



Externalizing Large SharePoint 2010
Objects with EMC VNX Series and
Metalogix StoragePoint

Proven Solution Guide

EMC Unified Storage Solutions



Copyright © 2011 EMC Corporation. All rights reserved.

Published March, 2011

EMC believes the information in this publication is accurate 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.

Part Number: h8185

Table of Contents

Chapter 1: About this Document.....	5
Overview	5
Audience and purpose	6
Scope	6
Reference architecture	7
Introduction to the EMC VNX series	8
Prerequisites and supporting documentation.....	9
Terminology.....	9
Chapter 2: SharePoint 2010 Farm Design.....	11
Overview	11
SharePoint farm design.....	12
Chapter 3: Virtualization.....	15
Overview	15
Concepts	16
Advantages of virtualization	16
Considerations	17
Implementation.....	17
Chapter 4: Network Design.....	19
Overview	19
Considerations	20
Implementation.....	20
Chapter 5: Storage Design.....	22
Overview	22
Design considerations.....	23
Storage design implementation.....	24
Chapter 6: SharePoint Data Backup with Replication Manager	25
Overview	25
Design considerations.....	26
Design implementation.....	26
Best practices.....	28
Chapter 7: Testing and Validation.....	29
Overview	29
Testing approach	30
Methodology.....	31
Testing overview	33
Test results analysis.....	34
BLOB externalization	34
SharePoint indexing results	35
Benefits of BLOB externalization	36

Impact of BLOB externalization.....	37
VNX5300 performance summary.....	40
Backup and restore of the externalized SharePoint farm	43
Using EMC file deduplication and compression with externalized BLOBs	47

Chapter 1: About this Document

Overview

Introduction

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 realworld 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 Proven Solution Guide summarizes a series of best practices that were discovered or validated during the testing of the Microsoft SharePoint 2010 BLOB Externalization solution with the following products:

- EMC® VNX5300™
- Microsoft SharePoint 2010
- Microsoft Windows 2008 R2 Hyper-V
- Metalogix StoragePoint 3.0
- EMC Replication Manager 5.3 SP1

The implementation of the SharePoint 2010 BLOB Externalization solution was carried out in a virtualized environment to leverage the existing equipment and to streamline the testing process.

Use case definition

This solution demonstrates the functional and performance aspects of the external BLOB storage solution for Microsoft SharePoint 2010 built on a virtual infrastructure that uses Microsoft Windows 2008 R2 Hyper-V.

Contents

This chapter contains the following topics:

Topic	See page
Audience and purpose	6
Scope	6
Reference architecture	7
Introduction to the EMC VNX series	8
Prerequisites and supporting documentation	9
Terminology	9

Audience and purpose

Audience The intended audience for this Proven Solution Guide is:

- Customers
 - Internal EMC personnel
 - EMC partners
-

Purpose The purpose of this reference architecture is to demonstrate how to use EMC's expertise and proven technologies to externalize and manage binary large objects (BLOBs) in the SharePoint farm. This solution is built on an EMC VNX5300 platform with multiprotocol support, which enables Fibre Channel (FC) block-based storage for SharePoint databases and CIFS-based storage for BLOB data.

The purpose of this solution is to:

- Design a Microsoft SharePoint Server 2010 environment with external BLOB storage on an EMC VNX5300.
- Realize space-saving benefits in the SharePoint farm content databases as a result of external BLOB storage.
- Back up and restore SharePoint databases and the BLOB store by using VNX™ SnapView™ and EMC Replicator technologies through EMC Replication Manager.

This Proven Solution Guide 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 the EMC VNX series.

Scope

Scope This document contains the results of testing a SharePoint 2010 server with externalized BLOBs by using an EMC VNX5300 and the StoragePoint 3.0 RBS provider from Metalogix. The objectives of this testing are to establish:

- A reference architecture of validated hardware and software that permits easy and repeatable deployment of SharePoint 2010 with BLOB externalization by using an EMC VNX platform.
 - The storage best practices to configure SharePoint 2010 by using EMC VNX platforms in a manner that provides optimal performance, recoverability, and protection in the context of the midtier enterprise market.
-

Not in scope Information on how to install, configure, design, and build architecture for a SharePoint 2010 farm and the required EMC products is outside of the scope of this document. However, links are provided on where to find all the required software for this solution.

Reference architecture

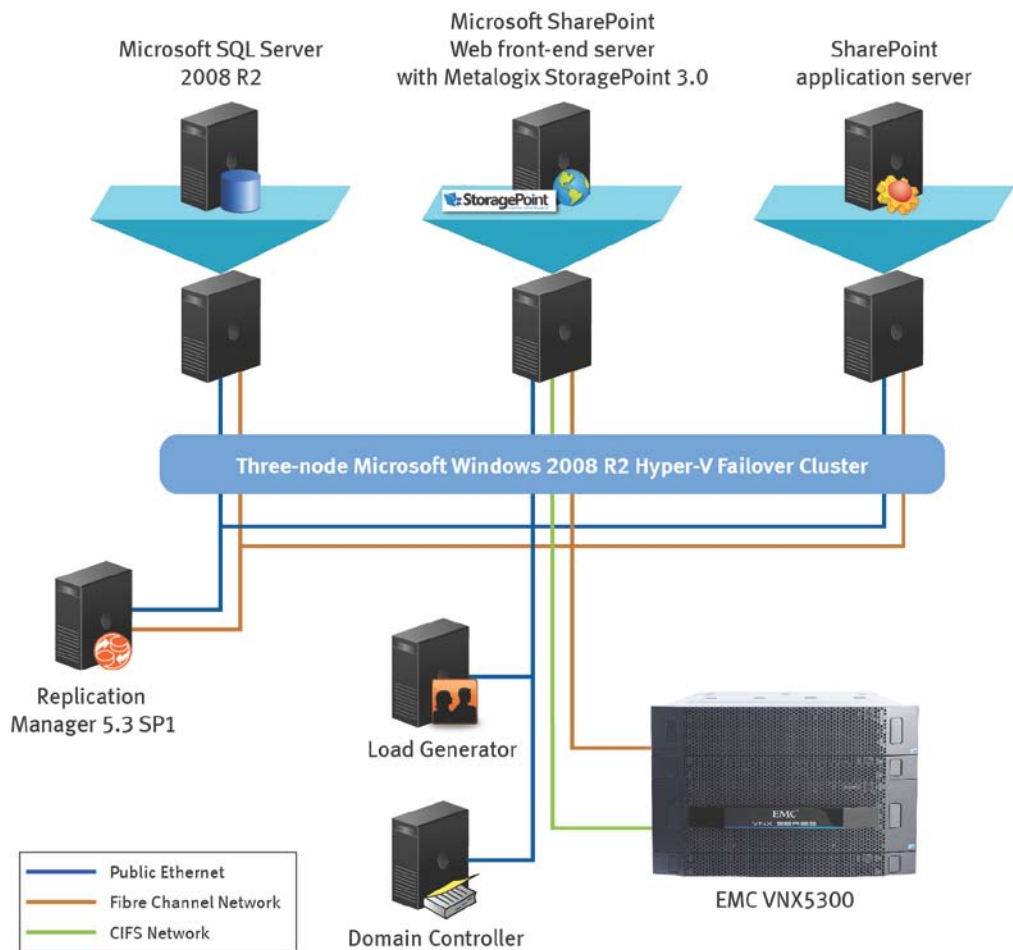
Corresponding reference architecture

This Proven Solution Guide has a corresponding Reference Architecture document that is available on Powerlink® and EMC.com. *Externalizing Large SharePoint 2010 Objects with EMC VNX Series and Metalogix StoragePoint—Reference Architecture* provides more details.

If you do not have access to the document, contact the EMC representative.

Reference architecture diagram

The following diagram is the logical architecture of the use case. The validated environment uses virtual servers. However, this is not a requirement of the solution. The use case can also be built from physical servers with no loss of validity.



VNX-000065

Introduction to the EMC VNX series

EMC VNX series

The VNX series delivers uncompromising scalability and flexibility for the midtier while providing market-leading simplicity and efficiency to minimize total cost of ownership. Customers can benefit from the new VNX features such as:

- Next-generation unified storage, optimized for virtualized applications
- Extended cache using Flash drives with FAST Cache and Fully Automated Storage Tiering for Virtual Pools (FAST VP), which can be optimized for the highest system performance and lowest storage cost simultaneously on both block and file.
- Multiprotocol support for file, block, and object with object access through Atmos™ Virtual Edition (Atmos VE).
- Simplified management with EMC Unisphere™ for a single management framework for all NAS, SAN, and replication needs.
- Up to three times improvement in performance with the latest Intel multicore CPUs, optimized for Flash.
- 6 Gb/s SAS back end with the latest drive technologies supported:
 - 3.5" 100 GB and 200 GB Flash, 3.5" 300 GB, and 600 GB 15k or 10k rpm SAS, and 3.5" 2 TB 7.2k rpm NL-SAS
 - 2.5" 300 GB and 600 GB 10k rpm SAS
- Expanded EMC UltraFlex™ I/O connectivity—Fibre Channel (FC), Internet Small Computer System Interface (iSCSI), Common Internet File System (CIFS), Network File System (NFS) including parallel NFS (pNFS), Multi-Path File System (MPFS), and Fibre Channel over Ethernet (FCoE) connectivity for converged networking over Ethernet.

The VNX series includes five new software suites and three new software packs, making it easier and simpler to attain the maximum overall benefits.

Software suites available

- VNX FAST Suite—Automatically optimizes for the highest system performance and the lowest storage cost simultaneously (FAST VP is not part of the FAST Suite for the VNX5100™).
- VNX Local Protection Suite—Practices safe data protection and repurposing.
- VNX Remote Protection Suite—Protects data against localized failures, outages, and disasters.
- VNX Application Protection Suite—Automates application copies and proves compliance.
- VNX Security and Compliance Suite—Keeps data safe from changes, deletions, and malicious activity.

Software packs available

- VNX Total Efficiency Pack—Includes all five software suites (not available for the VNX5100).
- VNX Total Protection Pack—Includes local, remote and application protection suites.
- VNX Total Value Pack—Includes all three protection software suites and the Security and Compliance Suite (the VNX5100 exclusively supports this package).

Prerequisites and supporting documentation

Technology

It is assumed that the reader has a general knowledge of the following products:

- An EMC VNX5300 platform
 - Microsoft SharePoint 2010
 - Microsoft Windows 2008 R2 Hyper-V
 - Metalogix StoragePoint 3.0
 - EMC Replication Manager 5.3 SP1
-

Supporting documents

The following documents, located on Powerlink.com, provide additional, relevant information. Access to these documents is based on your login credentials. If you do not have access to the following content, contact the EMC representative.

- *Externalizing Large SharePoint 2010 Objects with EMC VNX Series and Metalogix StoragePoint—Reference Architecture*
- *EMC Unified Storage for Microsoft Office SharePoint Server 2007 BLOB Externalization Enabled by EMC Celerra and Metalogix StoragePoint—Reference Architecture*
- *Celerra Network Server 5.5 Best Practices for Performance—Best Practices Planning*
- *EMC Solutions for Microsoft Office SharePoint Server with EMC Celerra Unified Storage Platforms—Reference Architecture*
- *EMC Unified Storage for Microsoft Office SharePoint Server 2007 BLOB Externalization Enabled by EMC Celerra and Metalogix StoragePoint—Proven Solution Guide*

The following document on the StoragePoint website (www.storagepoint.com) provides further information about StoragePoint administration:

- *StoragePoint Installation and Administration Guide version 2.2*
-

Terminology

Introduction

This section defines the terms used in this document.

Term	Definition
Automatic File System Extension	Extends a file system without operator intervention when a high water mark is reached.
binary large object (BLOB)	An unstructured binary data stream in SQL Server. The files in a SharePoint farm are stored in the database as BLOBs.
BLOB store	External storage location for BLOBs in the SharePoint farm or environment. For example, a CIFS share.
data deduplication	Only one unique copy of each file is stored if the file data appears more than once in the file system. The file is compressed to further improve storage efficiency.

Term	Definition
Data Mover	A Network Server cabinet component running the data access in real time (DART) operating system that retrieves files from a storage device and makes the files available to a network client.
Hyper-V	Microsoft Hyper-V Server 2008 R2 is a stand-alone product that provides a reliable and optimized virtualization solution and enables organizations to improve server utilization and reduce costs.
logical unit (LUN)	A LUN is a logical disk object presented from the storage array to a host. It is identified by a LUN number, which uniquely identifies the LUN for that host. The LUN number is not globally unique.
Microsoft SQL Server 2008 R2	A database application by Microsoft Corporation that stores the data of a SharePoint 2010 server.
Microsoft Visual Studio Team System (VSTS) 2008	Microsoft VSTS 2008 Test Edition provides a comprehensive suite of testing tools for Web applications and services that are integrated in the Visual Studio environment.
Orphaned BLOB	File in the BLOB file system that does not have a corresponding row in the SQL database.
RAID 1	A Redundant Array of Inexpensive Disks (RAID) method that provides data integrity by mirroring (copying) data to another disk. This RAID type provides the greatest assurance of data integrity at the greatest cost in disk space.
RAID 5	A RAID method where data is copied across disks in large stripes. Parity information is stored so that data can be reconstructed, if required. One disk can fail without data loss. Performance is good for reads but slow for writes.
Storage endpoint	Storage endpoint defines the location of the externalized BLOBs, such as storage area network (SAN), network-attached storage (NAS), content-addressed storage (CAS), or cloud storage.
storage processor (SP)	SP on a CLARiiON® storage system. On a CLARiiON storage system, a circuit board with memory modules and control logic that manages the storage system input/output (I/O) between the host's Fibre Channel (FC) adapter and the disk modules.
Storage profile	Storage profile dictates when and how content BLOBs are externalized. They can be web applications, site collections, or content databases scoped and optionally configured to compress or encrypt the externalized BLOBs.
Virtual Provisioning™	Allows a logical view of available capacity, which is independent of the actual physical capacity. Enables administrators to allocate capacity on-demand instead of allocating all the required space initially.

Chapter 2: SharePoint 2010 Farm Design

Overview

Introduction

Designing and implementing the layout for any environment is critical. Correcting layout errors can be expensive and time-consuming. Therefore, doing it right the first time should be the primary goal. This chapter explains the design process employed in building the SharePoint 2010 infrastructure for this use case. The information provided here can be used as a starting point to design and implement a similar environment.

This document does not provide complete planning details and architecture guidelines for a SharePoint 2010 farm. The customer should determine the set of portal sites, Internet presence sites, team sites, and specialized sites that the organization and its customers need.

For more information on SharePoint farm planning and architecture, refer to the books available on the Microsoft TechNet website.

Contents

This chapter contains the following topic:

Topic	See page
SharePoint farm design	12

SharePoint farm design

SharePoint 2010 farm components

In this use case, the SharePoint farm was designed for optimized performance, reduced bottlenecks, and ease of manageability.

The SharePoint farm consists of:

- A web application that was created by using the Enterprise Portal Collaboration template.
- Ten enterprise document center site collections created on 10 content databases.
- A service application associated with the web application.
- StoragePoint 3.0 software installed on the web front-end (WFE) server.

The following table describes the individual server roles in the SharePoint farm created for this use case.

Server role	Configuration details
SharePoint WFE server	<ul style="list-style-type: none">• The WFE server is deployed on a Hyper-V virtual machine hosted on a dedicated node of the Hyper-V failover cluster.• StoragePoint 3.0 is installed.• Internet Information Service (IIS) is configured to provide web content to SharePoint clients.• IIS web garden threads are set at the default value of one for optimal performance and ease of management.• IIS logging is disabled to limit the unnecessary growth of log files and to optimize performance.
SQL 2008 R2 database server	<ul style="list-style-type: none">• The database server is deployed on a Hyper-V virtual machine hosted on a dedicated node of the Hyper-V failover cluster.• SQL Server 2008 R2 Enterprise application server is installed.• The Fibre Channel (FC) LUNs from the VNX platform are attached to the virtual machine through the disk pass-through method.
SharePoint application server	<ul style="list-style-type: none">• The SharePoint application server is deployed on a Hyper-V virtual machine hosted on a dedicated node of the Hyper-V failover cluster.• A server that runs all the SharePoint application services including index crawling and search query services, and also hosts the SharePoint Central Administration website.• The Excel calculation and document conversion services are not tested as part of this use case. As a result, the services are inactive.• A dedicated FC LUN is attached through NTFS disk pass-through to store the content index (CI) files from the crawling operation.

SharePoint 2010 search configuration

The new Search server in SharePoint 2010 introduces new design considerations in the SharePoint farm.

In SharePoint 2010, a single SharePoint server can run Crawl and Query server roles. The Crawl services indexes the contents of the SharePoint farm, populates the crawl and property stores on the SQL database server, and then adds the CI files that will be used by the Query services.

In the test environment, 10 content databases were populated with 4,359,366 documents, which added up to 978 GB of data.

An FC LUN of 100 GB in size was created to store the CI files resulting from the crawl operations. A full crawl operation on the content databases created content index files of size 12.36 GB.

To use the virtual machine resources effectively and conveniently, the Crawl and Query server roles were enabled on a single virtual machine. The Query server role can also be enabled on the WFE server, which is also a valid configuration. In this case, another FC LUN that stores the content indexes will have to be provisioned to the Query server.

The Search server in SharePoint 2010 server is architected to provide greater redundancy in a single farm and to enable scalability in multiple directions. Each component that makes up the query architecture (query servers, index partitions, and property database) and the crawling architecture (crawl servers, crawl database, and property database) can be scaled out separately based on the needs of an organization.

The articles on the Microsoft website (www.microsoft.com) provide further information on scaling and sizing search architecture for the SharePoint 2010 server farm.

SharePoint databases and logs

The sequence of building a SQL environment for SharePoint 2010 requires three different subsets of databases in the following order:

1. SQL internal databases that are created during installation (master, model, msdb, and tempdb)
2. SharePoint databases created from the administration web page during the creation/configuration of the SharePoint farm.
3. SharePoint content databases (manually created).

The SharePoint configuration and content databases and log files were placed in the FC LUNs configured on the VNX series. In this solution, the databases are stored on a single 1.5 TB FC LUN provisioned from a RAID 5 (4+1) group of SAS disks. The log files use a 100 GB FC LUN provisioned from a RAID 1 (1+1) group of SAS disks.

**Master, model,
and msdb
databases**

The master, model, and msdb SQL system databases are created when SQL Server is installed. The master database records all the system-level information for a SQL server system. It records all login accounts and system configuration settings. The model database is used as the template for all databases created on a system. The msdb database is used by the SQL Server agent to schedule alerts and jobs, and to record operations. These databases rarely grow beyond 1 GB cumulatively. For this use case, these databases were stored in the default location on the local SQL Server system drive.

Tempdb

The tempdb database was placed on FC LUNs that were created on two separate spindles configured with RAID 1. Performance testing on this medium-size SharePoint farm did not show much activity on these temporary database files. About 1 GB out of 100 GB of allocated space was used by the temporary database files.

Chapter 3: Virtualization

Overview

Introduction

This chapter provides guidelines to install and configure the Hyper-V virtual environment that makes up the validated solution scenario.

Microsoft Hyper-V is a virtualization solution that enables users to turn their infrastructure into an efficient and flexible internal cloud. Therefore, this solution enables users to:

- Decrease their capital and operating costs by over 50 percent
- Run a greener data center and reduce energy costs
- Control application services levels with advanced availability and security features
- Streamline IT operations and improve flexibility

Microsoft Hyper-V is a key feature in Windows Server 2008. Hyper-V can transform or virtualize the hardware resources of an x64-based computer, including the central processing unit (CPU), random access memory (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.

Scope

The virtualization guidelines presented in this chapter apply to the specific components used during the development of this solution.

Contents

This chapter contains the following topics:

Topic	See page
Concepts	16
Advantages of virtualization	16
Considerations	17
Implementation	17

Concepts

Virtualization layer The virtualization layer abstracts the processor, memory, storage, and network resources of a physical server into multiple virtual machines. This enables multiple operating systems to run simultaneously and independently on a single physical server.

Pass-through disk Using a pass-through disk, a Hyper-V host disk can be exposed to the guest virtual machine without even placing a volume on it. Hyper-V will enable you to avoid the file system of the host and access a disk directly.

Cluster shared volume Cluster shared volume provides a single, consistent storage space that enables the hosts in a cluster to concurrently access virtual machine files on a single shared LUN. It enables the live migration of virtual machines without impacting other virtual machines sharing the same LUN.

Cluster shared volume eliminates the restriction of placing one virtual machine on one LUN and coordinates the use of storage with greater efficiency and higher performance.

Advantages of virtualization

Reduced costs One of the main challenges faced by the customer is to reduce costs by using the infrastructure effectively. A virtualized SharePoint environment reduces the cost by using the infrastructure effectively. Virtualization enables a reduction in the number of physical servers that are required to implement a SharePoint environment.

Reduced downtime A running virtual machine production database can be moved from one physical server to another physical server without any downtime.

Performance and scalability In a scale-out context, virtualization can provide superior performance and scalability when compared to physical server configurations, even when using hardware that is identical to that used in the physical server configuration.

Ease of use The single user interface enables administrators to monitor and manage multiple virtual machines from one console. Therefore, with virtualization, administrators can easily and conveniently manage virtual machines rather than physical servers.

Considerations

Hosting multiple servers

It is important to consider all aspects of the solution before virtualizing the environment. When hosting multiple virtual servers on a single physical server, consider the following:

- Do not place all high-priority virtual servers on a single physical server.
 - Do not place processor-hungry and memory-hungry virtual machines on a single physical server.
 - Configure high-availability features for disaster recovery.
-

Implementation

Hosting multiple servers

In this solution, a three-node Microsoft Windows 2008 Hyper-V cluster was used. The application virtual machines were hosted on this cluster.

To simplify management and resource planning, it is recommended to cluster physical systems with similar, if not identical, hardware. In this case, three identical Dell PowerEdge 2950 systems with the following specifications were used:

- Four 3.00 GHz Intel Xeon processors
- 16 GB of memory
- One 60 GB 15k internal SCSI disk
- Two on-board 10/100/1000 MB Ethernet NICs
- Four additional 10/100/1000 MB Ethernet NICs
- One dual-port QLogic HBA card

The following table shows the Hyper-V cluster at the production site and the initial placement and configuration of the virtual machines.

Server	Virtual machines	Virtual machine configuration
Hyper-V server 1	WFE server	<ul style="list-style-type: none">• 4 vCPUs• 8 GB memory• 2 virtual NICs
Hyper-V server 2	SQL Server 2008 R2	<ul style="list-style-type: none">• 4 vCPUs• 16 GB memory• 1 virtual NIC
Hyper-V server 3	SharePoint application server	<ul style="list-style-type: none">• 4 vCPUs• 8 GB memory• 1 virtual NIC

Storage considerations

Cluster-shared volumes provide a single, consistent storage space that enables hosts in a cluster to concurrently access virtual machine files on a single shared LUN. All application boot volumes are utilized as cluster-shared volumes. It enables the live migration of virtual machines without impacting other virtual machines sharing the same LUN. In this solution, the storage for all the application data and the backups are utilized as Hyper-V pass-through disks.

Network considerations

The network settings were configured identically on all nodes of the cluster and on all virtual machines. The WFE server was additionally configured with a dedicated NIC to access the CIFS share used as the BLOB store file system.

Chapter 4: Network Design

Overview

Introduction This chapter describes the network architecture of the EMC VNX5300 platform and Microsoft SharePoint server virtual machines.

Scope System-wide network design and architecture are outside the scope of this document and solution. This chapter presents recommendations for proper functionality that are consistent with industry-accepted best practices and should be compatible with the existing network infrastructure and policies.

Contents This chapter contains the following topics:

Topic	See page
Considerations	20
Implementation	20

Considerations

Physical design considerations EMC recommends switches that support the fastest and most reliable connection methodologies that are currently available.

Logical design considerations This validated solution uses VLANs on the IP network and soft zones on the FC network to segregate network traffic of different types and to improve throughput, manageability, application separation, high availability, and security.

Consider the following VLANs that are recommended::

- A client VLAN network that supports connectivity among the SharePoint servers, Replication Manager Server, and the client workstations.
- A file share VLAN for connectivity between the EMC VNX platform and the SharePoint WFE server to provide dedicated access to the CIFS share used as a BLOB store file system.
- A management VLAN that supports connectivity to the Hyper-V hosts for server administration and to the EMC VNX5300 platform for storage management.

The FC network zoning of the Hyper-V Cluster nodes and the VNX series should be consistent with the EMC best practice for single initiator zoning.

Implementation

Physical design implementation The 1 Gb Ethernet IP network is designed using Cisco Catalyst 6509 switches, while the FC storage network is designed using Brocade DS5100 4 Gb FC switches.

Logical design implementation Both the recommended VLANs are configured on the IP network. The FC network is set up for single initiator zoning, which is consistent with the industry best practices.

VNX for file network configuration The VNX for file consists of two Data Movers. These Data Movers can be configured in an active/active or active/passive configuration. In the active/passive configuration, the passive Data Mover serves as a failover device for the active Data Mover. In this solution, the Data Movers operate in the active/passive mode.

The VNX5300 Data Mover was configured for four 1 Gb interfaces on a single SLIC. Only a single port cge0 was used to support the SharePoint BLOB access on the CIFS share. Ports cge1, cge2, and cge3 are left free for further expansion.

**Fibre Channel
network
configuration**

A Brocade DS5100 FC switch was used to provide connectivity for storage network in this solution. Each node of the Hyper-V cluster has a single connection to the fabric. Each storage processor has two links to the SAN fabric for a total of four available front-end ports. The zoning is configured so that each server has four available paths to the storage array.

Chapter 5: Storage Design

Overview

Introduction Storage design is an important element to ensure the successful deployment of the SharePoint 2010 BLOB Externalization solution by using EMC VNX and Metalogix StoragePoint to protect the SharePoint databases and the BLOB store.

Scope The storage design layout instructions presented in this chapter apply to the specific components used during the development of this solution.

Contents This chapter contains the following topics:

Topic	See page
Design considerations	23
Storage design implementation	24

Design considerations

Overview

When planning storage for the SharePoint BLOB externalization solution, the storage is designed in two parts:

- Storage layout planning for production databases and BLOBs—When planning for databases and BLOBs, select the right configuration with optimum performance and higher capacity.
 - Storage layout planning for the backup destination—When planning for replication destinations, plan a storage capacity that is equal to the production databases and BLOB stores.
-

Performance

Many customers gather data while the application is running, and then use a 90th percentile to determine the level that should be planned for. The four primary variables used to determine the number of spindles for storage are:

- IOPS (or sometimes MB/s, if it is a serial workload)
 - Latency goals based on the application requirements
 - RAID level—When planning for performance, striped RAID 1 require fewer spindles than RAID 5 for almost all read/write workloads. They are approximately equal in a read-only workload.
 - Drive type—The drive type can dramatically decrease or increase the number of drives required to satisfy the workload. As a general best practice, database-type applications are hosted on FC drives.
-

Capacity

SharePoint users continually add content to their farms and this places limits on how much space is available to new users. Administrators must plan ahead for these capacity issues.

One of the advantages of this solution is that storage demands are moved away from SQL Server and placed with the VNX array, which enables the use of advanced features such as deduplication, Virtual Provisioning, and Automatic File System Extension. These features save space and enable an incremental growth pattern for the addition of spindles in the array. Storage administrators can begin with an initial storage configuration and confidently grow their arrays as necessary when the users increase their storage demands.

With Automatic File System Extension in place, administrators are not required to act as VNX extends the file system when a configurable percentage of available space has been consumed. The capacity is allocated on demand by VNX only to those file systems that reach a high water mark and need additional disk space. Deduplication removes duplicates and compresses files to slow the use of available space in the file systems. This delays the additional equipment needs in the current storage setup.

Storage design implementation

Introduction

The VNX platform stores the following:

- SharePoint Virtual Machine Virtual Hard Disks (VHD) files, index files, databases, logs, and tempdb
- BLOBs offloaded from SQL databases to the BLOB store file system, which is accessed through the CIFS share
- Snapshots of SharePoint databases and a replica of the BLOB store file system

SharePoint VHD and index files: A RAID 5 (4+1) group of spindles stores VNX system volumes, virtual machine VHD files, and index files.

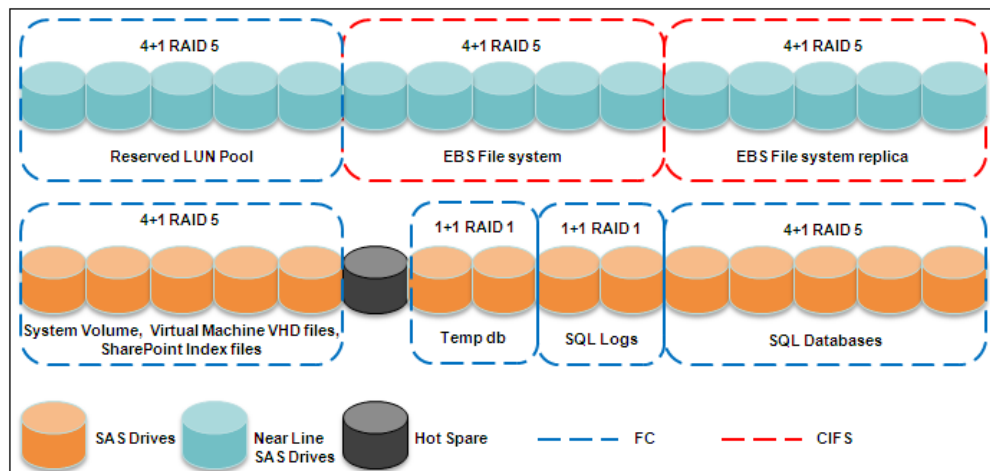
SharePoint logs and tempdb: The configuration tested in this solution uses RAID 1 (1+1) groups to store tempdb and log files.

SharePoint databases and the BLOB store file system: The validated solution uses a RAID 5 (4+1) group of SAS spindles for SQL databases and a RAID 5 (4+1) group of NL-SAS drives to store Externalized BLOB content.

Replica of the BLOB store file system: A RAID 5 (4+1) group of NL-SAS spindles to store the replica of the BLOB store file system.

Reserved LUN Pool: A RAID 5 (4+1) group of NL-SAS spindles to store the reserved LUN pool used by SnapView snapshots.

The following diagram shows the disk layout of the storage design.



Performance

To increase the performance and protection of data due to disk failures, the tempdb, logs, and databases reside on separate spindles.

Capacity

The database disks are set up on RAID 5 groups to increase the storage space for the SharePoint solution.

Chapter 6: SharePoint Data Backup with Replication Manager

Overview

Introduction

Customers are managing their IT environments by ensuring that it meets the corporate directives and adheres to strict service-level agreements. Because SharePoint environments have become business-critical applications, it is essential to implement an efficient backup solution.

The presence of external BLOBs in the SharePoint infrastructure leads to changes in the current backup and restore strategies. These changes can easily be handled by using Replication Manager to take snapshots and replicas of SQL databases and the BLOB store file system.

SnapView and Replicator functionalities are used in this solution to protect SQL databases and the BLOB store created on the VNX file system.

Scope

This solution provides options to create SQL-consistent snapshots of the SQL databases and the BLOB store file system. But it does not protect the contents of the SharePoint WFE server such as the Internet Information Server (IIS) metabase, virtual directories of portal sites, and the search index partitions residing in the Index and Query servers.

Contents

This chapter contains the following topics.

Topic	See Page
Design considerations	26
Design implementation	26
Best practices	28

Design considerations

Performance The backup operation must not unduly affect the production workload. Data backups must occur within the recovery time objective (RTO) and the recovery point objective (RPO) windows set by administrators.

Capacity There must be adequate space for all active data in the environment to be replicated on the target. The size requirement for snapshots depends on the data size, frequency of snapshots, number of retained snapshots, and rate of change in data.

Manageability Administrators must be able to schedule replication jobs and quickly perform recovery operations to recover data, if required. Backup jobs should be automated to save time and effort for administrators. Compliance issues should be addressed by the protect solution so that there is no missing information as a result of the data recovery operations.

Design implementation

Performance Generally, replication jobs do not significantly affect the performance of SharePoint server. However, it is not recommended to run replication jobs during heavy loads on the server because they will extend the replication window further.

In this solution, replication jobs were scheduled depending on the network over which the replication was done. For a network with low latencies, it takes less time to complete a replication update. Therefore, the jobs were scheduled at frequent intervals.

Because this solution uses in-frame (in the same VNX platform) file system replica and snapshots as a backup solution, administrators can schedule more frequent snapshots and replica jobs.

Capacity In this solution, the SQL databases and logs were backed up using SnapView snapshots. Hence, space for the reserved LUN pool was configured to accommodate 10 percent of data change. The VNX file system that was used as the BLOB store was replicated by using a Replicator job to an equal-sized read-only file system.

Manageability Replication Manager provides an easy-to-use GUI to manage, schedule, and monitor replication jobs. It also provides the flexibility of scripting using command line interface (CLI).

Backup and recovery considerations StoragePoint involves an additional step when backing up content because the BLOB store also needs to be backed up. At first, this may appear cumbersome but it actually creates opportunities that were not available to SharePoint administrators earlier.

Externalizing the content to a BLOB store by using StoragePoint reduces the content databases by at least 90 percent. This can address some of the item-level recovery challenges that SharePoint administrators face. SharePoint has a built-in mechanism for item-level recovery through the Recycle Bin feature. There are two levels of recycle bins. Level one is allowed for users and level two is controlled by the administrator.

When a user deletes a document, it is available in the recycle bin of the site for a period defined by the administrator. By default, it is 30 days. Without StoragePoint, the administrators will not increase the retention period because of the increase in the size of the content databases. However, with StoragePoint, a large amount of the available content database space can be used to set a higher retention period.

An additional option provided by the StoragePoint provider is the orphan BLOB retention policy. This setting dictates how long the StoragePoint provider retains the content BLOB files in the BLOB store after the SharePoint list item or document associations are removed from the databases.

BLOB store replica considerations

Replication Manager 5.3 enables the creation of the application-consistent snapshots of SQL databases. However, the replicas of VNX file systems used as BLOB stores cannot be created or managed from the Replication Manager GUI.

In this solution, the VNX file system local replication was set up by using the Unisphere Manager GUI. The following figure shows the settings used for the asynchronous file system replication to protect the BLOB store. The administrator needs to perform this replication setup only once. After the replication session is activated, the replica file system is updated according to the update policy set by the administrator.

Replication Name:	replication_blob_fs
Source File System:	blob_fs (2048000 MB)
Source File-level Retention:	Off
Storage System:	CLARiiON C125 FNM00103200530
Source Checkpoint Storage Pool:	cifs_blob_pool 5.2 TB (5457274 MB)
Destination Celerra Network Server:	rtpsol41a
Source Data Mover Interconnect:	loopback
Source Interface:	127.0.0.1
Destination Interface:	127.0.0.1
Create Destination:	<input type="radio"/> Create Using Storage Pool <input checked="" type="radio"/> Use Existing File System
Destination File System:	replica_blob_fs (2048000 MB)
Storage System:	CLARiiON C125 FNM00103200530
Destination Checkpoint Storage Pool:	fs_replica_pool 5.2 TB (5457274 MB)
Update Policy:	<input checked="" type="radio"/> Time Out of Sync <input type="radio"/> Manual Refresh
Time Out of Sync:	1 Minutes
Discard Changes on Destination Since Last Copy:	<input type="checkbox"/>
Use Tape Transport:	<input type="checkbox"/>

As shown in the figure, **Time Out of Sync** was selected as the **Update Policy** for the file system replication with an interval of 1 minute. This means that the destination file system was updated after every 1 minute with all the content BLOBs of the production BLOB store. Though the BLOB contents were synchronized with a time period of 1 minute, the actual RPO for the SharePoint environment is determined by the frequency of the SQL database snapshot job. In this solution, after file system replication was activated, the initial silvering of 942 GB of data in the file system was completed in 3 hours and 38 minutes.

Best practices

Use the latest available VNX Operating Environment (OE)

Install the latest VNX release code or patch to take advantage of new features, functionalities, and bug fixes.

Plan storage operations for minimal disruption

Some operations on the storage array can consume resources and may affect the production system if executed during heavy production loads. For this reason, storage-based operations that may consume resources must be scheduled during the nonpeak hours to minimize the potential for such occurrences.

Chapter 7: Testing and Validation

Overview

Introduction This chapter presents testing information that is common to all the tests presented in the document. It outlines the test tools, methods, workload used, common setup procedures, and architectural considerations.

Contents This chapter contains the following topics:

Topic	See page
Testing approach	30
Methodology	31
Testing overview	33
Test results analysis	34
BLOB externalization	34
SharePoint indexing results	35
Benefits of BLOB externalization	36
Impact of BLOB externalization	37
VNX5300 performance summary	40
Backup and restore of the externalized SharePoint farm	43
Using EMC file deduplication and compression with externalized BLOBs	47

Testing approach

Introduction To properly test the SharePoint 2010 environment, a set of tools designed by KnowledgeLake (a Microsoft Gold Certified Partner) was used. KnowledgeLake has its own data population tool that uses the SharePoint front end to insert unique documents into the SharePoint farm.

Data population The data population tool uses a set of sample documents. The document type and size are listed in the following table. Altering the document name and metadata (before insertion) makes each document unique. One load agent was used at the WFE server to load 978 GB of content data. The data was spread evenly across 10 sites (each site was created on a unique content database).

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

Load generation To generate and emulate the client load, Microsoft Visual Studio Team Suite (VSTS) was used in conjunction with the KnowledgeLake code to simulate realworld SharePoint user activity.

User operations A mix of three user operations such as browse, search, and modify was generated for SharePoint testing.

Browse

In a browse test, the code simulates a user browsing a site until the user reaches the end of the document listing, which does not contain sub-pages.

Search

In a search test, the code simulates a user running a stored procedure in the SQL database to find a unique number. In this case, the unique number is the Social Security number (SSN). The code then performs a web request to search for that unique number.

Modify

In a modify test, the code simulates a user retrieving a document. The document name is extracted from the database before each test run. The code then modifies the metadata for that document before saving it back to the farm in its modified form.

VSTS counters Several counters were measured to provide the time for each user action and an accurate indication of the user experience. Initially, the Response Time counter was

used, which is an average measure of all the page contents being loaded in parallel. From a user's point of view, the average time of all components is not the time a page takes to load. It was decided to use the Page Time counter, which is the time required for all page components to load. While this discounts parallel page loading, it illustrates a negative user experience scenario. This shows the slowest possible response that a user can experience.

The Microsoft article titled *How to: View Web Page Response Time in a Load Test* explains how the different counters calculate response times in the Page Time Warning section. Using Average Page Time was ideal for the search and modify profiles, but it could not be used for the browse profile. VSTS cannot average the web pages; it averages all components. The two options to calculate the browse time are to use either Average Total Time or Average Test Time counter.

Average Test Time is the time taken to send a request and complete the full test, including the computation time. It is slightly slower than the Average Page Time due to the computation time, but the impact is so minor (about 1 millisecond) that the counter was accurate enough for testing purposes.

Methodology

Introduction

This section explains the testing methodology used.

Testing methodology

The data population tool from KnowledgeLake was used to insert a set of sample documents into the customized document libraries of the 10 document center site collections created in the SharePoint farm. The data was spread evenly across 10 sites (each site was created on a unique content database).

To generate and emulate the client load, Microsoft VSTS 2008 was used in conjunction with the KnowledgeLake code to simulate real-world SharePoint user activity. The VSTS team test rig in the testing environment consisted of a single controller and an agent on a single host machine. The following steps describe the testing procedure used with the VSTS tool:

Step	Action		
1	In the VSTS test editor, create a scenario with a particular test mix of browse, search, and modify operations.		
2	Run the VSTS load test with an initial value for the maximum user count in a stepwise load pattern for one hour.		
3	During the test, monitor the WFE server CPU usage and the average test time for browse, search, and modify test modules.		
4	After the test completes, check whether the average test time values are within the Microsoft-defined acceptable range as shown in the following table and whether the average CPU usage of the WFE server is within the threshold of 75 percent. If these two conditions are met, increase the maximum user count by one and go to step 2, or else proceed to step 5.		
	Type of operation	Example	Acceptable user response time*
	Common	Browse	< 3 seconds
	Common	Search	< 3 seconds

	Uncommon	Modify	< 5 seconds
	* As per the Microsoft-recommended performance guidelines		
5	Stop the test runs when one of the following conditions are met: <ul style="list-style-type: none"> • Average test time values reach the acceptable limits. • Average CPU utilization of the WFE server reaches 75 percent. 		

Similarly, run the tests for the other two user profile mixes. After the test runs are complete, the VSTS counter (passed tests/sec), which is equivalent to the farm throughput called requests per second (RPS), is used to calculate the maximum users supported by the SharePoint farm.

Estimate throughput targets

Throughput is the number of operations (browse, search, or modify) that a server farm can perform each second. Ideally, the number of operations that are requested per second is lower than the number targeted for a given level of performance. If the number of operations requested exceeds the target number, user actions and other operations will take longer to complete.

Throughput is measured in RPS. RPS measurements can be converted to the total number of users by using a model of typical end-user behavior. Like many human behaviors, there is a broad range for typical behavior. The user model for SharePoint environment has the following two variables:

- Concurrency—The percentage of users that are actively using the system.
- Request rate—The average number of requests per hour that an active user generates.

Four levels of user load are shown in the following table.

User load	Request rate	Supported users
Light	20 requests per hour. An active user generates a request every 180 seconds.	Each response per second of throughput supports 180 simultaneous users and 1,800 total users.
Typical	36 requests per hour. An active user generates a request every 100 seconds.	Each response per second of throughput supports 100 simultaneous users and 1,000 total users.
Heavy	60 requests per hour. An active user generates a request every 60 seconds.	Each response per second of throughput supports 60 simultaneous users and 600 total users.
Extreme	120 requests per hour. An active user generates a request every 30 seconds.	Each response per second of throughput supports 30 simultaneous users and 300 total users.

Calculate maximum user capacity

The following calculation determines the maximum user capacity of a SharePoint farm for a typical user load:

$$\frac{\text{Typical user RPH}}{\text{Seconds per hour}} = \text{Typical user RPS}$$

$$\frac{\text{Farm RPS}}{\text{Typical user RPS}} = \text{Simultaneous users}$$

$$\frac{\text{Simultaneous users}}{\% \text{ concurrency}} = \text{Maximum user capacity}$$

For example:

$$\frac{36}{3600} = 0.01$$

$$\frac{5.67}{0.01} = 567$$

$$\frac{567}{0.10} = 5,670$$

The following table shows a sample calculation of the maximum user capacity of a SharePoint farm for different user loads with RPS as 5.35 and concurrency as 10 percent.

User load	Requests per hour (RPH)	Concurrency	Maximum user capacity
Light	20	10%	10,206
Typical	36	10%	5,670
Heavy	60	10%	3,402
Extreme	120	10%	1,701

For further information about how to determine maximum user capacities, refer to the Microsoft document *Estimate performance and capacity requirements for Windows SharePoint Services collaboration environments (Office SharePoint Server)* available on the TechNet Microsoft website.

Testing overview

Tested scenarios

The test series was designed to understand the behavior of SharePoint 2010 with the externalization of BLOBs from SQL databases to external CIFS shares on the VNX platform.

The tested scenarios can be broadly classified under the following categories:

- Study externalization and its impact on Content Indexing
- Study the benefits of externalization on:
 - SharePoint database management
 - Scalability of the SharePoint farm
 - Performance for different user profile load
- Backup and restore of SharePoint data
- Study the benefits of leveraging storage array features

Test results analysis

Key results summary

Some of the important results that were obtained after validating this solution are:

- Moved all the BLOBs from the content databases, which is up to 92.2 percent of total content database size, to a high-capacity storage tier with no loss of performance.
 - Increased scalability of SharePoint data by 65-70 percent approximately. The usable space is increased from 2.4 TB for Content and Configuration databases to 2.4 TB for the databases and additional space of 4 TB for SharePoint content or BLOBs without performance degradation.
 - Simplified SharePoint database management by allowing fewer and smaller content databases to manage, while adhering to Microsoft's guidelines.
 - Improved storage efficiency by compressing the externalized data and saving 20 percent of capacity with minimal degradation in performance.
 - Enabled rapid recovery of the SharePoint farm by using the advanced snapshot technology while maintaining database consistency.
-

BLOB externalization

Summary

Before externalization

In the test environment, 10 content databases were populated with 4,359,366 documents aggregating to 978 GB of data. Each content database contained up to 97.8 GB of data. The content databases in SharePoint stored files, metadata (information about the files), context information (file location, security information, and so on), and other structured relational data in its tables.

After externalization

After externalization using the StoragePoint 3.0 provider, all the BLOBs in the SQL content databases were offloaded to the VNX CIFS share except for the metadata, context information, other structured data, and reference IDs of BLOBs in the content databases. After completing the externalization process, 925.12 GB (94.59 percent content database data) was offloaded to the BLOB store as shown in the following figure.

Central Administration > StoragePoint Dashboard

StoragePoint® I Like It Tags & Notes

Farm Hierarchy

- Farm - SharePoint_Config
 - Web App. - SharePoint - 80
 - Web App. - SharePoint - website.nas.c
 - Content - WSS_Content_cb1e19556c
 - Content - WSS_Content_5c4e92f344
 - Content - ContentDB01
 - Content - ContentDB02
 - Content - ContentDB03
 - Content - ContentDB04

Storage Profile Filters Select All

- websitenascom (WebApplication)

StoragePoint Health and Metrics

Total BLOBs Externalized

4359366

Total Raw Size (MB)

947,326.81

Average BLOB Size (KB)

222.52

By Endpoint

Endpoint	Total BLOBs	Raw Size (MB)	Average Size (KB)	Compressed Size (MB)	Avg. Compressed Size (KB)	% Compressed
CIFS_share	4359366	947,326.81	222.52	N/A	N/A	N/A
SYSTEM CACHE	45275	12,191.61	275.74	N/A	N/A	N/A
Totals	4404641	959,518.42	223.07	0.00	0.00	

Close

After an externalization job is completed, it is strongly recommended to run the following steps to reclaim space from the content databases and rebuild the tables and indexes:

1. Execute DBCC_shrinkdb command (if required, execute this step multiple times to fully reclaim unused space)
2. Defragment tables
3. Rebuild indexes
4. Update statistics

In addition to these steps, once the space from the content databases is reclaimed, the storage administrators can also reclaim space from the VNX LUNs that were provisioned for the SQL database files. This process should be done during nonpeak periods to avoid performance degradation for SharePoint users. *EMC CLARiiON LUN Shrinking with Microsoft Exchange 2007 - Applied Technology* available on Powerlink provides information on the LUN shrinking process and the impact on performance.

SharePoint indexing results

Full farm indexes

The following table lists the duration of a full crawl on the SharePoint farm.

Storage configuration	Full index duration (hh:mm)	Items in index
BLOB store on NL-SAS disks	34:55	4,361,346

The indexing performance is determined by the time taken to index the items in the SharePoint farm, which will be used later for processing search queries.

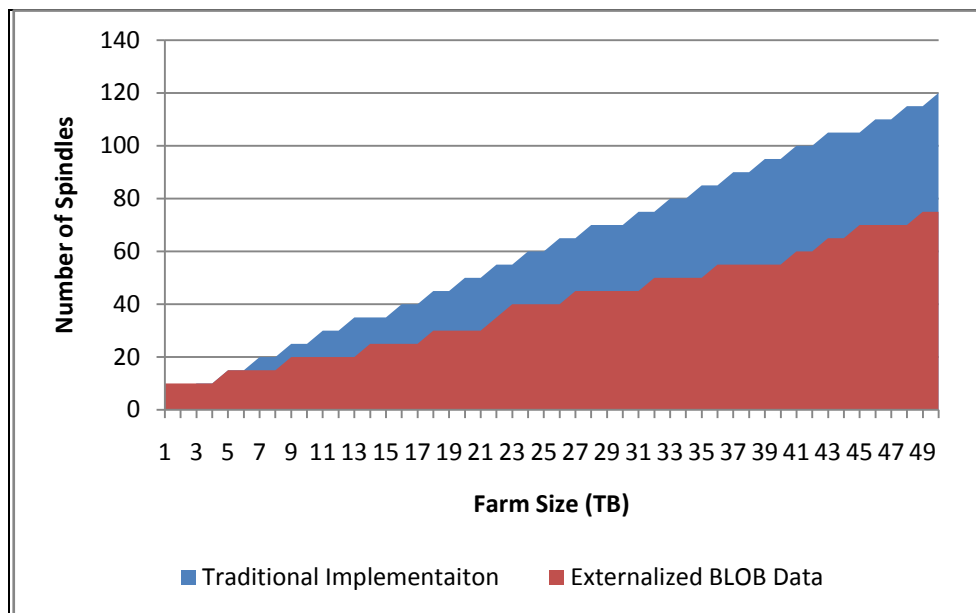
Benefits of BLOB externalization

SharePoint database management

If BLOBs are offloaded from SharePoint databases to the external file system, almost 92 percent of the content database space is freed and available for database administrators to reclaim. The content databases now contain only metadata information of the files because the actual files are stored on an external file share. Consequently, the growth of the content databases will be very less compared to the external storage. With smaller databases, administrators spend less time running maintenance operations on the databases, and hence reduce the downtime of database availability. Additionally, with a BLOB store provider, database administrators do not have to worry about performance degradation due to the 200 GB recommended limit on the content database size. Thus, using a BLOB store provider in SharePoint environments highly simplifies database management and increases the scalability of the SharePoint farm.

SharePoint farm scalability

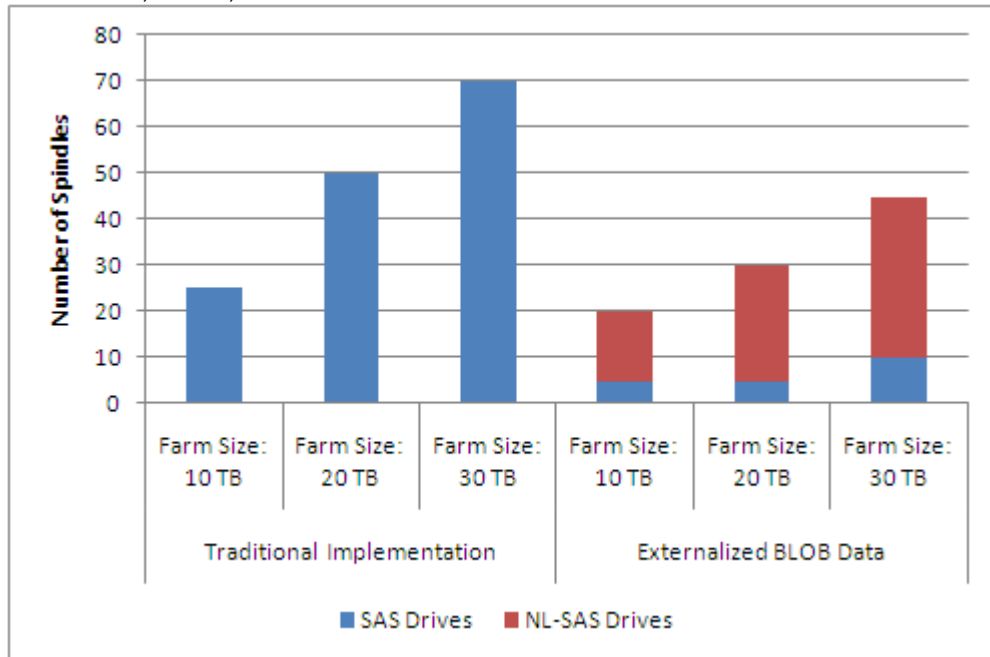
The scalability of the SharePoint farm data increases by externalizing SharePoint BLOBs to high-capacity drives. Because the SharePoint environments generally have a high growth rate of data, the externalization solution when compared to the traditional implementation can accommodate a large data growth by adding fewer spindles at the back-end storage system. The following graph shows the number of spindles or drives that are required for a SharePoint farm in an externalized solution when compared to the traditional solution using a VNX5300 platform.



This graph compares the number of spindles that are required to build a traditional SharePoint solution with 600 GB SAS drives and an externalized solution using 1 TB NL-SAS drives on a VNX5300 platform. Based on the graph, as the SharePoint farm size increases, the externalized solution requires approximately 40 percent less spindles at the back end.

The following graph shows the comparison of the number of spindles required in a traditional SharePoint solution using SAS drives and an externalized solution using SAS drives and NL-SAS drives. The comparison is shown for SharePoint farms of

sizes 10 TB, 20 TB, and 30 TB.



Impact of BLOB externalization


Introduction

This section presents the test results of the SharePoint 2010 BLOB externalization solution using the Celerra NS-480 storage system to show that the externalization of BLOBs in the SharePoint farm does not negatively affect the user experience.

This testing has been completed before the release of the VNX series, but the solutions engineering team believes that the results are representative and the conclusion remains valid when applied to the new platform.

Throughput (RPS) comparison

The following graph shows the comparison of RPS values that were obtained during the performance tests on a SharePoint 2010 server with and without BLOB externalization using the Celerra NS-480 storage system.




Conclusion

Based on the results, the following can be concluded:

- BLOB externalization using the StoragePoint software did not have a negative impact on the performance of the SharePoint farm. Moving files from SQL databases to a dedicated file share on a separate network increased the I/O performance of the files and the overall throughput of the SharePoint farm.
- When the baseline configuration is compared with the BLOB externalization on SATA disks of Celerra NS-480, there is an 8 percent and 10 percent increase in the throughput for the 80/10/10 and 70/05/25 profiles, respectively, and an 8.5 percent decrease in the throughput for the 50/20/30 profile.

Response time comparison

The following graph shows the response time for browse, search, and modify operations that were obtained during the performance tests on a SharePoint 2010 server with and without BLOB externalization using the Celerra NS-480 storage system.




Conclusion

The response time for browse, search, and modify operations at the maximum RPS value remained well within the acceptable limits set by Microsoft for SharePoint users. Generally, the browse operation, which is the most common SharePoint operation, had a faster response time after externalization, whereas the search and modify operations took a slightly longer amount of time.

Comparison of the CPU utilization of SharePoint servers

The following graph shows the comparison of the average CPU utilization of servers at the maximum RPS value obtained during the performance tests on SharePoint 2010 server with and without BLOB externalization using the Celerra NS-480 storage system.



Conclusion

The CPU utilization of the WFE server remained approximately constant in the SharePoint environment with and without BLOB externalization. However, the SQL server and application server CPU utilization was higher with BLOB externalization enabled. This is because with BLOB externalization enabled, the SQL server CPU is effectively used for SharePoint operations to achieve the optimum SharePoint farm throughput.

VNX5300 performance summary

Summary

This section covers the performance test results for the externalized SharePoint farm with the following user profile mixes:

- 80% browse/10% search/10% modify
- 70% browse/5% search/25% modify
- 50% browse/20% search/30% modify

The following table summarizes the user profile test results. For user profiles under normal conditions, the following table shows the RPS values that were obtained for the three user profiles along with the average user response times for browse, search, and modify operations. The maximum user capacity was determined by using a defined concurrency of 10 percent for all three user profiles.


User profile (browse/search/modify)	User load profile (36 RPH)	Requests per second (RPS)	Concurrency (%)	Maximum number of users	Average user response time in seconds (browse/search/modify)
80/10/10	Typical	5.67	10	5,670	2.35/2.0/1.82
70/05/25	Typical	7.91	10	7,910	1.78/1.45/1.74
50/20/30	Typical	8.63	10	8,630	1.95/1.69/2.14

In this configuration, the SharePoint databases were stored on an FC LUN that was created by using one RAID 5 (4+1) configuration of SAS disks, and the BLOBs were externalized to a VNX CIFS share that was created by using one RAID 5 (4+1) configuration of NL-SAS disks. The BLOBs in the SharePoint databases were externalized by using the StoragePoint provider. Subsequently, the SharePoint content databases contained metadata, context information, and reference IDs for the externalized BLOBs (which are very small in size).

The SharePoint user load was simulated on the farm to determine the RPS value keeping the average user response time for browse, search, and modify within the acceptable limits specified by Microsoft.

RPS and average user response time

The following graph shows the performance results of a SharePoint 2010 server for the three different user profiles.




Conclusion

From the tests, the externalized SharePoint farm throughput was determined. During the tests, it was observed that the CPU usage of the WFE server was a bottleneck. Therefore, when the browse test percentages increased, the RPS value decreased. The throughput was the highest for the 50/20/30 profile.

Processor utilization of SharePoint servers

The following graph shows the saturation user load and the processor utilization of SharePoint servers for the different user profile loads.




Conclusion

The VNX storage processor CPU utilization was less than 10 percent and the Data Mover CPU utilization was less than 5 percent through all the user profile tests and hence, it is not shown. This also indicates that this parameter is not a bottleneck for the SharePoint farm performance.

The SQL server CPU utilization increased from 20.4 percent in the 80/10/10 baseline test to 51.6 percent in the 70/05/25 baseline test to 63.8 percent in the 50/20/30 baseline test. This is mainly due to an increase in the percentage of document modify operations. Hence, environments that are highly collaborative and involve sharing documents need greater processing power for the SQL database server.

Disk latency and IOPS

The following graph shows the SharePoint database disk latency and IOPS observed during the three different user profile loads.



Conclusion

The IOPS generated by the 50/20/30 user profile mix was higher than the other profiles due to a high percentage of document modify operations. The average disk latencies of database, log, and tempdb disks for read and write operations were less than the acceptable limit of 20 ms for all the user profile tests. This suggests that the storage system adequately handled the I/O requests generated by the SharePoint user profile load.

Backup and restore of the externalized SharePoint farm

Introduction

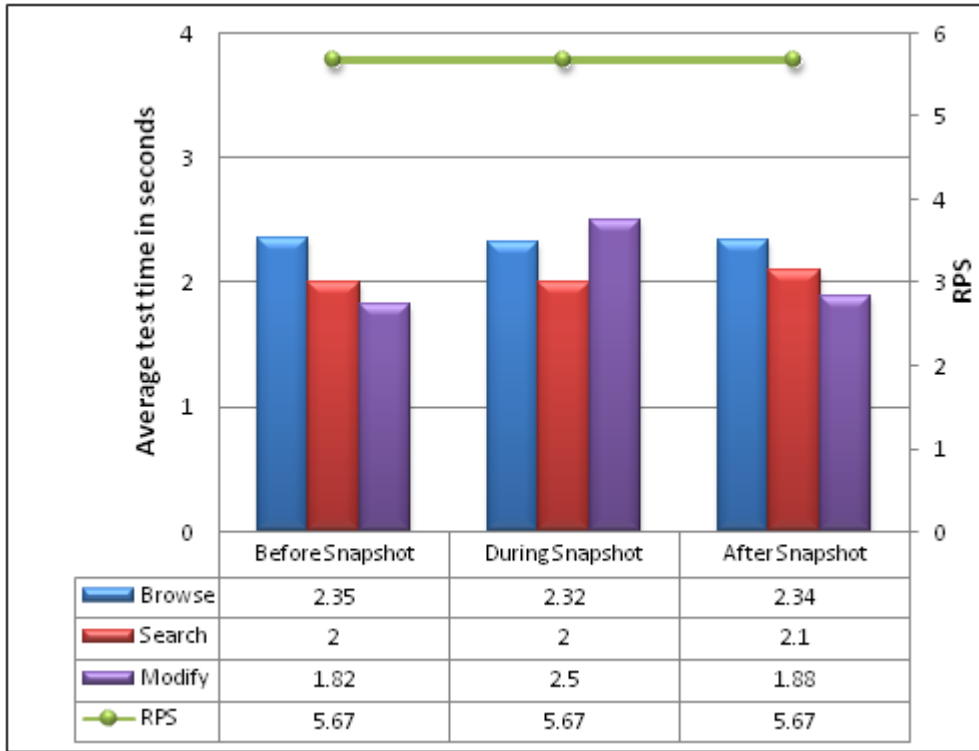
Offloading BLOBs to a separate file system will help SharePoint administrators to manage the backup and restore of SharePoint data by leveraging the snapshot and replication features of the storage system. With the help of Replication Manager, administrators can schedule the application-consistent snapshot of databases using SnapView snapshots and can use Unisphere Manager to schedule the replication of the BLOB store file system. Backup and restore using snapshots requires less than 6 minutes to accomplish the task. Whereas, the asynchronous file system replication update can be scheduled for 1 minute.

In this test scenario, to study the impact of snapshots and replication update, a maximum user load that is supported by the SharePoint environment is run along with SnapView snapshots of SQL databases scheduled after every 1 hour and asynchronous file system replication update occurring at an interval of 1 minute. The

RPS and the average test time during the load were measured.

RPS and average user response time

The following graph shows the average RPS and the average test time before, during, and after the SnapView snapshot session was created for the SharePoint farm using the 80/10/10 mixed workload. The snapshot creation process of the SQL database completed in less than 6 minutes.




Conclusion

The SharePoint farm sustained an RPS of 5.67 during the SQL database snapshot creation. The time taken to complete the snapshot job for SQL databases, every time it was scheduled to run, was less than 6 minutes. The snapshot and file system replication update did not affect the latency for browse and search operations. However, the modify operations took 0.68 seconds longer to complete when compared to the normal operation.

CPU utilization of SharePoint servers

The following graph shows the average CPU utilization of SharePoint before, during, and after the snapshot session created for the SharePoint farm databases using the 80/10/10 mixed workload. The snapshot creation process of the SQL database completed in less than 6 minutes.




Conclusion

The CPU utilization of the WFE server and the application server remained constant, whereas there was a slight increase in the CPU usage of SQL Server (from 20.5 percent to 29 percent) and the VNX Data Mover (from 4 percent to 21 percent) during the snapshot creation.

Disk latency and IOPS

The following graph shows the SharePoint database disk latency and IOPS during the three different user profile loads.



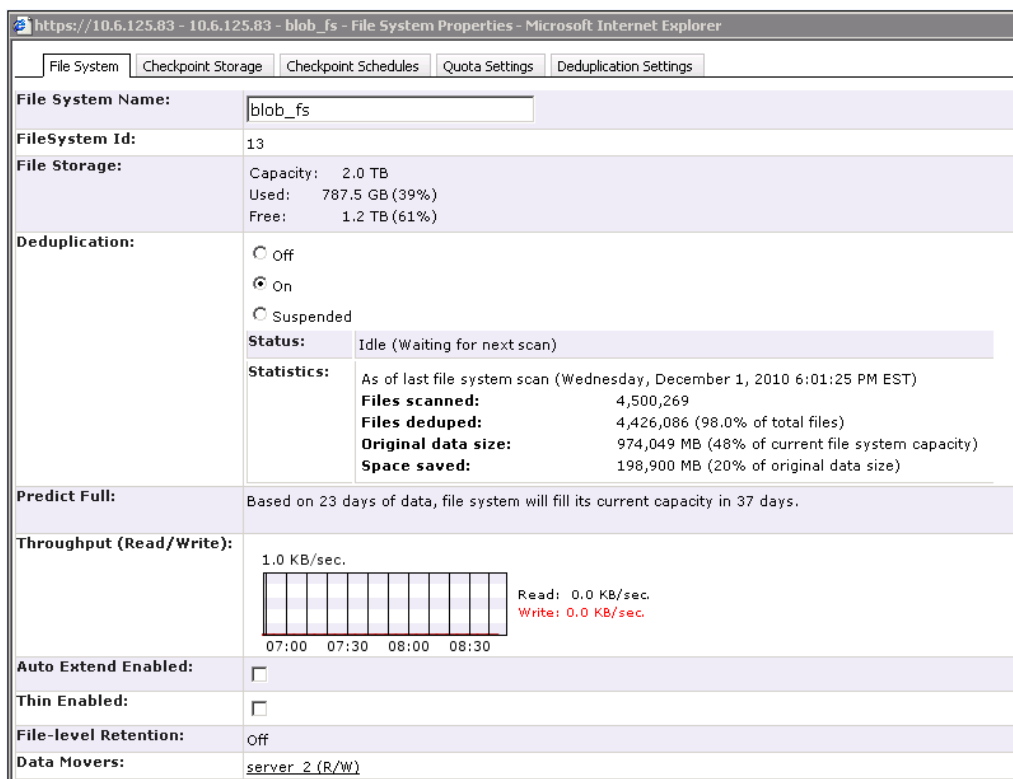
Conclusion

The graph shows that the SnapView snapshot operation had a slight impact (less than 1 ms increase) on the database and log disk-read and disk-write latencies. The overall average latencies for disk-read and disk-write remained well below 20 ms for all the user profile tests.

Using EMC file deduplication and compression with externalized BLOBs

Deduplication benefits

Deduplication in the test environment yielded 20 percent space savings. The following figure shows the NL-SAS shelf BLOB storage File System Properties page on Unisphere Manager after the entire file system is scanned and the deduplication process is completed.



Test summary

In this configuration, the SharePoint databases were stored on an FC LUN that was created by using one RAID 5 (4+1) configuration of SAS disks, and the BLOBs were externalized to a VNX CIFS share that was created by using one RAID 5 (4+1) configuration of NL-SAS disks. The BLOBs in the SharePoint databases were externalized by using the StoragePoint provider. Then deduplication was enabled for the BLOB store (CIFS share) file system.

The following table summarizes the user profile test results and shows the RPS values that were obtained for the three user profiles along with the average user response times for browse, search, and modify operations. The maximum user capacity was determined by using a defined concurrency of 10 percent for all the three user profiles.

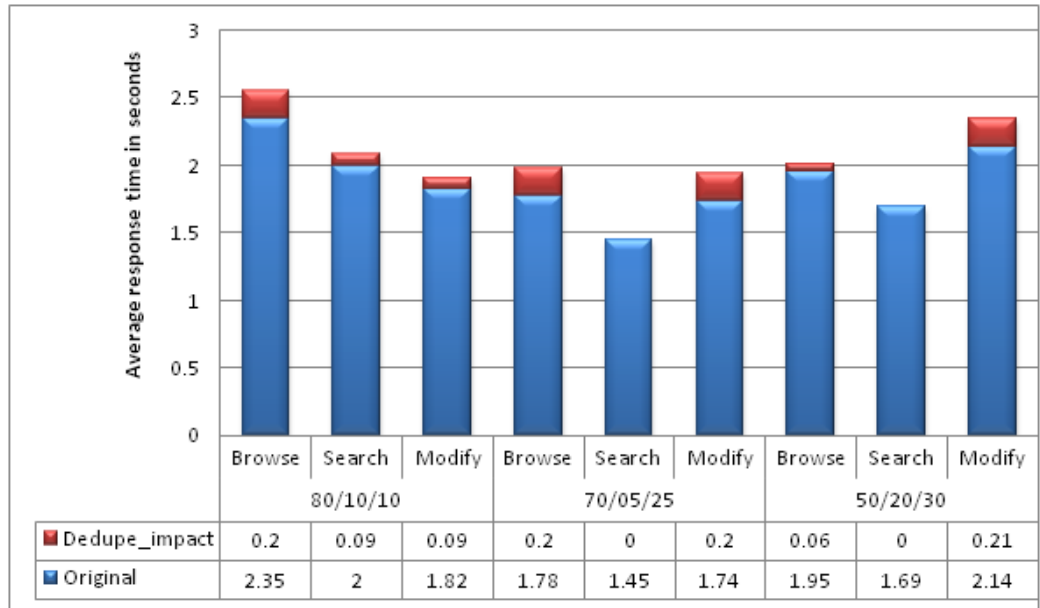
User profile (browse/search/modify)	User load profile (36 RPH)	Requests per second (RPS)	Concurrency (%)	Maximum user capacity	Average user response time (s) (browse/search/modify)
80/10/10	Typical	5.67	10	5,670	2.55/2.09/1.91

70/05/25	Typical	7.91	10	7,910	1.98/1.45/1.94
50/20/30	Typical	8.63	10	8,630	2.01/1.69/2.35

During the tests, it was observed that the SharePoint farm RPS remained unchanged for the three user profile tests. Deduplication for the BLOB store file system had a very low impact in the increase of average CPU utilization of VNX Data Mover. The increase in CPU utilization was 12 percent for all the three user profile tests.

RPS and average user response time

The following graph shows the performance results of SharePoint 2010 server for the three different user profiles.




Conclusion

Deduplication of the BLOB store file system did not affect the SharePoint farm throughput. Hence, the RPS values remained the same. However, at the maximum RPS, the response time for browse, search, and modify operations was slightly higher by the margin shown in the graph as dedupe_impact. But the overall average response time for these operations remained well under the acceptable limit set by Microsoft for SharePoint users.

Processor utilization of SharePoint servers

The following graph shows the saturation user load and the processor utilization of SharePoint servers for the different user profile loads.




Conclusion

Deduplication of the BLOB store file system had a slight impact on the processor utilization of the SharePoint servers, as shown in the graph. This impact did not affect the stability of the SharePoint environment as a whole. The SharePoint servers were able to handle the deduplication of the BLOB store file system.

Disk latency and IOPS

The following graph shows the SharePoint database disk latency and IOPS observed during the three different user profile loads.



Conclusion

Deduplication of the BLOB store file system had a slight impact on the disk-read and disk-write latency. The increase was less than 2 ms. But the overall average latencies for disk-read and disk-write remained well below the 20 ms mark for all the user profile tests.