For further information on this solution, please contact your EMC representative.

- Documenting the performance of the SharePoint 2007 farm under a heavy user load, and the performance impact under various system failure and recovery scenarios.

## 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. The CX3-40c offers leading flexibility with support for both Fibre Channel (FC) and iSCSI connectivity.

**Hyper-V virtualization technology** — Businesses are under pressure to ease management and reduce costs while retaining and enhancing competitive advantages, such as flexibility, reliability, scalability, and security. The fundamental use of virtualization to help consolidate many servers on a single system while maintaining isolation helps address these demands. One of the main benefits of server consolidation is a lower total cost of ownership (TCO), not just from lowering hardware requirements but also from lower power, cooling, and management costs<sup>3</sup>.

**System Center Virtual Machine Manager** — Microsoft System Center Virtual Machine Manager 2008 provides a management solution for the virtualized data center that helps enable centralized management of IT infrastructure, increased server utilization, and dynamic resource optimization across multiple virtualization platforms<sup>4</sup>. **SCVMM PRO** is a subset of System Center Virtual Machine Manager. It uses System Center Operations Manager Monitoring to provide administrators with performance-related tips to maximize hardware utilization.

---

3. Microsoft website source.

4. Microsoft website source.

## Environment profile

The SharePoint farm was designed as a Publishing / Collaboration portal. Fifteen SharePoint sites (document centers) were populated with random user data configured as 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.

### Disk layout

The solution used the Microsoft maximum recommended sizing of 100 GB for each SQL Server/Office SharePoint Server 2007 content database. A total of 1.5 TB<sup>5</sup> of data was split into 15 x 100 GB content databases, which was spread across 15 x 300 GB disks using RAID 5 for redundancy. All virtual machines used the CX3-40c for operating systems and content drives, except for the SQL Server cluster, which had local operating system drives.

### Microsoft SQL Server 2005

An active / passive SQL Server cluster was used for redundancy. The SharePoint content databases were pre-configured (datafile disk space was pre-allocated for optimum performance). Separate data and log LUNs were also used.

### Hyper-V infrastructure

All servers in the environment were virtual machines except for the SQL Server cluster. Hyper-V uses Microsoft Failover Clustering to enable flexible migration of virtual machines between the Hyper-V cluster nodes. This functionality enabled scaling of the servers as required, based on usage and load.

### Microsoft Office SharePoint Server 2007 virtual machine resources

As the WFEs are CPU-intensive, each WFE virtual machine was allocated four virtual CPUs. The Microsoft Office Excel<sup>®</sup> calculation server was configured as the Central Administration server, and there was one generalized application server used to serve farm purposes, for example document conversions.

- 
5. In total, 1.5 TB (equating to 6,818,250 documents) of data was used in testing the environment.

As crawling is a CPU-intensive activity the Index server was configured with four CPUs. To improve the crawl speed, the Index server also ran a dedicated WFE role, so all Index requests were maintained within one machine. The incremental crawl schedule was set to 15 minutes. In addition, every WFE server was a Query server for search purposes.

## Hardware resources

Table 1 details the hardware resources used in the solution.

**Table 1 Hardware resources**

Equipment	Quantity	Configuration
Storage array	1	<ul style="list-style-type: none"> <li>EMC CLARiiON CX3-40c FC array with FLARE 26 (version 3.26.040.5.016)</li> <li>30 x 300 GB FC, RAID 5, 15k rpm disks</li> <li>6 x 300 GB FC, RAID 10, 15k rpm disks</li> </ul>
SAN	2	32-port 4 Gb/s FC switches
Network switches	2	48-port Gigabit Ethernet network switches (inter-linked)
Hyper-V farm	3	Infrastructure Servers: 4-socket quad core Xeon E7330 2.4 GHz (16 cores), 48 GB RAM
Virtual machines	1	<b>Index server (running WFE role dedicated for crawling):</b> 4 CPUs, 2.4 GHz, 6 GB RAM
	10	<b>WFEs (running query):</b> 4 CPUs, 2.4 GHz, 4 GB RAM
	2	<b>Application servers:</b> 2 CPUs, 2.4 GHz, 2 GB RAM
	2	<b>Domain controllers:</b> 2 CPUs, 2.4 GHz, 2 GB RAM
HBAs for SAN	2	4 GB dual-port PCI-e HBAs
NICs for iSCSI	3	Quad-port Gigabit Ethernet PCI-e network cards
SCOM host	1	2-socket quad core Xeon E7330, 2.4 GHz (8 cores), 16 GB RAM (also used as DC)
SCVMM host	1	2-socket quad core Xeon E7330, 2.4 GHz (8 cores), 16 GB RAM (also used as DC)
SQL Server active node	1	4-socket quad core Xeon E7330, 2.4 GHz (16 cores), 48 GB RAM
SQL Server passive node	1	2-socket quad core Xeon E7330, 2.4 GHz (8 cores), 16 GB RAM



## Software resources

Table 2 details the software resources used in the solution.

**Table 2 Software resources**

Software	Version
Windows Server 2008	Enterprise edition (64-bit)
Microsoft SQL Server 2005	SP2 Enterprise edition (64-bit)
Microsoft Office SharePoint Server 2007	SP1 plus infrastructure update 1 and August 2008 rollup pack
Windows Server 2008 Hyper-V technology	1.0 (GA KB 950050).
SCVMM 2008	2.03 (GA)
SCOM 2007	SP1

## Farm content types

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

**Table 3 Content types**

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

## 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.

The CLARiiON CX3-40c was not stressed during maximum user testing: SP utilization did not exceed **35%** and LUN utilization was within acceptable parameters.

**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.

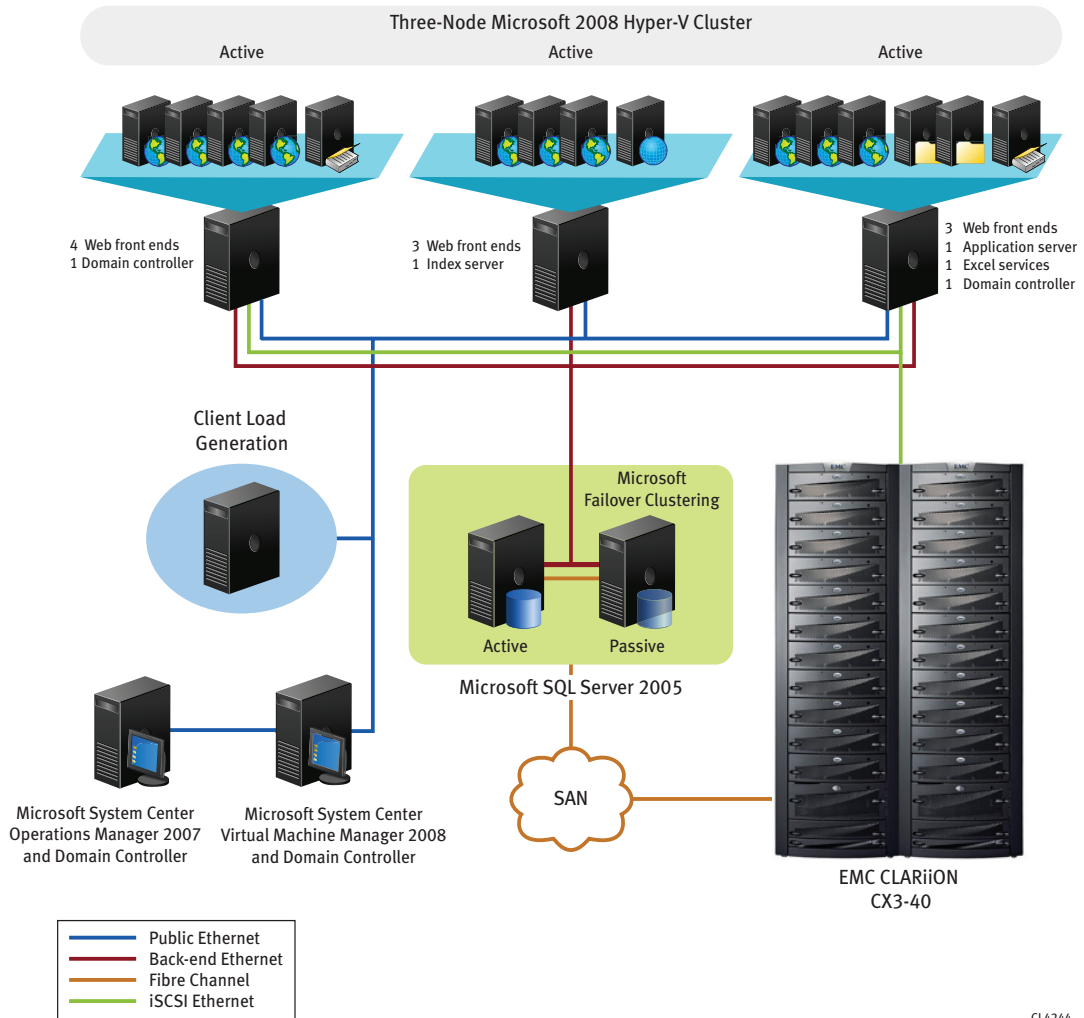
**Table 5** Performance profiles

Portal usage type	User profile <sup>a</sup> (browse / search / modify %)	Requests <sup>b</sup> per sec	Concurrency	Max user capacity <sup>c</sup>	Average user response time (secs) (browse / search / modify)
Heavy read	80 / 10 / 10	50	1%	300,000	< 3 / < 3 / < 5
Heavy publishing	70 / 5 / 25	52.8	1%	316,800	< 3 / < 3 / < 5
Heavy collaboration	50 / 20 / 30	54	1%	324,000	< 3 / < 3 / < 5

- All users were run against a Microsoft heavy user profile, that is, 60 requests per hour.
- 0% think time was applied to all tests. "0% think time" is the elimination of typical user decision time when browsing, searching, or modifying in Office SharePoint Server. For example, a single complete user request is completed from start to finish without user pause, therefore, creating a continuous workload on the system.
- The maximum user capacity is derived from the following formula:  
**# = seconds per hour / RPH / Concurrency% \* RPS**  
 Example: 3600 / 60 / 1% \* 54 = 324,000  
 Example: 3600 / 60 / 10% \* 54 = 32,400 (supported user capacity for 10% concurrency)

## Physical architecture

Figure 1 illustrates the physical architecture of the solution.



CL4244

**Figure 1 Office SharePoint Server 2007 farm environment enabled by Hyper-V**

## Conclusion

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

Sizing and configuring an Office SharePoint Server 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 Office SharePoint Server 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.

Windows Server 2008 Hyper-V technology and Microsoft System Center Virtual Machine Manager on EMC's CLARiiON CX3-40c provide a very valuable proposition in offering greater flexibility by evenly spreading the load and optimizing the utilization of the physical hardware.

The solution also provides the following benefits:

- Illustrates real-world expectations for realistic CLARiiON storage requirements and provisioning for a typical Office SharePoint Server 2007 enterprise farm.
- Illustrates how to design and scale an Office SharePoint Server 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 Hyper-V solution based on a given Office SharePoint Server 2007 farm configuration.
- Reduces configuration timelines by documenting the key design considerations.

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

To learn more about this SharePoint solution, contact an EMC representative or visit the Solutions section of EMC Powerlink (<http://powerlink.emc.com>).

## About Microsoft virtualization solutions

Microsoft provides a complete suite of technologies to enable an integrated, end-to-end, virtualized infrastructure. Using products that span the desktop to the data center, Microsoft technologies bring capacities online in real time, as needed; streamline and provision applications, services and data on-demand; accelerate backup and recovery; and enhance availability to protect against system failure and service interruptions.

Microsoft's extensive partner ecosystem complements and extends the Microsoft virtualization toolset with products for desktops, servers, applications, storage, and networks. Together with their partners, Microsoft delivers the most robust, complete solutions for the virtualized infrastructure.

For more information about Microsoft Joint Virtualization Solutions, visit [www.microsoft.com/virtualization/partners.msp](http://www.microsoft.com/virtualization/partners.msp).