




**EMC Unified Storage for
Microsoft Office SharePoint
Server 2007**

**BLOB Externalization Enabled by
EMC Celerra and Metalogix StoragePoint**

Reference Architecture

EMC Unified Storage Solutions



Copyright © 2010 EMC Corporation. All rights reserved.

Published August, 2010

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

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: H8013

Table of Contents

Reference architecture overview	4
Solution architecture.....	6
Key components.....	8
Validated environment profile	9
Hardware and software resources	10
Conclusion	11

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.

This document describes the reference architecture of the EMC Unified Storage for Microsoft Office SharePoint Server 2007 BLOB Externalization Enabled by EMC[®] Celerra[®] and Metalogix StoragePoint solution, which was tested and validated by EMC Information Infrastructure Solutions.

Solution purpose

The purpose of this reference architecture is to demonstrate how to leverage EMC's expertise and proven technologies to relocate and manage binary large objects (BLOBs) in the SharePoint farm, which is known as external BLOB storage (EBS).

The purpose of this solution is to:

- Design a Microsoft Office SharePoint Server (MOSS) 2007 environment with external BLOB storage on a Celerra file system.
- Realize space-saving benefits in the SharePoint farm content databases as a result of external BLOB storage. An 88 percent reduction in the amount of space used by content databases was observed.
- Use Celerra deduplication to save BLOB storage file system space. In this solution, 18 percent file system space savings was observed.

This reference architecture is not intended to be a comprehensive guide to every aspect of the solution, but only to offer an overview of the architecture and discuss the benefits of deploying an external BLOB storage solution by using Celerra unified storage systems.

The business challenge

Storing BLOBs in a SharePoint database results in a number of challenges as the content volume grows. As large files are uploaded to SharePoint sites, the increased size of the database requires extended disaster recovery, indexing, and maintenance timeframes. Also, an additional SQL I/O overhead is created by the retrieval of BLOBs. Externalizing SharePoint BLOBs eases the workload of SQL servers in the farm. Another benefit of external BLOBs is that full farm crawls and backups complete three to six times faster.

SharePoint 2007 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 files and documents are routinely stored in SharePoint farms. Therefore, enterprises have a greater need than ever to manage SharePoint sites and the files they hold. Databases are an integral part of this environment. Managing these databases is critical to the end-user experience on the SharePoint farm.

The technology solution

This solution demonstrates how to use an EMC Celerra unified storage platform to hold external BLOBs for a SharePoint 2007 environment. In this solution, Metalogix StoragePoint externalizes the SharePoint documents from SQL Server to a file system on Celerra.

This solution includes file system access to SharePoint documents instead of database retrieval of BLOBs.

Solution benefits

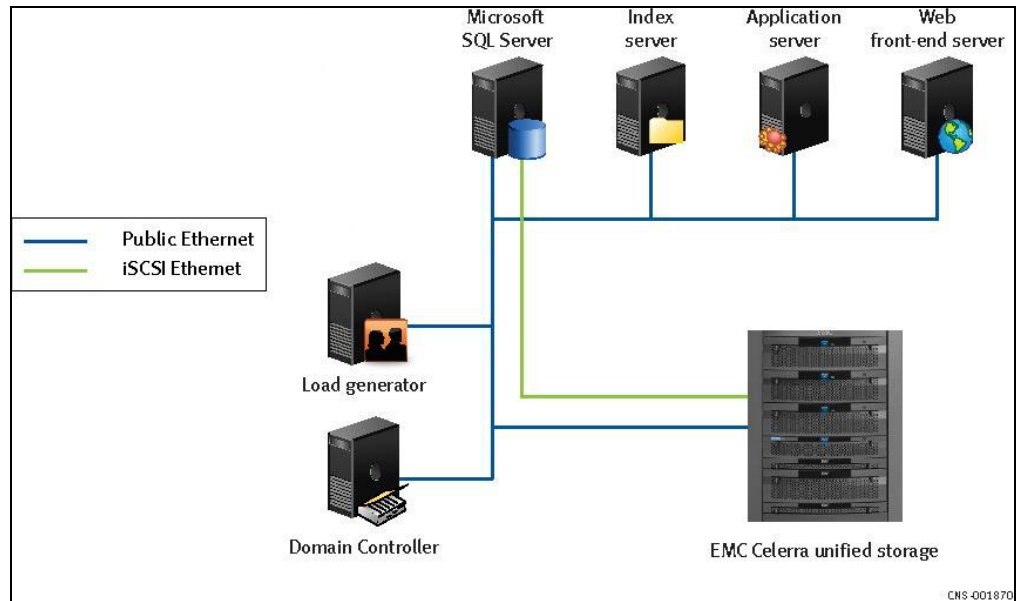
The benefits of this solution are as follows:

- **External BLOB storage frees database resources:** Storing BLOBs in file systems outside the SQL database frees valuable space and speeds indexing and SharePoint farm backup operations. An 88 percent reduction in space was observed in the SQL content database disk usage report.
- **File system storage introduces features such as deduplication:** With SharePoint documents in dedicated file systems, advanced features of the Celerra unified storage system can be used by the farm.
- **One platform to hold and back up the BLOB file systems:** Celerra can perform NDMP backup-to-disk operations on the BLOB file system. One platform acts as both the source and the target for backups of the BLOB file system, which simplifies the protection of SharePoint BLOBs.
- **StoragePoint provides features such as compression and encryption:** External BLOBs can be compressed to save space and can be encrypted for added security.

Solution architecture

Architecture diagram

The following illustration depicts the logical architecture of the solution.



Reference architecture overview

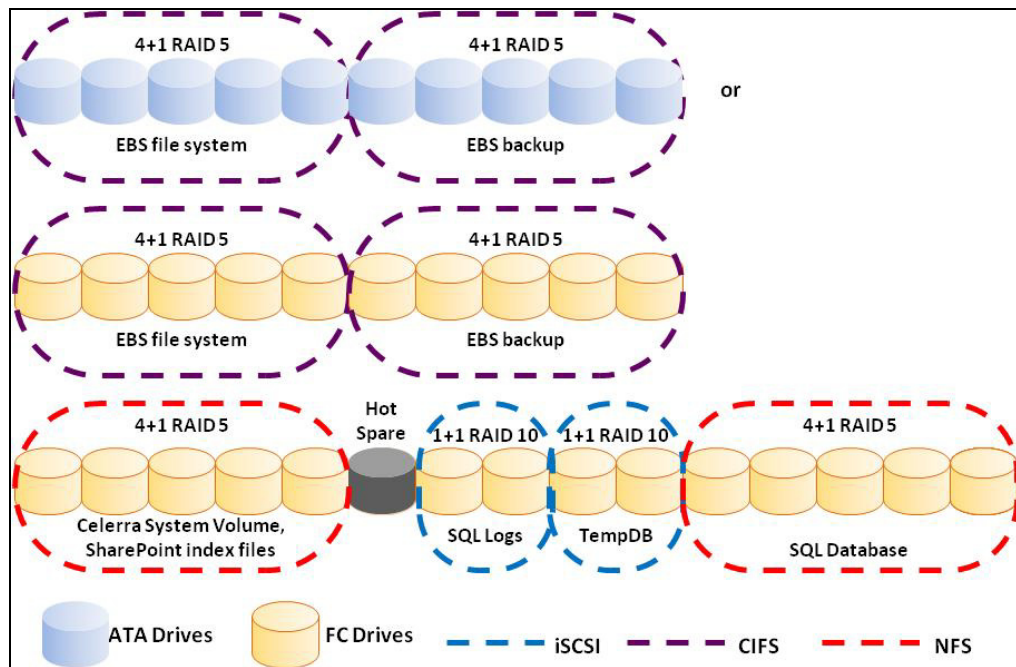
The validated solution is a SharePoint 2007 farm that uses the Celerra platform for SQL Server storage.

The key components of this architecture are:

- A web front-end server running StoragePoint software
- A SharePoint index server
- A SharePoint application server
- A SQL Server 2008 server
- A Celerra storage array to store the SharePoint configuration databases, content databases, search databases, content index files, and SQL BLOBs, which are the SharePoint documents

Storage layout

The following illustration depicts the overall storage layout of the solution. The EBS file system and backups can reside on ATA or Fibre Channel (FC) spindles.



Storage layout overview

The Celerra storage array stores the following:

- SharePoint index files, databases, logs, and tempdb
- BLOBs from SQL Server in the EBS file system
- EBS backups

SharePoint index files: A RAID 5 (4+1) group holds system volumes and SharePoint index files.

SharePoint databases logs and tempdb: The configuration tested in this solution uses RAID 5 (4+1) to store databases, and a separate RAID 1 (1+1) group to store logs and tempdb.

EBS file system: A RAID 5 (4+1) group holds the file system that StoragePoint uses to copy BLOBs from SQL databases.

EBS backup: StoragePoint needs to have its CIFS shares backed up on a regular basis.

Key components

Introduction

This section briefly describes the following key components of the SharePoint 2007 external BLOB storage solution:

- EMC Celerra unified storage platform
- StoragePoint

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

EMC Celerra unified storage platform

The EMC Celerra unified storage platform is a dedicated network server optimized for file and block access, delivering high-end features in a scalable and easy-to-use package. For high scalability, Celerra unified storage platforms leverage both the innovative EMC CLARiiON[®] Fibre Channel RAID storage, delivering best-in-class availability and data protection, and the availability, performance, and ease of management of EMC Celerra.

Celerra unified storage systems deliver a single-box block and file solution offering a centralized point of management for distributed environments. This makes it possible to dynamically grow, share, and cost-effectively manage multi-protocol file systems and provide multi-protocol block access. Administrators can take advantage of simultaneous support for NFS and CIFS protocols by allowing Windows and Linux/UNIX clients to share files using the Celerra system's sophisticated file-locking mechanisms and by leveraging iSCSI or FC for high-bandwidth or latency-sensitive applications

The Celerra unified storage platform provides native CLARiiON value-added functionality such as EMC MirrorView[™]/Synchronous, EMC MirrorView/Asynchronous, EMC SnapView[™], and clones.

StoragePoint

StoragePoint reduces the size of SharePoint content databases by relocating (externalizing) content BLOBs from the database and maintaining only the metadata within SQL Server. BLOBs are relocated to the storage area network (SAN), NAS, or cloud platforms where they can be partitioned by SLA, isolation requirements, or retention policies. The BLOB I/O is moved from SQL Server to the SharePoint web front ends, which are easier to scale than SQL Server. Scalability is addressed by adding more web servers to the farm.

Validated environment profile

Profile characteristics

The solution was validated with the following environment profile.

Profile characteristic	Value
SharePoint farm user data	162 GB
Concurrency	10%
Enterprise Portal Collaboration Site Collections	1
Document Library sites	10
Number of documents in the SharePoint farm	2,462,346
Size of content index files on index and WFE-query server	2.36 GB
Web front-end server (virtual machine)	1 (also running a query role)
Index server (virtual machine)	1 (also running Web services dedicated for crawling)
SQL Server (virtual machine)	1
Application server (virtual machine)	1

Farm user load profiles

A typical Microsoft user load profile was used to determine the maximum user count that the SharePoint farm could sustain while ensuring average response times remained within acceptable user limits. As per the Microsoft standard, a typical user performs 36 requests per hour (RPH). The following table details the acceptable user limits for SharePoint operations.

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

Hardware and software resources

Hardware

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

Hardware	Quantity	Servers
Dell PowerEdge 2950 4*1 CPU, 3.0 GHz, 16 GB of RAM	1	Microsoft SQL Server 2008
Dell PowerEdge 2950 4*1 CPU, 3.0 GHz, 16 GB of RAM	1	Application server
Dell PowerEdge 2950 4*1 CPU, 3.0 GHz, 16 GB of RAM	1	Web front-end server
Dell PowerEdge 2950 4*1 CPU, 3.0 GHz, 16 GB of RAM	1	Index server
Dell 1850 2*1 CPU, 2.9 GHz, 4 GB of RAM	1	Domain controller
Dell 1950 2*1 CPU, 2.9 GHz, 2 GB of RAM	1	Load generator
EMC Celerra unified storage One shelf of 268.4 GB (15k rpm) FC disks Five FC 268.4 GB (15k rpm) disks Five SATA II 917.2 GB (7.2k rpm) disks	1	FLARE® (version 04.29.000.5.003) NAS version (5.6.47.11)

Software

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

Software	Version
Microsoft Windows Server	Windows 2003 x64 Enterprise Edition R2 SP2 (4) Windows 2003 x32 Enterprise Edition R2 SP2 (2)
Microsoft Office SharePoint Server 2007 Enterprise Edition	SP2
Microsoft Office SQL Server 2008 Enterprise Edition	R2
EMC Celerra DART	5.6.47.11
EMC CLARiiON FLARE	R29 04.29.000.5.503

Conclusion

Summary

This reference architecture depicts a validated solution, EMC Unified Storage for Microsoft Office SharePoint Server 2007 BLOB Externalization Enabled by EMC Celerra and Metalogix StoragePoint. This solution decreases the size of SharePoint databases and subsequently reduces maintenance windows and SharePoint indexing operations. It also provides the benefits of Celerra file system storage, such as deduplication.

Next steps

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

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