

**EMC Virtual Infrastructure for
Microsoft Office
SharePoint Server 2007**

Enabled by EMC Celerra and VMware vSphere 4

Reference Architecture

EMC NAS Product Validation



Copyright © 2009 EMC Corporation. All rights reserved.

Published August, 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.

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

Part number: H6470

Table of Contents

Reference architecture overview.....	4
Solution architecture	6
Key components	8
Validated environment profile.....	11
Hardware and software resources	12
Conclusion	13

Reference architecture overview

Document purpose EMC's commitment to consistently maintain and improve quality is led by the Total Customer Experience (TCE) program, which is driven by Six Sigma methodologies. As a result, EMC has built Customer Integration Labs in its Global Solutions Centers to reflect real-world deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges currently facing its customers.

Solution purpose The purpose of this reference architecture document is to demonstrate functional and performance aspects of the Microsoft Office SharePoint Server (MOSS) 2007 solution built on a VMware vSphere 4 cluster using an EMC® Celerra® storage system. This reference architecture validates the performance of all aspects of the solution and provides guidelines for building similar solutions.

This reference architecture is not intended to be a comprehensive guide to every aspect of the EMC Virtual Infrastructure for Microsoft Office SharePoint Server 2007 enabled by EMC Celerra and VMware vSphere 4 solution. For more detailed information on this solution, please see *EMC Virtual Infrastructure for Microsoft Office SharePoint Server 2007 Enabled by EMC Celerra and VMware vSphere 4 Proven Solution Guide*.

Business challenge Microsoft Office SharePoint Server 2007 is a server application that facilitates collaboration, provides content management features, and implements business processes. It provides an integrated platform to plan, deploy, and manage intranet, extranet, and Internet applications across and beyond the enterprise. SharePoint uses multiple servers in various roles to organize and deliver website collaboration and information sharing across organizations.

Critical business information stored in SharePoint sites is rapidly increasing. Therefore, the need for enterprises to manage these SharePoint sites is greater today than ever before. When knowledge workers collaborate on projects and author documents:

- Storage requirements increase
- Data compliance and security become important
- Processes and policies require high-availability servers

In addition, these enterprises must manage IT costs and reduce the risk of business disruption. An increasing number of companies are exploring ways to virtualize SharePoint environments to reduce costs, increase availability, add flexibility, and use resources more efficiently.

Technology solution

SharePoint systems are environments and not applications by themselves. Many factors must be considered when deploying solutions into an existing environment, or planning a new application deployment. EMC proven solutions allow you to start from a known reference configuration and then customize it for your needs by examining the similarities and differences between various methods of accomplishing common tasks in a SharePoint environment.

This solution uses VMware virtualization technology to address cost reduction and disruption avoidance. By decreasing the required number of physical servers, enterprises can achieve substantial power and hardware savings. Virtualized servers include features that increase the availability of virtual machines (VMs) and help balance workload across ESX servers. VMware HA clustering allows VMs to automatically recover from ESX host failure and restart on a surviving host in the cluster. VMware Distributed Resource Scheduler (DRS) balances workloads across the cluster for better performance.

This solution addresses each of these challenges by using tested and proven solutions validated by EMC Engineering in the NAS Product Validation Lab in Research Triangle Park, North Carolina.

Solution benefits

Virtualizing SharePoint Environments Reduces Hardware Cost:

Virtualized SharePoint environment reduces cost by utilizing infrastructure effectively. Virtualization enables reduction in the number of physical servers required to implement a SharePoint environment. The tested configuration will support a medium SharePoint farm servicing between 9 and 11 requests per second representing 9,000 to 11,000 users at a 10% concurrency rate.

Minimal Downtime for your SharePoint Environment: An ESX Server clustered environment with High Availability features constantly monitors the physical hardware failure and provides automated failover and recovery of virtual machines with minimal downtime of less than 3 minutes. Test results of the Web Front End server recovery averaged 2 minutes and SQL Server recovery averaged 2 minutes and 42 seconds. We were unable to recognize any downtime when the Application/Index Server VM migrated.

No Stress Server Load Balancing: VMware Distributed Resource Scheduler (DRS) continuously monitors the resource utilization across the cluster and intelligently balances the load with zero downtime. The DRS automation level can be set to automatic or manual migration of virtual machines, freeing administrators from ever having to worry about the placement of VM's.

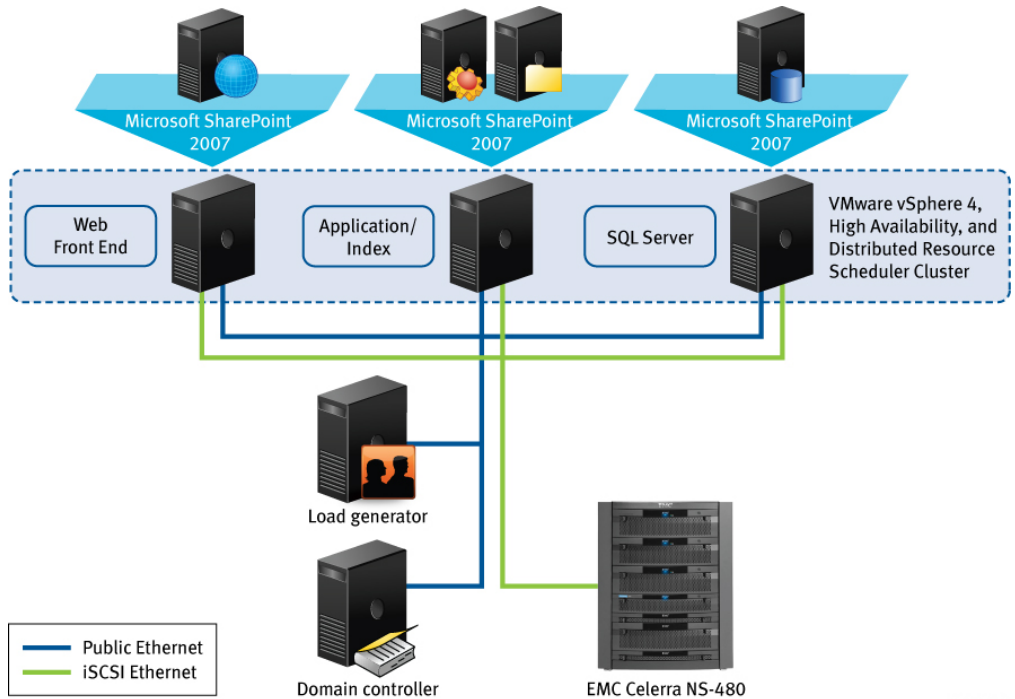
Storage Optimization via Virtual Provisioning: VMware vStorage Thin Provisioning of Virtual disks used along with VMotion (which enables dynamic migration of VMDKs) or VMFS volume grow (which provides the ability to dynamically increase the size of your datastore) provides:

- Higher utilization of virtual machine storage by letting you dedicate more storage capacity than the actual purchased capacity.
- Simplified storage capacity management.

Solution architecture

Architecture diagram

The following illustration depicts the three-node VMware vSphere 4 High Availability and Distributed Resource Scheduler cluster solution.



Reference architecture overview

VMware ESX 4.0 enables guest operating systems to connect to storage by using several methods. This reference architecture specifies that database storage be provided to the server using the iSCSI protocol at the ESX server level by using initiator software and the Virtual Machine File System (VMFS) protocol on the ESX servers. The VMs connect to iSCSI logical unit numbers (LUNs) as VMFS datastores on the ESX server.


A cluster is a collection of ESX server hosts and associated VMs with shared resources and a shared management interface. When you add a host to a cluster, the host's resources become part of the cluster's resources. When you create a cluster you can choose to enable the cluster for HA, DRS, or both. For this configuration, a VMware cluster was created using three ESX 4.0 servers with both HA and DRS features enabled for the cluster. A resource pool was also created on the cluster to centrally manage all the virtual machine resources. Four VMs running MOSS 2007 and SQL Server were balanced across the ESX servers and automatically restarted on a

EMC Virtual Infrastructure for Microsoft Office SharePoint Server 2007 Enabled by EMC Celerra and VMware vSphere 4 Reference Architecture

surviving ESX server during a physical server failure.

Storage layout

The following illustration depicts the overall storage layout of the solution.

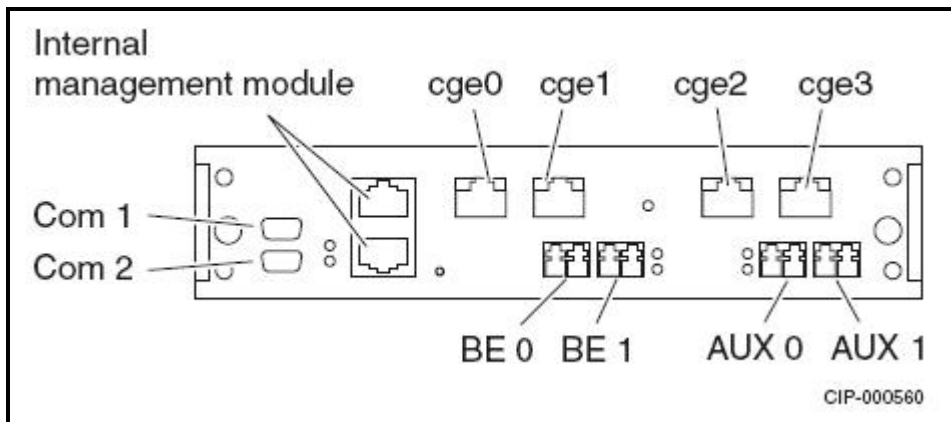


Storage layout overview

The validated solution uses the first shelf of a Celerra NS-480 storage system to host SharePoint databases, logs, and Index files. The database disk was set up in a RAID 5 (4+1) group to increase storage space for the SharePoint solution. The logs and tempdb files were placed on separate RAID 1 groups to maximize performance of the SharePoint databases.

VMware vStorage thin provisioning technology was used while creating the database, log, and virtual machine operating system disks for higher utilization of storage space. All of the virtual servers are stored on Celerra iSCSI storage shared among the three ESX hosts in the VMware HA and DRS cluster used in this solution. The RAID 5 (4+1) group that holds the Celerra system volumes has been used for this. The SharePoint Index files were also stored on the RAID 5 (4+1) group that holds the Celerra system volumes.

Network layout The following illustration shows the ports on an EMC Celerra NS-480.



Network layout overview The EMC Celerra NS-480 contains two Data Movers that can operate independently. They can also operate in the active/passive mode, with the passive Data Mover serving as a failover device for the active Data Mover. In this solution, Data Movers operate in the active/passive mode. An NS-480 Data Mover has four Gigabit Ethernet ports.

Ports cge0 and cge1 handle storage traffic. Ports cge2 and cge3 remain open for future use. The Data Mover supports several types of link aggregation for IP traffic. However, for this configuration, no link aggregation was used.

As a best practice, the Data Mover network ports connected to the storage network should be dedicated to storage traffic. However, if the ports are not heavily used, they can be shared with non-storage network traffic. EMC recommends monitoring the network to avoid bottlenecks.

Key components

Introduction The key components of the Microsoft Office SharePoint Server 2007 solution include:

- EMC Celerra unified storage
- VMware vSphere 4

For details on all the components that make up the reference architecture, see the section entitled [Hardware and software resources](#).

EMC Celerra unified storage

EMC Celerra unified storage provides access to block and file data using the iSCSI, CIFS, NFS, and Fibre Channel (FC) protocols. Using these network services, EMC Celerra platforms deliver a complete multi-protocol storage solution. EMC Celerra unified storage allows storage consolidation via existing IP infrastructures along with the power of high-speed FC networks.

VMware vSphere 4 with ESX 4.0

The new VMware vSphere 4 provides significant performance enhancements that make it easier for organizations to virtualize their most demanding and intense workloads.

vSphere 4 includes ESX 4.0 and the management interface, vCenter. VMware ESX 4.0 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.

This validated solution utilizes VMware HA, VMware DRS, and vStorage thin provisioning advanced features of vSphere 4 infrastructures to provide a comprehensive solution for Microsoft Office SharePoint Server 2007.

The following table provides a summary of the advanced features and benefits of this solutions

Feature	Benefits
VMware High Availability (HA)	<ul style="list-style-type: none">• Automatically detects physical machine failure of ESX servers and restarts the virtual machines on other ESX servers in a shared storage environment.• Ensures that capacity is always available in order to restart all virtual machines affected by ESX server failure. HA continuously monitors capacity utilization and “reserves” spare capacity to be able to restart virtual machines.• Protects against operating system failures with virtual machine failure monitoring in VMware HA.• Suspends failover actions during maintenance operations on servers.
VMware Distributed Resource Scheduler (DRS)	<ul style="list-style-type: none">• Aggregates resources across many servers into shared resource pools.• Manages resources independently of the physical servers that contribute the resources.• Organizes resource pools hierarchically to match available IT resources to the business organization. DRS ensures

	<p>that resource utilization is maximized while business units retain control and autonomy of their infrastructure. Resource pools can be flexibly added, removed, or reorganized as business needs or organization change.</p> <ul style="list-style-type: none"> • Aligns computing resources with business goals while ensuring flexibility and efficient utilization of hardware resources. VMware DRS continuously monitors utilization across resource pools and intelligently allocates available resources among virtual machines based on pre-defined rules and policies. VMware DRS dynamically responds to changing virtual machine requirements using VMware VMotion to move virtual machines without disruption between servers. This automates the operational management of virtual machine environments.
VMware vStorage Thin Provisioning	<ul style="list-style-type: none"> • Allows overallocation of storage capacity for increased storage utilization. • VI administrators gain the ability to allocate virtual disk files as “thick” or “thin”. • Allows virtual machines on VMware ESX hosts to provision the entire space required for the disk’s current and future activities, but initially commits only as much storage space as the disk needs for its initial operation. This is achieved with zero performance impact, continuous service availability and complete data integrity. • Enables organizations to provision heterogeneous storage pools, increase utilization, and reduce administration costs.
VMotion and Storage VMotion	<ul style="list-style-type: none"> • VMotion migrates running virtual machines from one server to another with no disruption or downtime. This can be used to avoid application disruptions due to planned hardware maintenance. • Storage VMotion migrates running virtual machines from one storage location to another with no disruption or downtime. It can be used across different types of storage (FC, iSCSI,

	NFS, and even local storage) allowing customers to allocate the right tier of storage to applications based on their lifecycle.
--	---

Validated environment profile

Profile characteristics The solution was validated with the following environment profile.

Profile characteristic	Value
SharePoint farm user data	668 GB
Concurrency	10%
Enterprise portal collaboration site collections	1
Document Library sites	10
Number of documents in the SharePoint farm	3,344,625
Size of Content Index files on Index and WFE-Query servers	17.26 GB
Web front-end servers (VMs)	1 (also running query role)
Index servers (VMs)	1 (also running Web services dedicated for crawling)
SQL servers (VMs)	1
Application servers (VMs)	1

Hardware and software resources

Hardware The following table lists the hardware used to validate the solution.

Equipment	Quantity	Configuration
Storage array	One	EMC Celerra NS-480 unified storage DART: 5.6.44.4 Fifteen 300 GB disks (15k rpm)
Servers	Three	Dell 2950 server class Four 3.0 GHz dual core processors 16 GB of RAM Four Gigabit Ethernet NICs
Servers	Two	Dell 1850 Server Class Two 2.8 GHz dual core processors 4 GB RAM 2 Gigabit Ethernet NICs
Enterprise network switch	One	Gigabit Ethernet network switch
Virtual machines	Four	<ul style="list-style-type: none"> • WFE server (running query service): Four CPUs, 3.0 GHz, 4 GB RAM • SQL server: Four CPUs, 3.0 GHz, 14 GB RAM • Index server (running WFE role dedicated for crawling): Two CPUs, 3.0 GHz, 4 GB RAM • Application server: Two CPUs, 3.0 GHz, 4 GB RAM

Software The following table lists the software used to validate the solution.

Software	Version
VMware ESX 4.0	4.0.0
VMware vSphere	4.0.0
Microsoft Windows Server	Windows 2003 x64 Enterprise Edition R2 SP2 (4) Windows 2003 x32 Standard Edition R2 SP2 (2)
Microsoft Office SharePoint Server 2007 Enterprise Edition	SP1 through FEB 24 2009 Updates
EMC Celerra DART	5.6.44.4
EMC CLARiiON® FLARE®	R23 02.23.050.5.705

Conclusion

Summary

SharePoint is a highly configurable and variable application environment, and each customer will have a different set of requirements, performance goals, and recovery targets. This document attempts to provide a framework for implementing these solutions that addresses the inherent complexity of many environments, while retaining the ease of use that characterizes the EMC Celerra platform.

Next steps

EMC can help accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of a solution for a Microsoft Office SharePoint Server 2007 environment.

To learn more about this and other solutions, contact an EMC representative.
