

EMC CONTINUOUS PROTECTION FOR VIRTUALIZED SHAREPOINT 2010 FARMS

A Detailed Review

EMC GLOBAL SOLUTIONS

Abstract

This white paper describes the design, deployment, and validation of a comprehensive midsize Microsoft SharePoint 2010 solution. This solution is based on VMware vSphere™ 4.1 technologies for server virtualization and availability. EMC® CLARiiON® CX4-120 is used as the storage platform for virtual-machine data, content, and remote BLOB storage (RBS). Finally, EMC RecoverPoint, which is integrated with VMware vCenter™ Site Recovery Manager (SRM), ensures continuous availability of the Microsoft SharePoint 2010 farm.

June 2011

vmware[®]
PARTNER

TECHNOLOGY
ALLIANCE

Copyright © 2011 EMC Corporation. All Rights Reserved.

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.

VMware, ESX, vMotion, VMware vCenter, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h8139.2

Table of Contents

Executive summary	5
Overview	5
Business case	5
Key results.....	5
Introduction	6
Purpose.....	6
Scope.....	6
Audience.....	6
Terminology	7
Solution architecture	9
Solution architecture overview	9
Physical architecture	9
Validated environment profile	10
Hardware resources.....	10
Software resource	11
Mid-farm SharePoint 2010 design in a virtualized environment	13
Mid-farm SharePoint 2010 design overview	13
SharePoint farm configuration	13
Remote BLOB storage for SharePoint 2010 configuration	13
SharePoint 2010 search architecture.....	14
SharePoint index protection design.....	14
Design consideration for tempdb	15
Virtual machine configuration and resource allocation	15
VMware High Availability and Distributed Resource Scheduler design	17
Storage design	18
Storage design overview.....	18
Disk layout	18
Storage design and consideration for SharePoint	20
FAST Cache configuration	22
Data compression configuration.....	23
Storage design and consideration for Replication Manager	24
Storage design consideration for RecoverPoint.....	24
Network design	26
Network design overview.....	26
Network design consideration for VMware infrastructure	26
Network design consideration for RecoverPoint.....	26

Backup and recovery design for midsize SharePoint farm	27
Backup and recovery design overview	27
Backup and recovery consideration for content databases with remote BLOB storage	27
Static mount for Replication Manager	28
Integration with Metalogix Selective Restore Manager	29
Disaster recovery design for midsize SharePoint farm	30
Disaster recovery design overview	30
RecoverPoint continuous remote replication configurations	30
RecoverPoint consistency group configuration	30
Integrating RecoverPoint with VMware vCenter	32
RecoverPoint management within Unisphere	33
VMware vCenter Site Recovery Manager design consideration	34
Failback consideration	35
Best practices for VMware Site Recovery Manager	36
Test methodology and scenarios	37
Test methodology	37
Test tools	37
Load generation	37
SharePoint user profiles	37
Test scenarios	38
Test results	39
Test 1: Baseline SharePoint farm performance test	39
Test 2: SharePoint farm performance test with FAST Cache	40
Test 3: SharePoint farm performance test with data compression	41
Test 4: Test for local VMware HA and Distributed Resource Scheduler clusters	42
Test 5: Local data protection test using EMC Replication Manager and SnapView snap technology	43
Test 6: Local content database recovery using Replication Manager	44
Test 7: Single item recovery with Metalogix Selective Restore Manager	46
Test 8: Remote site failover and failback using SRM 4.1 and RecoverPoint CRR	48
Conclusion	51
Summary	51
Findings	51
Next steps	51
References	52
White papers	52
Product documentations	52

Executive summary

Overview

This white paper documents validated test results for a virtualized SharePoint 2010 environment using EMC and VMware infrastructure technology. The SharePoint production environment runs on a CLARiiON® CX4-120 storage system and is protected locally with EMC® Replication Manager snapshots and VMware® High Availability. Site resiliency is provided by EMC RecoverPoint and VMware vCenter™ Site Recovery Manager (SRM).

These features provide a complete solution for SharePoint protection on the market.

Business case

Microsoft SharePoint Server is quickly becoming more than just a collaboration platform, but a core platform within the mission-critical framework for many organizations. Availability and recoverability become critical factors when designing the supporting IT infrastructure.

EMC is a leader in providing scalable infrastructure and data protection solutions for SharePoint. VMware vSphere™ integration is commonly used across the EMC storage and replication product portfolio, with first-to-market technologies. EMC Proven Solutions leverages these advantages to create robust solutions that help maintain your service-level agreement.

Key results

In this solution, an efficient and fully virtualized SharePoint environment with high availability is demonstrated:

- The SharePoint farm in the solution (including the SQL server) was virtualized depicting a midsize SharePoint environment with high availability.
- The sustained maximum user capacity was 13,080 at 10 percent user concurrency on the virtualized SharePoint farm.
- The average search and browse response time is improved after the EMC FAST Cache technology is enabled.
- Thirty percent of physical disk space was saved by using EMC LUN compression features on the SharePoint server BLOB store LUNs. With SQL data file storage space, 91.2 percent of space was freed after enabling SQL RBS FileStream.
- A full-site disaster caused only 15 minutes 33 seconds of downtime while all SharePoint virtual machines were failed over from the production site to the disaster recovery site because of the use of EMC RecoverPoint/SE and VMware SRM.
- It took 27 minutes to replicate a 1.5 TB SharePoint 2010 farm by using EMC Replication Manager 5.3.1 with the EMC SnapView™ snap technology for local protection with no downtime to the application and negligible host resource impact.
- With EMC Replication Manager 5.3.1, it took only 6 minutes to restore a 100 GB content database from a replica.

Introduction

Purpose

The purpose of this solution is to:

- Present a realistic implementation that includes cost-effective configuration achieved by optimized storage configuration, server virtualization, storage-based replication for rapid backup and restore, and multi-site availability with low recovery point objective (RPO) and recovery time objective (RTO).
- Optimize the data-center efficiency by minimizing the physical footprint, connectivity requirements, and energy consumption through virtualization and consolidation.
- Free valuable database resources by storing BLOBs outside the SQL database and leveraging EMC data compression technology.
- Show rapid replication and recovery using EMC Replication Manager and EMC SnapView snap technology for SharePoint 2010.
- Protect the Microsoft Office SharePoint 2010 farm from a site disaster by leveraging RecoverPoint/SE and VMware vCenter SRM.
- More granular recovery is enabled by Metalogix Selective Recovery Manager, a product of an EMC partner, Metalogix.

Scope

This white paper covers:

- Presenting an overview of the concepts and technologies in the solution.
- Documenting the backup and restore performance by using EMC Replication Manager and EMC SnapView snap technology for a midsize SharePoint 2010 farm.
- Providing the minimum disruption of service and the minimum data loss that is limited by the asynchronous nature of the replication (minimum RPOs and RTOs) for a site disaster.

This white paper does not provide detailed installation instructions. Actual implementations can vary based on customer-specific environmental factors and the parameters that the test results show.

Audience

This white paper is intended for Microsoft SharePoint 2010 database administrators and storage architects who are involved in planning, architecting, or administering an environment with EMC CLARiiON as the storage platform, and also for those who are planning to implement the replication solutions.

Terminology

Table 1 defines terms used in this document.

Table 1. Terminology

Term	Definition
Asynchronous replication	Asynchronous replication ensures that data replicated at a secondary site is an identical copy of the primary site with some data lag. The writes on the primary site can continue without waiting for the acknowledgment from the secondary site.
Consistency group	A consistency group is a set of replicated volumes that are managed as a single entity and whose secondary images always remain in a consistent and restartable state with respect to their primary image and each other.
EMC data compression	EMC data compression is an EMC CLARiiON efficiency feature that allows users to store information using as little as half the storage capacity that they use today. All compression and decompression processes are handled by CLARiiON, so no server cycles are consumed in the process, and no additional server software is required.
EMC Replication Manager	EMC Replication Manager automates and simplifies the management of replicas. It orchestrates critical business applications, middleware, and underlying EMC replication technologies to create and manage replicas at the application level for a variety of purposes, including operational recovery, backup, restore, development, and simulation. Customers interested in reducing manual scripting efforts, improving recovery, and creating parallel access to information can implement Replication Manager to put the right data in the right place at the right time.
Journal	The RecoverPoint journal is used to track the changes to the source volume. A journal consists of one or more journal volumes. A journal is required for both the source and the target/destination side for a RecoverPoint CRR setup in order to support the failover from the active production side to the DR side then failback while necessary. You must assign the journal volume while creating the consistency group replication set.
RecoverPoint continuous remote replication (CRR)	RecoverPoint continuous remote replication (CRR) provides bidirectional, heterogeneous block-level replication across any distance using asynchronous, synchronous, and snapshot technologies over FC and IP networks. It periodically distributes updates to the remote copy of the data with all the changes that occurred on the local copy since the last update.
Recovery point objective (RPO)	Recovery point objective (RPO) is the point in time to which systems and data must be recovered after an outage. This defines the amount of data loss a business can endure.
Recovery time objective (RTO)	Recovery time objective (RTO) is the period of time within which the systems, applications, or functions must be recovered after an outage. This defines the amount of downtime that a business can endure and survive.

VMware vCenter Site Recovery Manager (SRM)	VMware vCenter Site Recovery Manager (SRM) delivers advanced capabilities for disaster recovery management, non-disruptive testing, and automated failover. VMware vCenter SRM can manage failover from production datacenters to disaster recovery sites, as well as failover between two sites with active workloads. Multiple sites can even recover into a single shared recovery site. Site Recovery Manager can also help with planned datacenter failovers such as datacenter migrations.
VMware Distributed Resource Scheduler (DRS)	VMware Distributed Resource Scheduler (DRS) continuously monitors CPU utilization across cluster nodes and resource pools, and intelligently allocates available resources among virtual machines according to business needs.
VMware High Availability (VMHA)	VMware High Availability (VMHA) provides cost-effective high availability for any application running in a virtual machine, regardless of its operating system or underlying hardware configuration.

Solution architecture

Solution architecture overview

This solution demonstrates a fully virtualized SharePoint 2010 environment by using VMware vSphere and EMC CLARiiON CX4-120 with EMC efficient storage technology. The SharePoint 2010 farm with high availability is protected locally with Replication Manager integrating with EMC SnapView snap technology and remotely with RecoverPoint/SE CRR controlled by VMware vCenter SRM.

Physical architecture

Figure 1 depicts the overall physical architecture of the validated solution environment.

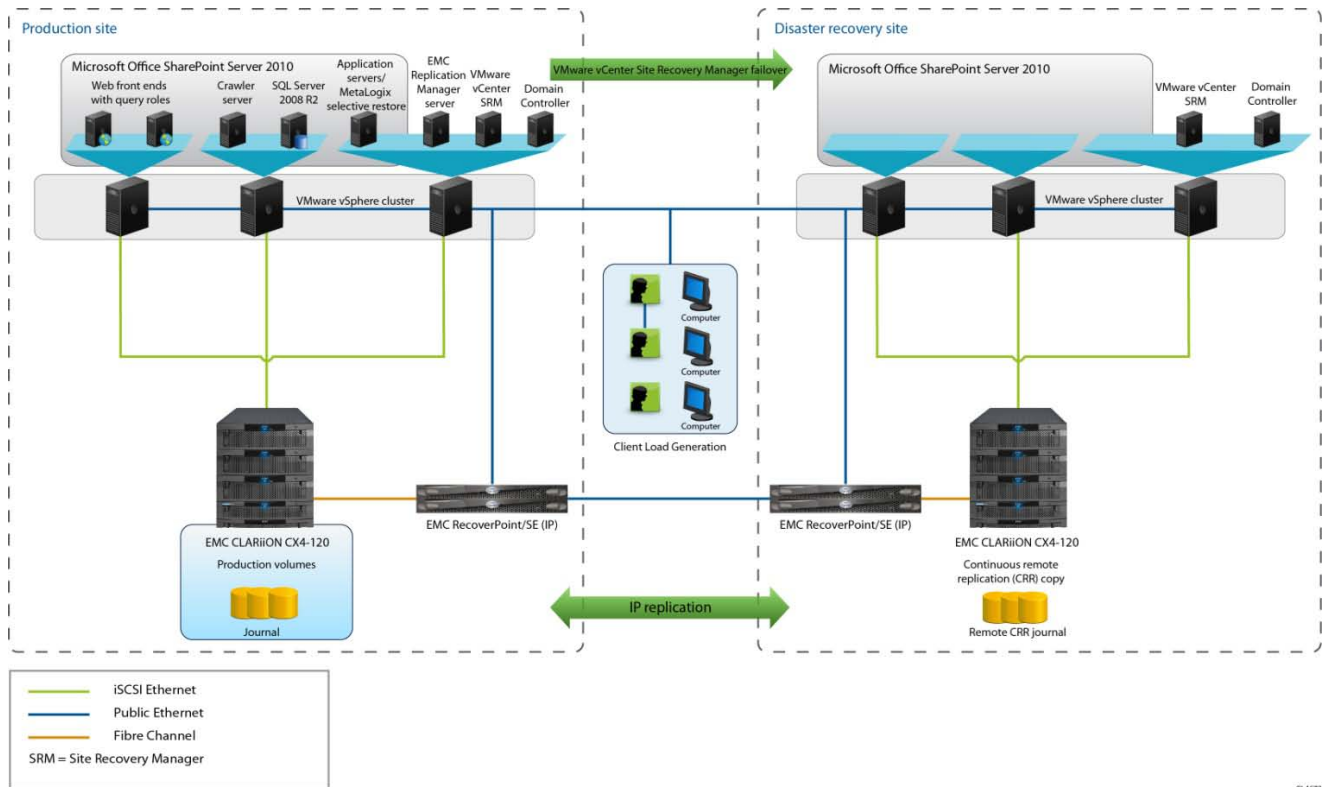


Figure 1. Physical architecture of the solution environment

Validated environment profile

The solution was validated with the environment profile as shown in Table 2.

Table 2. Profile characteristics

Profile Characteristics	Value
SharePoint farm user data	1 TB
User concurrency	10%
User profile	Heavy (60 requests in each hour)
Sites (for each site collection)	10
SQL Server 2008 R2 virtual machine	1
Web front-end (WFE) virtual machine	2
Application server virtual machine (hosting central administrator)	1
Crawler virtual machine	1
VMware vSphere cluster (physical)	6 nodes (3 for productions / 3 for DRs)

Hardware resources

Table 3 lists the details of the hardware required resources.

Table 3. Hardware required resources

Equipment	Quantity	Configuration
Storage array (production site)	1	EMC CLARiiON CX4-120 with: <ul style="list-style-type: none"> • 2 x Fibre Channel (FC) ports and 2 x iSCSI ports for each storage processor (SP) • FLARE® 30 – 4.30.000.5.507 • 9 x 450 GB 15,000 rpm FC disk • 5 x 600 GB 15,000 rpm FC disk • 30 x 1 TB 7,200 rpm SATAII disk • 2 x 73 GB Flash drive
Storage array (disaster recovery site)	1	EMC CLARiiON CX4-120 with: <ul style="list-style-type: none"> • 2 x FC ports and 2 x iSCSI ports for each SP • FLARE 30 – 4.30.000.5.507 • 9 x 450 GB 15,000 rpm FC disk • 5 x 600 GB 15,000 rpm FC disk • 30 x 1 TB 7,200 rpm SATA II disk

VMware ESX® servers 2 socket quad-core (8 cores) 2.27 GHz, 36 GB RAM (production site)	3	Virtual machines include: <ul style="list-style-type: none"> • 1 x SP crawler server (4 vCPUs/8 GB) • 2 x SP WFE servers (4 vCPUs/8 GB) • 1 x SP application server (2 vCPUs/4 GB) • 1 x SQL server (4 vCPUs/16 GB) • 1 x Replication Manager server/Mount host (2 vCPUs/4 GB) • 1 x vCenter SRM server (2 vCPUs/4 GB) • 1 x domain controller (2 vCPUs/4 GB)
VMware ESX servers 2 socket quad-core (8 cores) 2.27 GHz, 36 GB RAM (disaster recovery site)	3	Virtual machines include: <ul style="list-style-type: none"> • 1 x vCenter SRM server (2 vCPUs/4 GB) • 1 x domain controller (2 vCPUs/4 GB)
Network switches	2	1 Gb IP switch (48 ports)
Fibre Channel switches	2	4 Gb FC SAN switch (16 ports)
EMC RecoverPoint/SE appliances	4	Gen-3 RecoverPoint/SE 3.3 SP1 patch 1

Software resource

Table 4 lists the details of the software required resources.

Table 4. Software required resources

Equipment	Version
EMC Replication Manager	5.3.1
EMC CLARiiON Navisphere® Command Line Interface (CLI)	FLARE 30
EMC Solutions Enabler	7.2
EMC Admsnap	FLARE 30
EMC PowerPath®/VE	5.4 SP2
EMC RecoverPoint Site Recovery Manager failback plugin	1.0.0.2
EMC RecoverPoint/SE	3.3 SP1
EMC RecoverPoint Adapter for VMware vCenter Site Recovery Manager	1.0 SP3
Microsoft Windows Server 2008 R2	7600

Microsoft SQL Server 2008 R2	10.50.1600.1
Microsoft SharePoint 2010	14.0.4762.1000
Microsoft SQL Server 2008 R2 feature pack (Remote BLOB store)	10.50.1600.1
Visual Studio test suite 2010	SP1
VMware vCenter	4.1
VMware vCenter Site Recovery Manager	4.1
VMware vSphere	4.1
MetaLogix Selective Restore Manager	4.1.9.4
KnowledgeLake document loader	1.1

Note It is strongly recommended to use the latest EMC FLARE Operating Environment for the EMC CX4 storage.

Mid-farm SharePoint 2010 design in a virtualized environment

Mid-farm SharePoint 2010 design overview

In this midsize SharePoint 2010 environment design, the major configuration highlights include:

- The SharePoint farm uses three ESX servers at the production site. Two web front-end (WFE) servers are also configured with query roles for load balancing.
- The entire SharePoint farm is fully virtualized and deployed on the CLARiiON CX4-120 platform with IP storage (iSCSI technology).
- Remote BLOB storage (RBS) is implemented to move binary large objects (BLOB) data from SharePoint 2010 content databases to external SATA disks with EMC data compression technology.
- The query components have been scaled out to two partitions. Each query server contains a part of the index partitions and a mirror of another index partition for better query performance and fault tolerance.

SharePoint farm configuration

The SharePoint farm is designed as a publishing portal. There is 1 TB of user content, consisting of 10 SharePoint site collections (document centers) with 10 content databases, each populated with 100 GB random user data. Microsoft network load balancing (NLB) was enabled on two WFEs for load balancing and failover consideration.

In this solution, VMware High Availability is used to provide high availability for applications running on virtual machines, regardless of the operating system or the underlying hardware configurations. SQL Server clustering can also be used to configure availability within a SharePoint 2010 farm. To use SQL Server clustering on ESX server clusters, a failover cluster must be combined with two nodes and two or more shared disks.

Remote BLOB storage for SharePoint 2010 configuration

This implementation of the FileStream provider is known as the *local FileStream provider*. You can conserve storage resources by using the local RBS FileStream provider to place the extracted BLOB data on a lower-cost storage tier (SATA or NL-SAS). You cannot use RBS with the local FileStream provider on remote storage devices, such as network-attached storage (NAS). The FileStream provider is supported when it is used on locally attached hard disk drives only (FC/iSCSI/FCoE).

In this solution, RBS with a SQL server FileStream provider is implemented on eight content databases with 800 GB data. BLOB stores are stored on SATA disks with the EMC data compression feature enabled.

For more information about the RBS installation and configuration on SharePoint 2010, refer to [Plan for remote BLOB Storage \(RBS\)](http://technet.microsoft.com/en-us/library/ff628569.aspx) (<http://technet.microsoft.com/en-us/library/ff628569.aspx>).

SharePoint 2010 search architecture

SharePoint 2010 search architecture improves scalability for both crawl and query components compared with MOSS 2007.

The search server consists of crawler servers with the function to crawl and propagate the indexes on the query server and update the property stores on the SQL server. In SharePoint 2010, the crawler server no longer stores a copy of index files. They are propagated to the query component during the crawl operation. Because of this, the crawler server is no longer a single point of failure.

The query servers split the content between themselves so that each of the query servers holds only a subset of the content index files. The property store is the authoritative source for all indexed content properties and does not need to be synchronized with the crawler servers.

For more information on scaling and sizing search architecture for the SharePoint Server 2010 farm, refer to [Plan the topology for enterprise search](#) at [Microsoft Technet](#).

SharePoint index protection design

SharePoint index protection is always a big challenge to IT administrators. Administrators can always do a full crawl to regenerate index files. However, a full crawl of 1 TB data can take more than one or two days to complete. Moreover, it requires more resources than an incremental crawl does. In this solution, SharePoint index protection is improved to reduce the necessity of full crawls and to help maintain business service-level agreement.

In this solution, the query function is running on the WFE servers. The query components have been scaled out to two partitions for load balancing. Each query component also holds a mirror of another index partition for fault-tolerance consideration. Figure 2 shows the topology of the query partitions in this solution.

Index Partition - 0 - SQL\Search_PropertyStoreDB_79b9309493e74d15bec7297da44c9b6f		
Query Component 1	WFE2	Online
Query Component 5	WFE1	Online - Failover Only
Index Partition - 1 - SQL\Search_PropertyStoreDB_79b9309493e74d15bec7297da44c9b6f		
Query Component 3	WFE1	Online
Query Component 7	WFE2	Online - Failover Only

Figure 2. Topology of query components in the search service application

A 40 GB iSCSI LUN stores the index partition and is provisioned to each query server. Another 40 GB iSCSI LUN is attached to the crawler server to store the temporary index files during the crawl operation.

Data and log files of both SharePoint 2010 property and crawl databases resided on dedicated RAID 1/0 LUNs with FAST Cache enabled because of a potentially high I/O requirement.

For further information on scaling and sizing search architecture for the SharePoint Server 2010 farm, refer to [Plan the topology for enterprise search](#) at [Microsoft Technet](#).

Besides the local protection provided by index partitions, the crawler server along with the entire SharePoint farm is protected by VMware vCenter Site Recovery Manager with RecoverPoint/SE when it fails over to the remote site. In most cases, with all these protections, customers do not need to redo a full crawl for terabytes of SharePoint content data even after a failover to the recovery sites.

Design consideration for tempdb

Microsoft SQL Server performance best practices suggest that the number of tempdb data files should equal the number of core CPUs, and each of the tempdb data files should be of the same size.

In this solution, four tempdb data files were created. The number is equal to that of SQL server core CPUs. The tempdb data and log files are placed on a dedicated iSCSI RAID 10 LUN with FAST Cache enabled for better performance.

Refer to [Optimizing tempdb performance](#) at Microsoft Technet for more information about the tempdb best practices for SQL Server 2008 R2.

Virtual machine configuration and resource allocation

The virtualization allocation of this solution is detailed in Table 5.

Table 5. Virtualization resource allocation

Server role	Quantity of servers	vCPUs	Memory (GB)	Boot disk (GB)	Raw device mapping (RDM) device (GB)	Quantity of RDM devices	RDM device
WFE	2	4	8	60	40	1	Query volume
Index	1	4	8	60	40	1	Index volume
Application (central administration)	1	2	2	60	NA	NA	NA
SQL Server 2008 R2 Enterprise	1	4	16	60	150	10	Content database data volume
					30	10	Content database log volume
					150	8	BLOB store volume
					50	1	Configuration database volume
					100 + 50	2 (1× 100 GB disk and 1× 50 GB disk)	SharePoint property data and log volumes
					100 + 50	2 (1× 100 GB disk and 1× 50 GB disk)	SharePoint crawl data and log volumes
					100	5	SQL temp database and log volumes

VMware High Availability and Distributed Resource Scheduler design

VMware High Availability (VMHA) and Distributed Resource Scheduler (DRS) features are enabled on the ESX clusters.

VMware High Availability provides high availability for applications running on virtual machines, regardless of their operating system or underlying hardware configurations. On a server failure, affected virtual machines are automatically restarted on other production servers with spare capacity.

VMware Distributed Resource Scheduler continuously monitors utilization across cluster nodes and resource pools, and intelligently allocates available resources among virtual machines according to business needs. DRS ensures maximum performance and load balancing across an ESX cluster in a reported or automated way.

In this solution, VMware ESX servers are configured to be part of VMware HA and DRS clusters with a cluster tolerance level of one host failure. The virtual machines are distributed across the three ESX servers, and the **Fully automated** setting is selected in the DRS setting as shown in Figure 3. As a result, VMware vCenter places virtual machines that join the cluster on the appropriate hosts. Running virtual machines are failed over between hosts by vCenter, as needed, to ensure the best use of cluster resources. vCenter displays a history of migrations in the **Migration** tab.

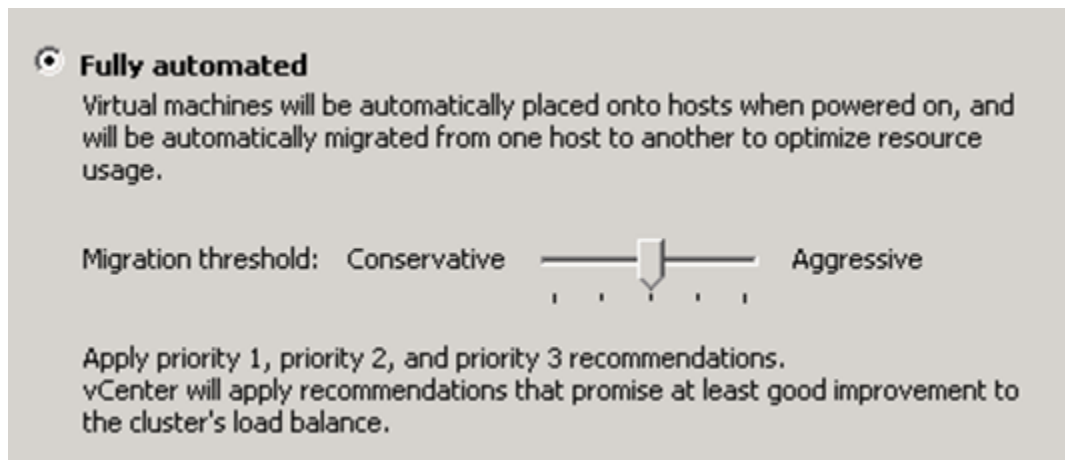


Figure 3. Production cluster settings

Storage design

Storage design overview

The storage layout on EMC CLARiiON CX4-120 is designed to ensure the optimal performance of the SharePoint farm, and rapid replication and recovery operations of the RecoverPoint components.

The following aspects are considered when designing the storage:

- Disks and pools layout planning for the production SharePoint data, BLOB store, and search database.
- Local snapshots and replication of production SharePoint data.
- RecoverPoint storage for both the production site and the disaster recovery site.

Disk layout

The following diagrams show the disk layout including SharePoint data, snap cache, and RecoverPoint journal LUNs on the production and DR sites.

Figure 4 shows the disk layout on the production site.

Production site storage layout

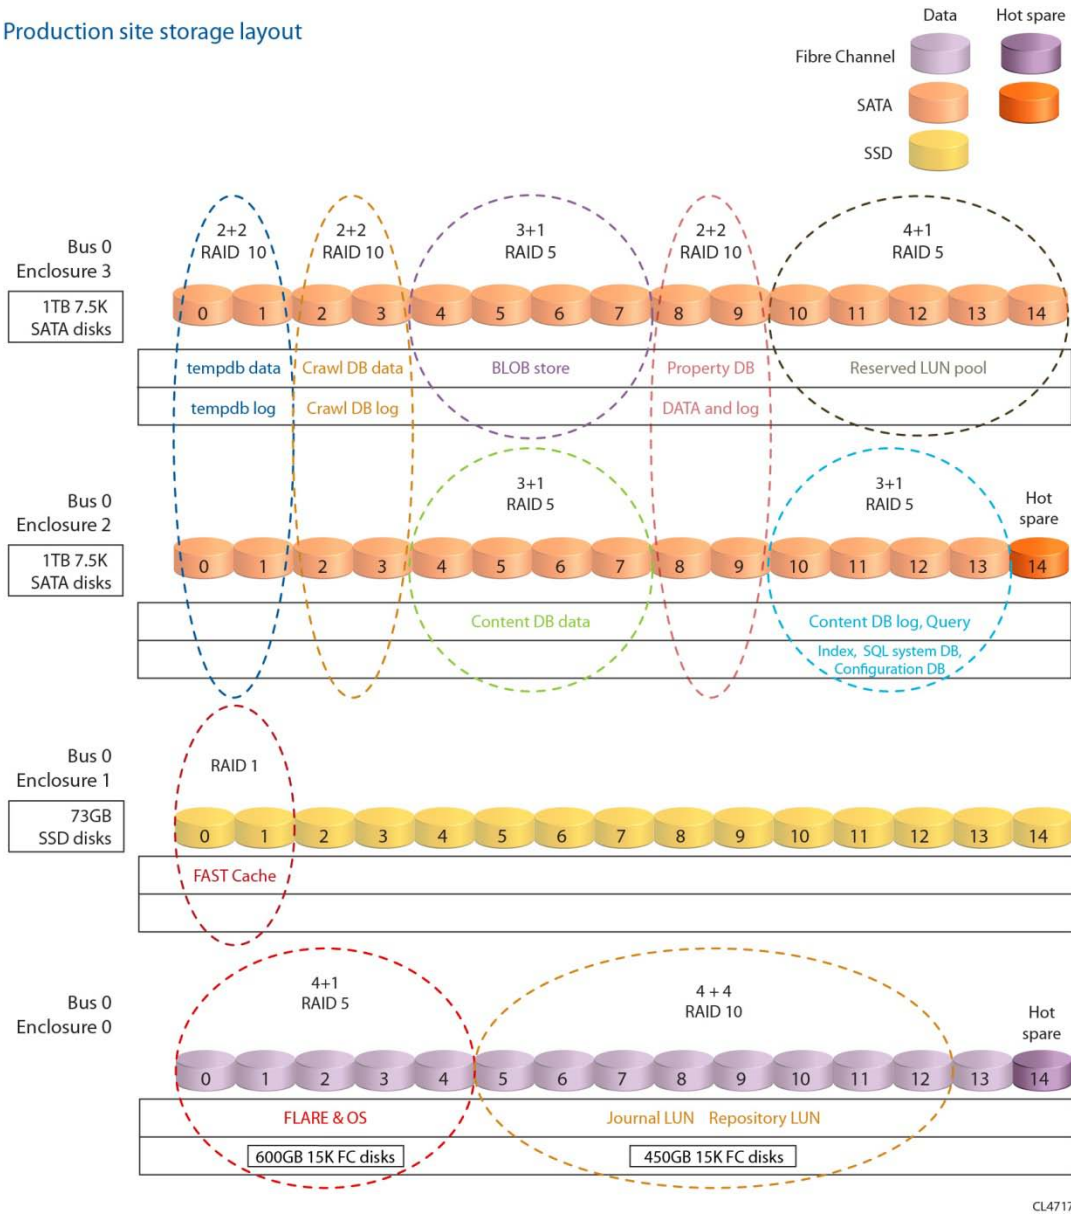


Figure 4. Production site storage layout

Figure 5 shows the storage layout of the DR site.

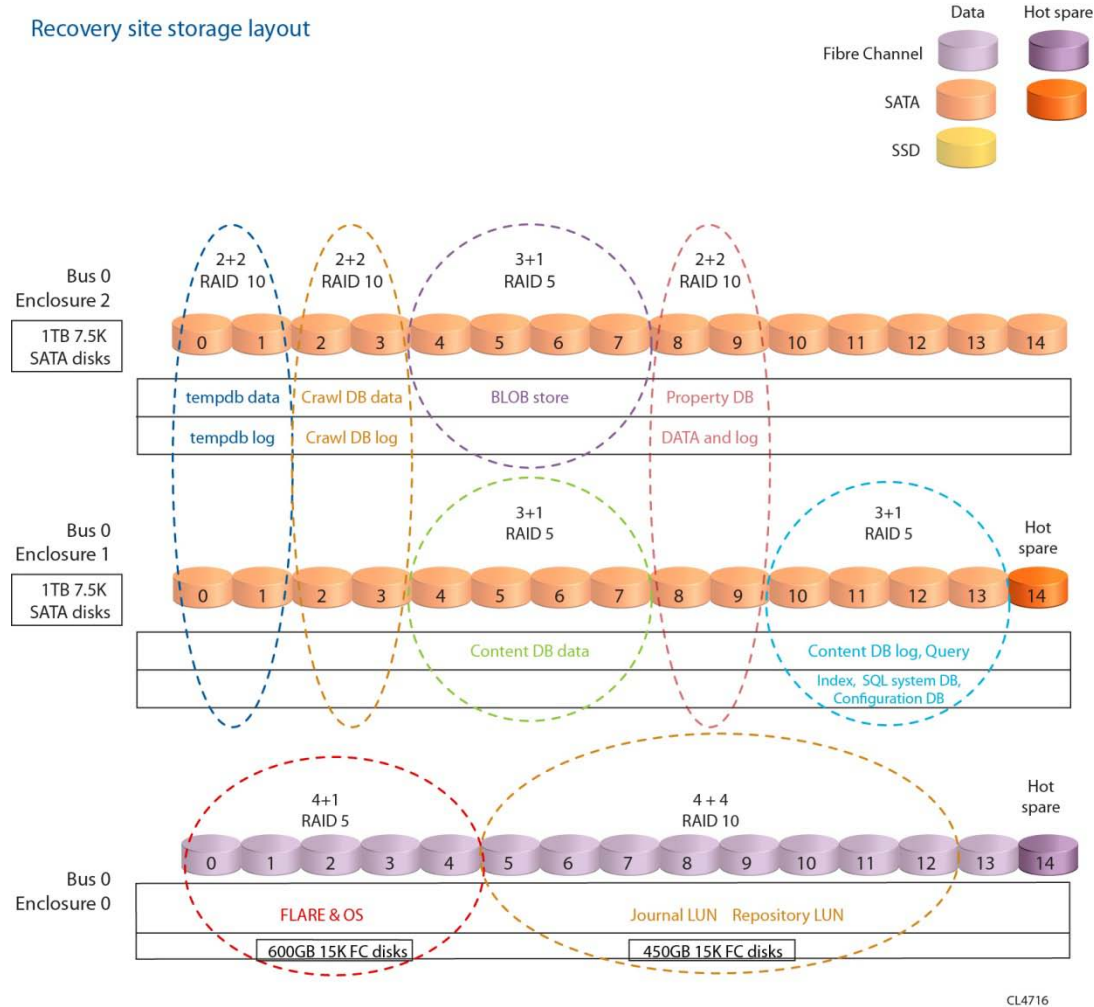


Figure 5. Recovery site storage layout

Storage design and consideration for SharePoint

The key design considerations for SharePoint 2010 in this solution are the following:

- SATA spindles are used to store all the SharePoint data including content databases, BLOBs, search databases, and tempdb on CLARiiON CX4-120.
- FAST Cache is enabled on the tempdb, search databases, and search database log RAID 1/0 LUNs as these are highly write-intensive.
- BLOB store is stored on pool-based thin LUNs with data compression enabled to offer reductions in used storage capacity.
- 800 GB of SharePoint data is externalized to the BLOB store by using RBS with the FileStream provider on SQL Server 2008 R2.

Table 6 lists the volume layout for the SharePoint data on the production and DR CLARiiON CX4-120 array.

Table 6. SharePoint volume layout

Volume	Disk Configuration	Description
Operating System volume	600 GB 15,000 rpm FC disks R5 (4+1)	Domain controller, WFE, SQL, application, crawler, vCenter, and Replication Manager server OS LUN
FAST Cache	73 GB Flash drives R1 (1+1)	To increase search and tempdb read/write performance
Tempdb	1 TB 7,200 rpm SATA disks R10 (2+2)	tempdb data and log LUNs
Crawl databases and log		Crawl database data and log LUNs
Property databases and log		Property database data and log
Content databases data volume	1 TB 7,200 rpm SATA disks R5 (3+1)	Content databases data LUNs
Content databases log volume	1 TB 7,200 rpm SATA disks R5 (3+1)	Content databases log LUNs
SQL system databases		SQL system databases LUN
Query volume		Query LUNs to store indexes on each WFE server
Index volume		Index LUN to store the temporary index on the crawler server
SharePoint configuration databases		SharePoint configuration databases LUN
BLOB stores	1 TB 7,200 rpm SATA disks R5 (3+1)	Eight BLOB store LUNs
Hot spares	1 TB 7,200 rpm SATA disks 600 GB 15,000 rpm FC disks	Hot spares for SATA and FC disks

FAST Cache configuration

FAST Cache is a new feature introduced in FLARE 30. It optimizes the performance by serving high-demand workloads from an extended cache of up to 2 TB.

In this solution, FAST Cache is enabled on the search crawl and property database data LUNs and log LUNs, because they can be highly write-intensive during the crawl and query operations. Two 73 GB Flash disks in RAID 1 are used as FAST Cache on the CLARiiON CX4-120 storage array as shown in Figure 6.

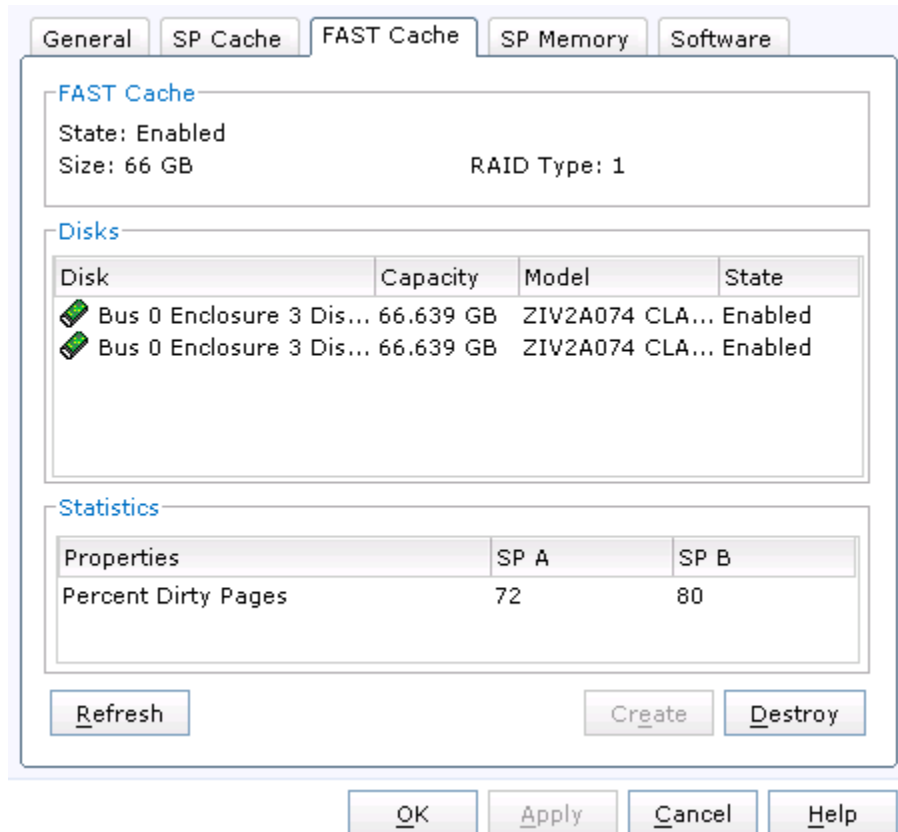


Figure 6. FAST Cache configuration on the CLARiiON CX4-120 system

It is easy to display the FAST Cache properties in any Unisphere™ table. In Figure 7, the FAST Cache property is displayed as **ON** in the tempdb LUNs.

Name	FAST Cache	ID	State	User Capacity (GB)
TEMP_DB_LUN_1	On		25 Ready	100.000
TEMP_DB_LUN_2	On		26 Ready	100.000
TEMP_DB_LUN_3	On		27 Ready	100.000
TEMP_DB_LUN_4	On		28 Ready	100.000
TEMP_LOG_LUN	On		29 Ready	100.000

Figure 7. FAST Cache property in the Unisphere table

Refer to the white paper *EMC CLARiiON, Celerra Unified, and VNX FAST Cache – A Detailed Review* for the configuration and best practices of FAST Cache.

Data compression configuration

In this solution, the EMC data compression feature is enabled on the BLOB store LUNs. It allows users to store information while using less storage capacity. The EMC data compression feature is available for CLARiiON platforms running release 30 of the FLARE operation environment. Compression is compatible with all CLARiiON replication features like SnapView.

All the BLOB stores are stored in the pool-based LUNs (thick LUNs). When the LUNs are compressed, they become thin LUNs and remain in the pool. Thin LUNs are the key element of data compression, because as data is compressed, thin-LUN blocks are freed and return to the pool, and then are able to be used by other LUNs.

In this solution, the default compression rate setting (medium) is applied. At the medium rate, up to 30 to 50 percent of CLARiiON SP CPU can be used with the maximum number of concurrent compression operations running. The initial compression is started without user load because the returned capacity to the pool can contribute to the CPU and write-cache usage.

For more information about the configuration and best practices of EMC data compression, refer to the white paper *EMC Data Compression – A Detailed Review*.

As shown in Figure 8, both **Thin** and **Compression** are set to **On**. That means all BLOB-store LUNs have been set to thin LUNs with the data-compression feature enabled.

Name	ID	State	Thin	Compression
BLOB_Store__1	52	Ready	On	On
BLOB_Store__2	53	Ready	On	On
BLOB_Store__3	54	Ready	On	On
BLOB_Store__4	55	Ready	On	On
BLOB_Store__5	56	Ready	On	On
BLOB_Store__6	57	Ready	On	On
BLOB_Store__7	58	Ready	On	On

Figure 8. Pool-based LUN property

Storage design and consideration for Replication Manager

Either the CLARiiON SnapView clone or snaps technology can be used for Replication Manager to create a replica for SharePoint data on a CLARiiON array. The CLARiiON SnapView clone technology is a mirror of the real volume. The CLARiiON SnapView snap technology uses the copy-on-first-write feature to create point-in-time snapshots.

In this solution, the CLARiiON SnapView snap technology is used for Replication Manager replication. The CLARiiON SnapView snap technology can create or destroy a snapshot in seconds, regardless of the LUN size, because it does not copy all of the data but only the changes. It significantly reduces the replication time when compared with the CLARiiON SnapView clone technology.

To configure CLARiiON SnapView snaps, a reserved LUN pool with the proper number and size of LUNs (also known as snapshot cache) should be allocated for the snapshot function. In this particular solution, snapshot cache is created on RAID 5 (3+1) of four SATA spindles, because all the SharePoint data are stored on SATA spindles for a mid-farm environment. Forty three LUNs with a total of 1.5 TB user data needs to be backed up, so 80 LUNs of 30 GB capacity are created to form the snapshot cache.

For more information about the SnapView snap design and consideration, refer to *EMC CLARiiON SnapView Snapshots and Snap Sessions Knowledgebook*.

Storage design consideration for RecoverPoint

The RecoverPoint repository and journal LUNs are on RAID 1/0 Fibre Channel disks. This is to ensure the best performance for RecoverPoint, especially for the repository LUN.

The journal capacity of each consistency group must be large enough to support the business requirements of the group. In order to calculate the total capacity required for the RecoverPoint journals, it is necessary to measure the data change rate at the production site.

EMC also offers other approaches for the journal volume capacity calculation. Customers can either submit the request to the Solution Validation Center (SVC), or use the sizing tools for the calculation such as EMC BCSD (Business Continuity Solution Designer) and the RecoverPoint Calculator. The RecoverPoint Calculator is located on EMC Powerlink:

http://powerlink.emc.com:80/km/live1/en_US/Sales_Support/Tool/RecoverPointSE_Calculator.xls.

As a general rule, when changed data statistics are not available, the journal size can take 20 percent capacity of the data being replicated. In this solution, the initial journal size is 20 percent of the capacity of the total replication sets.

For a more accurate design, the changed data statistic with user load is captured by monitoring the bandwidth on the CLARiiON array storage processor and RecoverPoint statistics.

The following equation is used for sizing the journal volume:

$$\text{Journal size} = \frac{(\text{changed data in Mbps}) \times (\text{required rollback time in seconds})}{(20\% \text{ reserved for image - access log}) \times \text{system factor (usually 5\% for internal system needs)}}$$

Refer to the latest *EMC RecoverPoint Administrator's Guide* for a detailed explanation on the equation.

In this solution, the required rollback time is 24 hours (86,400 seconds). The estimated change rate is 40 Mb/s with full user load. As a result, the estimated journal size was:

$$\text{Journal size} = \frac{40 \times 86400}{(1-0.20) \times 1.05} = 4536000\text{Mb} = 567\text{GB}$$

Note Multiple small-sized journal volumes are preferred compared with one large journal volume when further storage growth is planned.

The RecoverPoint array usage with LUN sizes and total LUN space on the production and DR array is shown in Table 7.

Table 7. Total LUN space on the CLARiiON CX4-120 array

CX4-120 array LUN	LUN size (GB)	Number of LUNs	Total LUN space (GB)
RecoverPoint repository	5	1	5
SharePoint 2010 journal	50	12	600

Network design

Network design overview

To ensure the optimal iSCSI performance for the SharePoint 2010 farm in a small or midsize environment, 1 Gb Ethernet connections for all network connections are strongly recommended.

The network components used include the following:

- Two 48-port GigE network switches in this solution and one per site for both the local site usage and the connection to the remote production/DR site.
- Two 16-port 4 Gb Fibre Channel switches in this solution and one per site for the data transferring between the CLARiiON array to the RecoverPoint appliance (RPA) nodes.

Network design consideration for VMware infrastructure

In the VMware ESX hosts network, two virtual switches are configured with the four connected network interface cards (NICs) of each VMware ESX server:

Virtual switch 0: Named vSwitch0 at the production and DR sites with two NICs connected for I/O. The port group for the virtual machines was configured for the public virtual machine network.

Virtual switch 1: Named vSwitch1 at the production and DR sites with two NICs connected for I/O. This vSwitch is used for VMkernel to run the iSCSI connection and vMotion® migration communication. The CLARiiON iSCSI ports are divided into two IP subnets in order to be aligned with CLARiiON best practices. Thus, the two VMkernel IP addresses are defined to enable the iSCSI multipath and communication between the iSCSI initiator of the vSphere hosts and the iSCSI targets of the CLARiiON CX4-120 array.

Network design consideration for RecoverPoint

RecoverPoint uses powerful bandwidth reduction and compression technologies to minimize bandwidth usage, which dramatically reduces the time lag on writing data to the storage from the source site to the target site. In this case, RecoverPoint uses the CLARiiON array-based splitter technology for data splitting. The following are considered when designing the RecoverPoint data replication network:

- Use two FC switches for data transfer and replication between the CLARiiON CX4-120 array and the RecoverPoint nodes.
- Use the existing IP switches for the remote site RecoverPoint CRR data replication.

Backup and recovery design for midsize SharePoint farm

Backup and recovery design overview

In this solution, EMC Replication Manager with SnapView snap technology is used to automate the creation of mountable production-data snapshots without taking SharePoint production offline for local backup and recovery. Replication Manager provides an easy-to-use graphical user interface (GUI) to manage, schedule, and monitor replication jobs. It is also flexible to script by using Command Line Interface (CLI). Metalogix Selective Restore Manager is also used to perform single-item recovery from snapshots. The replication jobs are set to run every 4 hours to protect the whole SharePoint farm.

The following aspects are considered when designing Replication Manager:

- Replication Manager 5.3.1 is used to create full replicas of an entire SharePoint 2010 farm (including content databases, configuration database, administration database, search database, and search index files).
- The Replication Manager agent is installed on all the SharePoint servers for replication jobs.
- All the SharePoint data LUNs are accessed through a VMware physical RDM device. Local storage (like the C drive) is not supported as an RDM.
- The Replication Manager server is also the mount host of the replica. Because the mount host is a virtual machine, a static mount is used in this solution.

Replication Manager uses the Windows VSS framework to obtain metadata about the SharePoint farm for replication from the SharePoint VSS writer. RM also utilizes the SQL Server Virtual Device Interface (VDI) technology to back up and restore the content databases with EMC array-based technology almost instantly. Replication Manager can pause the crawl operation and restart it so that the replicas are application consistent.

For more information about the detail configuration and installation for SharePoint application sets, refer to *EMC Replication Manager Guide 5.3.1 – Product Guide*.

Backup and recovery consideration for content databases with remote BLOB storage

In SQL Server 2008 R2 and SQL Server 2008, Replication Manager 5.3.1 automatically discovers FileStream filegroups that are associated with the database, if any. As a result, RBS FileStream can be detected and included as a part of replica automatically by Replication Manager 5.3.1 for the purpose of recovery.

Figure 9 lists the replica objects of the content databases including the RBS FileStream filegroup.

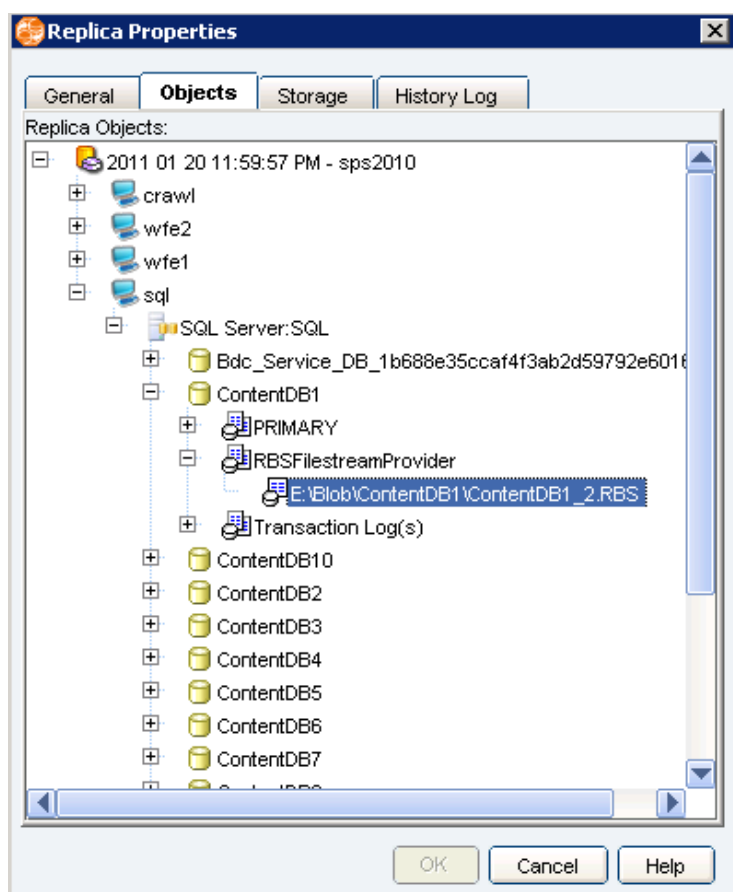


Figure 9. Replica objects of the content databases

There might be inconsistencies between the BLOB stores and the RBS databases. BLOBs can be present in the BLOB store replica while absent from the RBS database replica. It is strongly recommended to run RBS database and BLOB store consistency checks and clean up orphaned BLOBs by using RBS Maintainer after database restoration. You can refer to the Microsoft book [Remote BLOB storage for Microsoft SharePoint Server 2010](#) for more information on RBS Maintainer.

For more information on the configuration of replication jobs for SharePoint 2010 farm, refer to *EMC Replication Manager 5.3.1 – Product Guide*.

Static mount for Replication Manager

With Replication Manager, a virtual machine can act as a mount host for EMC CLARiiON iSCSI system. In this solution, static mount is used, with which, replication LUNs are manually placed into the ESX storage group, and then the LUNs are visible to the virtual machine as RDM devices before replication begins. Before running Replication Manager, manually exposing the snapshots to the mount host is needed.

During the replication, Replication Manager temporarily creates a snapshot. When Replication Manager gets to the mount phase, it detects the pre-exposed LUNs and removes the snapshot that has been created.

Consider the following when deploying a static mount in this solution:

- Do not install the Navisphere®/Unisphere agent on the virtual machine.
- Manually place the snapshot into the ESX storage group and then make the LUN visible to the virtual machine as raw device mappings before the mount job.
- Check **Maintain LUN visibility after unmount** in the mount options (see Figure 10).




Figure 10. Mount options

For detailed configuration steps for Replication Manager static mount, refer to *EMC Replication Manager Administrator's Guide*.

Integration with Metalogix Selective Restore Manager

In this solution, granular recovery is enabled by Metalogix Selective Restore Manager based on Replication Manager replicas. For more information, refer to [Metalogix Selective Restore Manager documentation](http://www.metalogix.net/helpsrsm/index.html) (<http://www.metalogix.net/helpsrsm/index.html>).

Disaster recovery design for midsize SharePoint farm

Disaster recovery design overview

In this solution, asynchronous data replication of RecoverPoint continuous remote replication (CRR) is used over IP networks between two sites. With RecoverPoint journal data buffering, the journal periodically sends updates to the remote copy of the data with all the changes based on RPO settings.

VMware Site Recovery Manager integrates with RecoverPoint through RecoverPoint storage replication adapter and makes disaster recovery rapid, reliable, manageable, and affordable. This integration centrally and automatically manages the traditional complex disaster recovery processes with the virtualized infrastructure environment.

RecoverPoint continuous remote replication configurations

In CRR configurations of this solution, the data is transferred between two sites over IP network. In addition, the CLARiiON splitter for RecoverPoint is used to replicate from the production site to the DR site.

The replication mode is set to asynchronous, and the snapshot granularity is set to dynamic. The bandwidth reduction feature of RecoverPoint is also enabled to allow for replication over slower IP links, longer distances, or both.

Refer to the latest *EMC RecoverPoint Administrator's Guide* for more information about the configuration for RecoverPoint CRR.

RecoverPoint consistency group configuration

A RecoverPoint consistency group is a set of SAN-attached storage volumes at the production and disaster recovery sites to maintain write-order consistency. Data consistency and write-order fidelity are maintained across all volumes assigned to a consistency group. The consistency group enables the SharePoint search, configuration, content, and BLOB data integration and consistency, and also avoids any data corruption issues across all farm components including SQL servers.

In this solution, the whole SharePoint farm is configured as a single consistency group in RecoverPoint. When the VMware SRM recovery plan starts, the consistency group is failed over to the DR site array so that the DR site node has access to all the iSCSI LUNs.

In this solution, 50 replication sets are created in the SharePoint 2010 consistency group to map the 45 RDM devices and five OS volumes of the SharePoint 2010 farm, as shown in Figure 11.

Name	Size	Prod_SharePoint	DR_SharePoint
✓ RSet 44	100GB	TEMP_LOG_LUN (29)	Temp_Log_LUN (29)
✓ RSet 43	100GB	TEMP_DB_LUN_4 (28)	Temp_DB_LUN_4 (28)
✓ RSet 42	100GB	TEMP_DB_LUN_3 (27)	Temp_DB_LUN_3 (27)
✓ RSet 41	100GB	TEMP_DB_LUN_2 (26)	Temp_DB_LUN_2 (26)
✓ RSet 40	100GB	TEMP_DB_LUN_1 (25)	Temp_DB_LUN_1 (25)
✓ RSet 39	40.0GB	SQL_System_LUN (67)	SQL_System_LUN (67)
✓ RSet 38	5.0GB	SQL_Mount_Host (68)	SQL_Mount_Host (68)
✓ RSet 37	50.0GB	Search_Admin_DB (62)	Search_Admin_DB (62)
✓ RSet 33	20.0GB	RP_Journal_LUN_App (164)	RP_Journal_LUN_App (69)
✓ RSet 47	40.0GB	Query_LUN_2 (23)	Query_LUN_2 (23)
✓ RSet 48	40.0GB	Query_LUN_1 (22)	Query_LUN_1 (22)
✓ RSet 6	50.0GB	Property_Log_2 (173)	Property_Log_2 (20)
✓ RSet 45	100GB	Property_DB_2 (172)	Property_Data_2 (70)
✓ RSet 49	150GB	OS_WFE2_New (168)	New_OS_WFE_2 (73)
✓ RSet 50	150GB	OS_WFE1_NEW (166)	New_OS_WFE_1 (72)
✓ RSet 34	100GB	OS_SQL_New (165)	New_OS (71)
✓ RSet 35	150GB	OS_Crawl_New (167)	New_OS_Crawl (74)
✓ RSet 32	80.0GB	OS_4_APP (4)	App_OS_LUN (5)
✓ RSet 36	40.0GB	Index_LUN (24)	Index_LUN (24)
✓ RSet 7	50.0GB	Crawl_log_2 (171)	Crawl_log_2 (21)
✓ RSet 46	100GB	Crawl_DATA_2 (170)	Crawl_Data_2 (93)
✓ RSet 31	50.0GB	Content_DB_Log_9 (50)	Content_DB_Log_9 (50)
✓ RSet 30	50.0GB	Content_DB_Log_8 (49)	Content_DB_Log_8 (49)
✓ RSet 29	50.0GB	Content_DB_Log_7 (48)	Content_DB_Log_7 (48)
✓ RSet 28	50.0GB	Content_DB_Log_6 (47)	Content_DB_Log_6 (47)
✓ RSet 27	50.0GB	Content_DB_Log_5 (46)	Content_DB_Log_5 (46)
✓ RSet 26	50.0GB	Content_DB_Log_4 (45)	Content_DB_Log_4 (45)
✓ RSet 25	50.0GB	Content_DB_Log_3 (44)	Content_DB_Log_3 (44)
✓ RSet 24	50.0GB	Content_DB_Log_2 (43)	Content_DB_Log_2 (43)
✓ RSet 23	50.0GB	Content_DB_Log_10 (51)	Content_DB_Log_10 (51)
✓ RSet 22	50.0GB	Content_DB_Log_1 (42)	Content_DB_Log_1 (42)
✓ RSet 21	150GB	Content_DB_Data_9 (40)	Content_DB_Data_9 (40)
✓ RSet 20	150GB	Content_DB_Data_8 (39)	Content_DB_Data_8 (39)
✓ RSet 19	150GB	Content_DB_Data_7 (38)	Content_DB_Data_7 (38)

50 replication sets.

Figure 11. Replication sets in the SharePoint 2010 consistency group

Managing the RecoverPoint consistency groups is an important part of this solution. RecoverPoint can manage image access of all devices in this consistency group. With VMware vCenter SRM sitting on top of this solution and orchestrating recovery from the production site to the DR site, the control of the RecoverPoint consistency groups must be passed over to the VMware vCenter SRM.

As shown in Figure 12, select **Use VMware Site Recovery Manager (SRM)** for the RecoverPoint consistency group control.

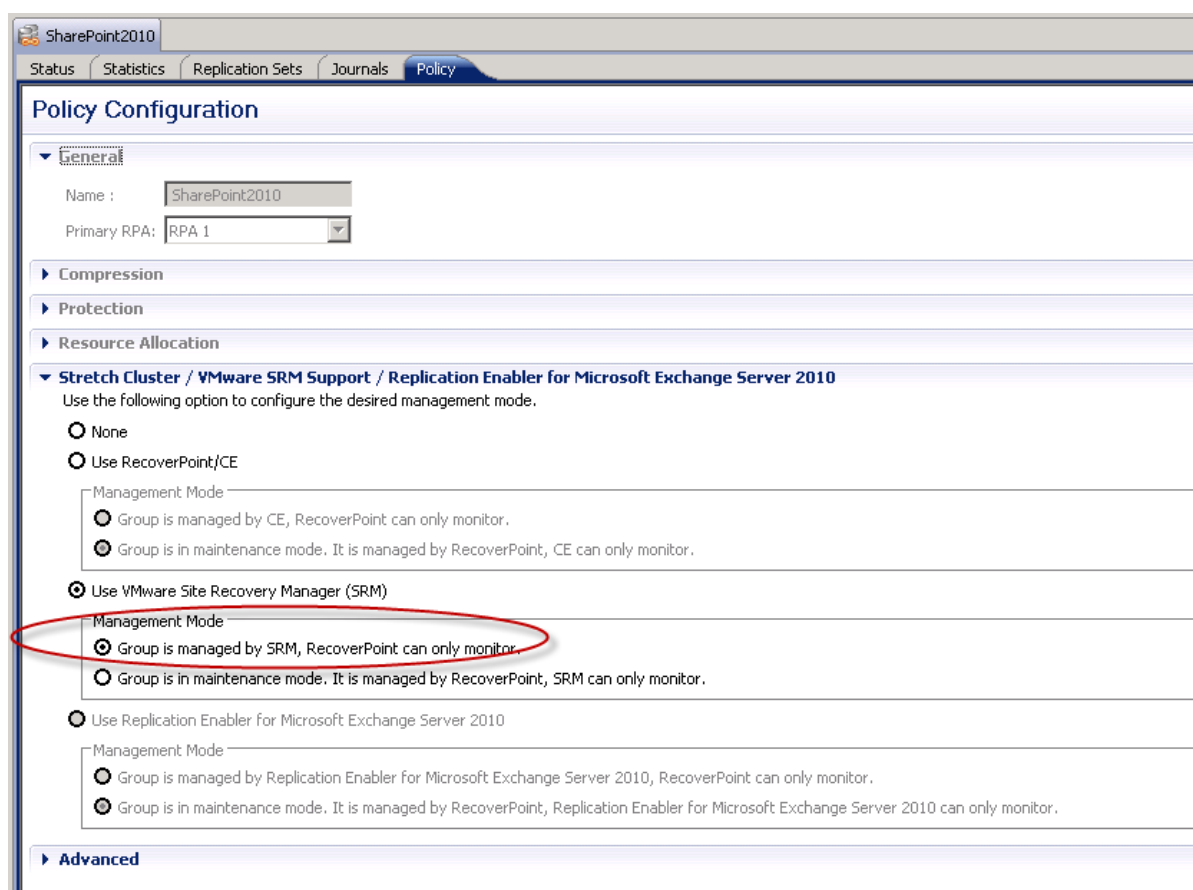


Figure 12. Configuration policy for asynchronous replication – VMware SRM support

Integrating RecoverPoint with VMware vCenter

With RecoverPoint’s integration with the VMware vCenter server, it is easy to display the relevant mappings between the RecoverPoint appliances, VMware vCenter servers, VMware vSphere 4.1 servers, and virtual machines by using the RecoverPoint management GUI.

This integration helps you to:

- Monitor VMware ESX server virtual machines that are managed by each specific vCenter server.
- Examine each ESX server (name and IP address) that is managed by the vCenter server.
- Display each virtual machine that belongs to the respective ESX server.
- View the VMware file system (VMFS) or Raw Device Mapping/P volume replication status (fully configured, partially configured, or not configured).
- Monitor additional local or remote vCenter servers.
- Add, update, rescan, or remove vCenter servers from the monitoring list.
- Display all the virtual machines that belong to a vCenter server in the tree view.

As shown in Figure 13, the virtual machines have been replicated by RecoverPoint, including the RDM devices and the OS volumes.

Name	IP	Consistency group	Copy	Replication set	Datastore
192.168.56.20	192.168.56.20				
192.168.56.210	192.168.56.210				
192.168.56.211	192.168.56.211				
WFE02	192.168.56.5				
naa.6006016082002b		SharePoint2010	Prod_SharePoint	RSet 47	
naa.6006016082002b		SharePoint2010	Prod_SharePoint	RSet 49	snap-62340bdb-WFE2_OS
Non VM Luns					
192.168.56.212	192.168.56.212				
App	192.168.56.8				
SQL	192.168.56.7				
Non VM Luns					

Figure 13. Virtual machine protected by RecoverPoint

RecoverPoint management within Unisphere

The architecture of Unisphere allows modules to be seamlessly integrated into its management interface. Because of this architecture, related EMC products can be managed from the same Unisphere session. With FLARE release 30, Unisphere provides RecoverPoint management. With Unisphere, the replication status as well as the storage environment can be managed. See Figure 14 for the RecoverPoint management interface in Unisphere.

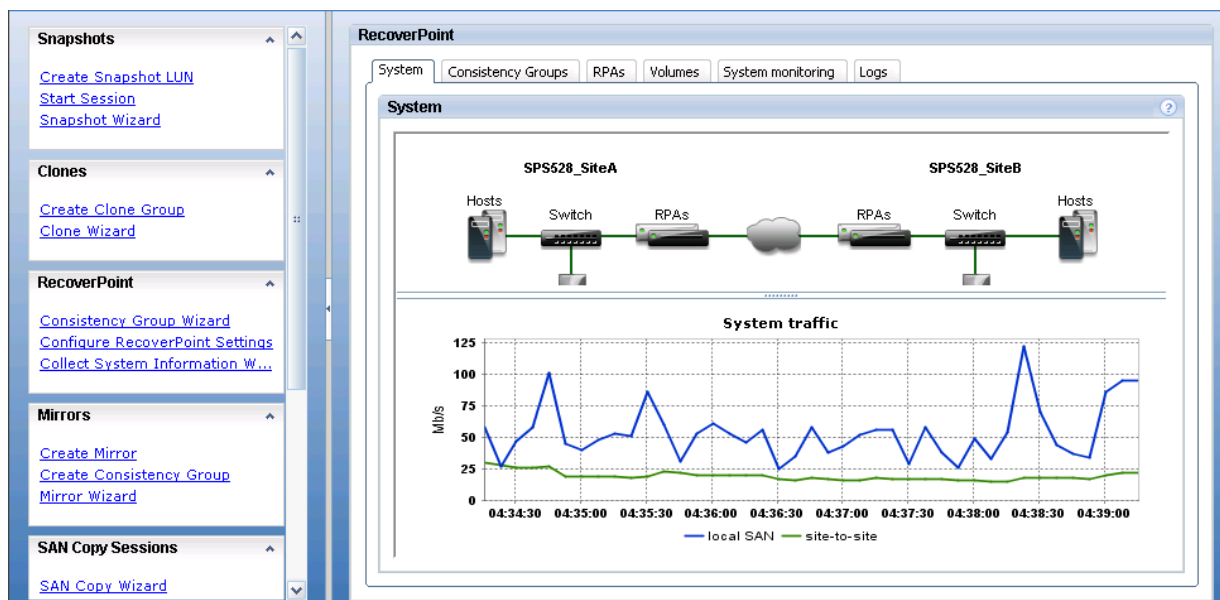


Figure 14. RecoverPoint management interface in Unisphere

VMware vCenter Site Recovery Manager design consideration

A VMware vCenter SRM protection group is a group of virtual machines that performs coordinated failovers. When VMware vCenter SRM performs a failover, it calls RecoverPoint to operate on all LUNs of virtual machines in the protection group. In this solution, all SharePoint virtual machines are configured as SRM protection groups to perform failovers to the DR site.

Figure 15 shows the array manager settings when datastore groups are detected by VMware SRM.

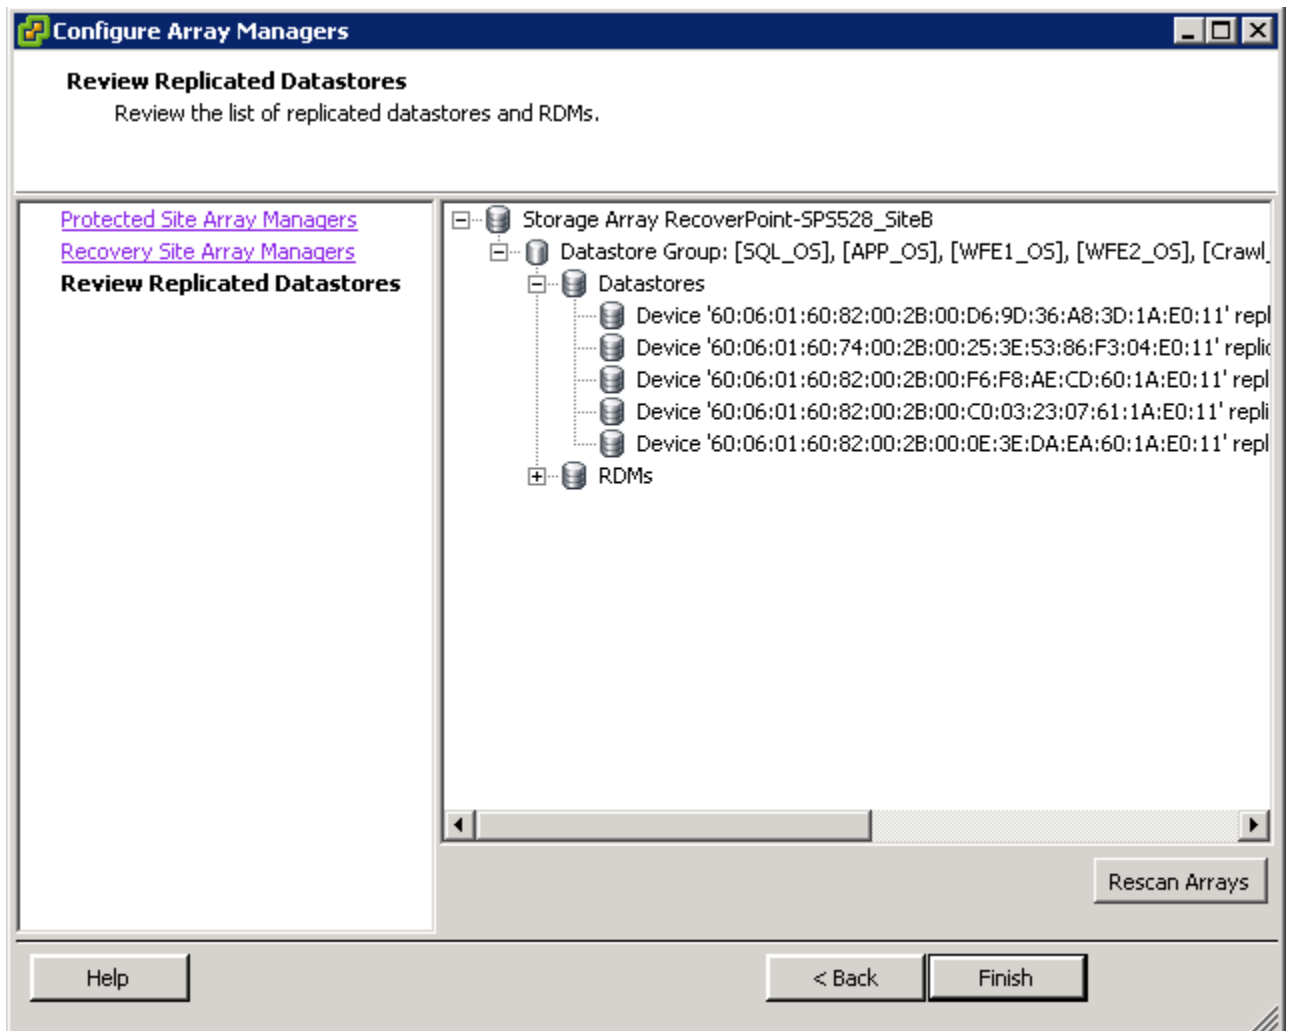


Figure 15. Datastore groups detected by VMware SRM in the array manager settings

Figure 16 shows five virtual machines are selected in this protection group to protect the SharePoint 2010 farm in the VMware SRM.

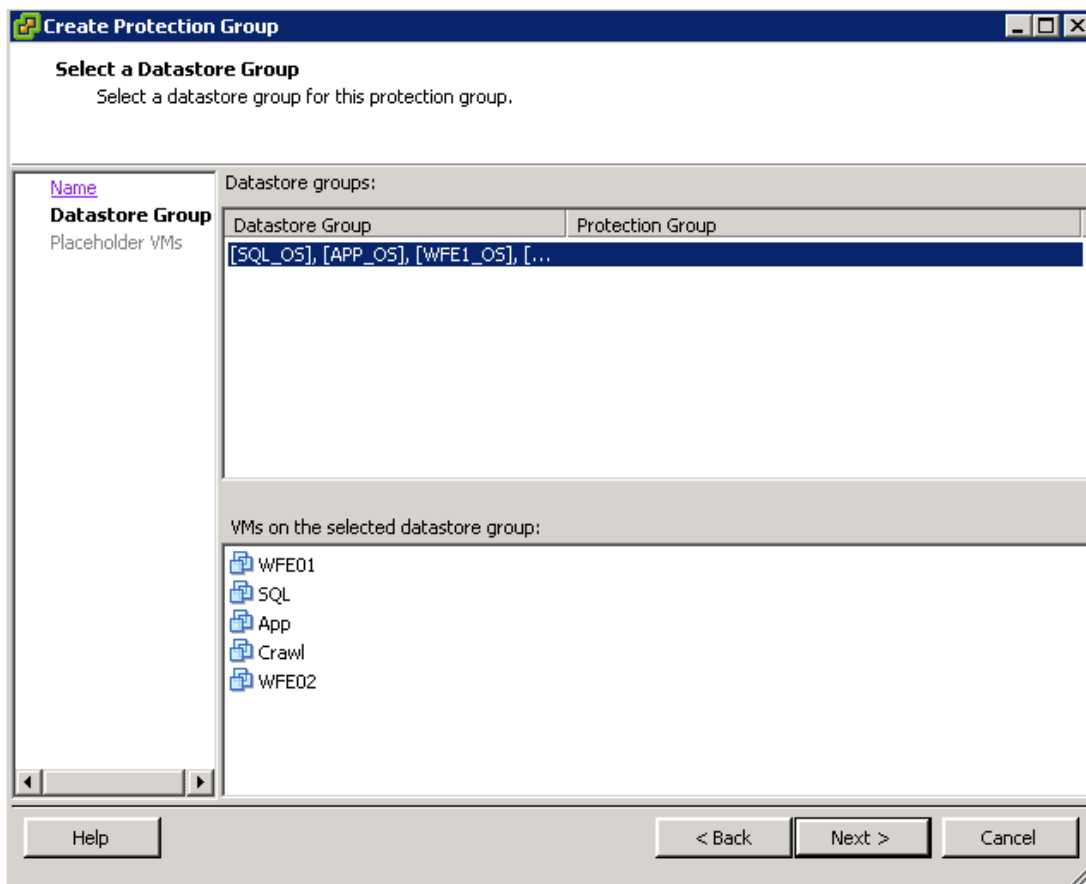


Figure 16. Protection group created to protect a SharePoint 2010 farm in VMware SRM

Failback consideration

When a recovery plan is running and the virtual machines in it are failed over to the recovery site, consider the following failback procedures:

- For a virtual machine that does not have RDM attached, use the EMC RecoverPoint Site Recovery Manager Failback plug-in as installed on the VMware vCenter server. For detailed configuration steps, refer to *EMC RecoverPoint Site Recovery Manager Failback Plug-in Technical Notes*.
- For other virtual machines, configure the original production site to be the recovery site, then run a failback recovery plan that can fail over the protected inventory from the original recovery site back to the original production site. Refer to *Site Recovery Manager Administration Guide vCenter Site Recovery Manager 4.1* for a detailed procedure.

Best practices for VMware Site Recovery Manager

Some best practices have applied to this solution to improve the recovery performance:

- Grouping virtual machines in fewer protection groups enables faster test and real recoveries, provided that those virtual machines have no constraints preventing them from being grouped under similar protection groups. The test results show that recovery time with five SharePoint virtual machines in one protection group was reduced by 2 minutes compared with the recovery time in using five protection groups (one SharePoint virtual machine for each protection group).
- In a recovery plan, the virtual machines being recovered can be assigned with high, normal, or low priority. The dependencies between virtual machines to be recovered should be clear so that only a certain number of required virtual machines can be assigned as high priority. It impacts the recovery time in tests as well as in real recovery. The recovery time increases as the number of virtual machines with high priority increases, because the operations of recovering high-priority virtual machines are executed sequentially. This applies to both real and tested recoveries. In this solution, all virtual machines are set to **Normal priority**, which decreased RTO by 1 minute, compared with setting the SQL virtual machine to **High priority** and others to **Normal priority**.
- The round trips from the Site Recovery Manager database to the Site Recovery Manager server may greatly impact the recovery time. It is recommended that the Site Recovery Manager database be installed as close to the Site Recovery Manager server as possible.
- In this solution, VMware DRS has been enabled on both protected and recovery sites to load balance the cluster during the recovery and failback operations. It is recommended to enable VMware DRS on the recovery cluster. By doing so, the site recovery manager can utilize VMware DRS for load balancing the recovered virtual machines across the hosts so as to ensure the optimal performance and RTO. It is easier to distribute shadow virtual machines across all hosts automatically.

Test methodology and scenarios

Test methodology

The data population tool uses a set of sample documents. Altering the document names and metadata (before insertion) can make each document unique.

One load-agent host is allocated for each WFE, allowing data to be loaded in parallel until the targeted 1 TB data size is reached. The data is spread evenly across the 10-site collections (each collection is a unique content database).

Table 8 lists the average size of the documents in the content databases.

Table 8. Average size of the documents in the content database

Type	Size (KB)
.doc	270.7
.docx	109.8
.xlsx	26.1
.mpp	254.5
.pptx	202.4
.jpg	95.0
.gif	76.5
.vsd	494.41

Test tools

To simulate load on the SharePoint farm, Microsoft Visual Studio Team System (VSTS) was used. A client load emulation tool provided by an independent third-party vendor (KnowledgeLake, Inc., an accredited Microsoft Gold Partner) was used to ensure that the SharePoint farm was operating at the optimal performance.

Load generation

To generate and emulate client load, Microsoft VSTS 2010 was used in conjunction with KnowledgeLake code to simulate real-world SharePoint user activity.

The VSTS team test rig consisted of a controller and two VSTS team agents. The controller evenly distributed client load across the agents, which evenly distributed the load across the WFEs.

SharePoint user profiles

The user profiles consist of a mix of three user operations: browse, search, and modify.

KnowledgeLake DocLoaderLite was used to populate SharePoint with random user data, while Microsoft VSTS 2010 emulated client user load. Third-party vendor code was used to ensure an unbiased and validated test approach.

During validation, a Microsoft heavy-user load profile was used to determine the maximum user count that the Microsoft SharePoint 2010 server farm could sustain while ensuring the average response times remained within acceptable limits. Microsoft standards state that a heavy user performs 60 requests in each hour; that is, there is a request every 60 seconds.

The user profiles in this testing consist of three user operations:

- 80 percent browse
- 10 percent search
- 10 percent modify

Note Microsoft publishes default service-level agreement (SLA) response times for each SharePoint user operation. Common operations (such as browse and search) should be completed within 3 seconds or less, and uncommon operations (such as modify) should be completed within 5 seconds or less. These response time SLAs were comfortably met and exceeded.

Test scenarios

Table 9 describes the design validation and the performance scenarios for this solution.

Table 9. Test scenario

Test scenario	Details
1	Baseline SharePoint farm performance test.
2	SharePoint farm performance test with FAST Cache.
3	SharePoint farm performance test with data compression
4	Local VMware High Availability and Distributed Resource Scheduler (DRS) clusters test.
5	Local data protection test using EMC Replication Manager and SnapView snap technology.
6	Local content database recovery using Replication Manager.
7	Single item recovery using Metalogix Selective Restore Manager.
8	Remote site failover and failback test using SRM 4.1 and RecoverPoint CRR.

Test results

Test 1: Baseline SharePoint farm performance test

The baseline performance test tested the maximum user load on the production SharePoint farm. This test was run for 8 hours with a full user load. FAST Cache was disabled and 800 GB of BLOB data was stored on the BLOB store.

The test shows that the maximum user capacity is 12,240 at 10 percent user concurrency. Refer to Table 10 for details of the test result.

Table 10. Test result of baseline VSTS performance

Test scenario	Requests per second (RPS)	Concurrency (%)	Maximum user capacity	Average user response time (in each second)		
				Browse	Search	Modify
8 hour test with FAST Cache disabled	20.4	10	12,240	2.42	1.37	1.27

Because of the full user load, CPU utilization on the WFE servers was the bottleneck. To achieve more user capacity, more WFE servers were scaled out in the test. Table 11 lists the CPU and memory usage of the production virtual machine and ESX servers.

Table 11. CPU and memory usage on the virtual machine and ESX servers

Server types	Average CPU process time (%)	Average memory usage (%)
SQL server	46.2	99
WFE01	85.5	27
WFE02	84.9	30
Crawler server	1.31	33
Application server	0.23	36
ESX server 1	51.6	20.55
ESX server 2	28.30	27.55
ESX server 3	47.29	19.17

Test 2: SharePoint farm performance test with FAST Cache

In this test scenario, FAST Cache was enabled on the search database and tempdb LUNs. The test was run for 8 hours with a full user load. It took 2.5 hours for FAST Cache to warm up. The write cache hit ratio of the property database is nearly 100 percent and the read cache hit ratio is 56 percent.

The test result is listed in Table 12.

Table 12. Matrix performance result with Fast Cache

Items	Performance with FAST Cache disabled	Performance with FAST Cache enabled	Improvement
Maximum user capacity at 10% user concurrency	12,240	13,080	6.8% capacity increase
RPS	20.4	21.8	6.8% RPS improve
Crawl performance	45 items/second	86 items/second	91% crawl performance improve
Browse average user response time (seconds)	2.42	2.31	4.5% average browse response time decrease
Search average user response time (seconds)	1.37	1.0	27% average search response time decrease

Besides the IOPS of the SharePoint search databases, the CPU threads of crawler servers and the slow responses of the crawled items can also affect the crawl performance.

Figure 17 shows the number of passed tests per second (passed tests/sec) in the VSTS performance test. The average number of tests passed tests per second was 21.8.




Figure 17. Number of passed tests/sec

Similar to Test 1, CPU utilization on the WFE servers was the bottleneck because of the full user load. The ESX servers and SharePoint virtual machine system performance is similar to

Test 1. To achieve more user capacity, more WFE servers were scaled out. By scaling out the WFE servers, more benefits of FAST Cache for crawling and querying were shown.

Test 3: SharePoint farm performance test with data compression

In this test scenario, data compression was enabled on the BLOB store LUNs with user load. The results showed:

- It took 7 hours to complete the initial compression for BLOB stores with four concurrent compression operations without user load. Data compression operation did not show any performance impact from the VSTS performance tests with a full user load.
- The reduction ratio reached 1.38:1 on average after the initial compression of the BLOB store. On average, 91.2 percent of the database resources were freed.

Table 13 lists the content databases consumed capacity on average after the initial compression.

Table 13. Consumed capacity after the initial compression

Database file size before RBS (GB)	Database file size after RBS (GB)	BLOB store capacity after RBS (GB)	Consumed capacity after initial compression (GB)	Reduction ratio	Percentage of original database size (%)
100	8.8	100	73	1.38:1	91.2

Figure 18 shows the compression LUN summary on CLARiiON CX4-120 after the initial compression.

The screenshot shows a window titled "Compressed LUNs Summary" with a "Compression Feature State: On" indicator. Below this is a "Filter for" input field. The main content is a table with the following data:

LUN Name	State	Percent Comp...	Rate	User Capacity ...	Consumed Ca...
BLOB_Store...	Compressed	100	Medium	150 GB	74 GB
BLOB_Store...	Compressed	100	Medium	150 GB	73 GB
BLOB_Store...	Compressed	100	Medium	150 GB	73 GB
BLOB_Store...	Compressed	100	Medium	150 GB	73 GB
BLOB_Store...	Compressed	100	Medium	150 GB	73 GB
BLOB_Store...	Compressed	100	Medium	150 GB	73 GB
BLOB_Store...	Compressed	100	Medium	150 GB	72 GB
BLOB_Store...	Compressed	100	Medium	150 GB	72 GB

Figure 18. Compression LUN summary

Test 4: Test for local VMware HA and Distributed Resource Scheduler clusters

This test was focused on the following major features of VMware in providing high availability to the SharePoint environment:

- vMotion through live virtual-machine migration
- VMware HA through sudden loss of one of the three ESX servers hosting SharePoint virtual machines
- DRS through hardware load balancing of poorly positioned virtual machines in a SharePoint farm

In this virtualized SharePoint 2010 farm, there were five virtual machines evenly distributed on three ESX cluster servers.

You can find a summary of the test results for vMotion, VMware HA, and DRS in Table 14.

Table 14. Test results for vMotion, VMware, and DRS

VMware Feature	Test Description	Test results
vMotion	One WFE virtual machine was failed over from one ESX server to another with full user load.	2 minutes 20 seconds.
High Availability	The ESX server on which the WFE virtual machine and the crawler virtual machine were running was restarted. The WFE and crawler were restarted on the other ESX server.	20 seconds
Distributed Resource Scheduler	To prove the value of DRS, the distribution of virtual machines was skewed across two of the three ESX servers and the DRS was set to Fully automated .	When the DRS recommendations were applied, one of the WFE virtual machines that was running on an ESX server with heavy load was failed over to an ESX server with lighter load.

Figure 19 shows the number of passed tests per second (passed tests/sec) after the redistribution of the virtual machines with VMware DRS feature.




Figure 19. Passed tests/sec after the redistribution of virtual machines

Test 5: Local data protection test using EMC Replication Manager and SnapView snap technology

In this test, the SharePoint 2010 farm replication synchronized 43 source LUNs to snapshots in a single job during full user load. The CLARiiON SnapView snap job was completed successfully in 27 minutes. The job includes the snapshot preparation, a SQL virtual device interface (VDI) online backup session across all 28 data volumes including BLOB stores, and one VSS replica job. The total source LUNs capacity is 4.1 TB.

During the full-farm replication, the crawl operation was also running. As shown in Figure 20, there was minimal-performance impact on the SharePoint farm during replication.





Figure 20. Passed tests/sec during the full-farm replication

The performance counter of virtual machines and ESX servers showed that there was no additional CPU or memory usage during the job. Figure 21 shows the CPU usage of the SQL server during the replication job. The average CPU usage was 37.695 percent during the replication. No additional SQL CPU utilization was caused by the backup job.



Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	SQL	Usage	Average	Percent	46.11	53.17	27.35	37.695

Figure 21. CPU usage of the SQL server during the full-farm replication

Test 6: Local content database recovery using Replication Manager

After the content database replica was created by using Replication Manager, the single content database can be easily restored to the production SQL server. The RTO for the content database varies depending on whether the RBS is enabled. The test shows the following results:

- Only 6 minutes were needed for restoring a 100 GB content database without RBS.
- It took 23 minutes to restore a content database with RBS FileStream enabled including 100 GB BLOB store and 9 GB content database data and log.

Before restoring the content database, all SharePoint timer services were stopped on the SharePoint servers. And then the targeted content database was selected from the replica on Replication Manager (see Figure 22).

Replica Time	Description	State	Source	Technology	Host	Job
2011 01 26 09:48:38 PM	sps2010	Successful	Primary Storage	SnapView Snap	Multi-host	sps201
2011 01 26 10:21:02 PM	sps2010	Successful	Primary Storage	SnapView Snap	Multi-host	sps201

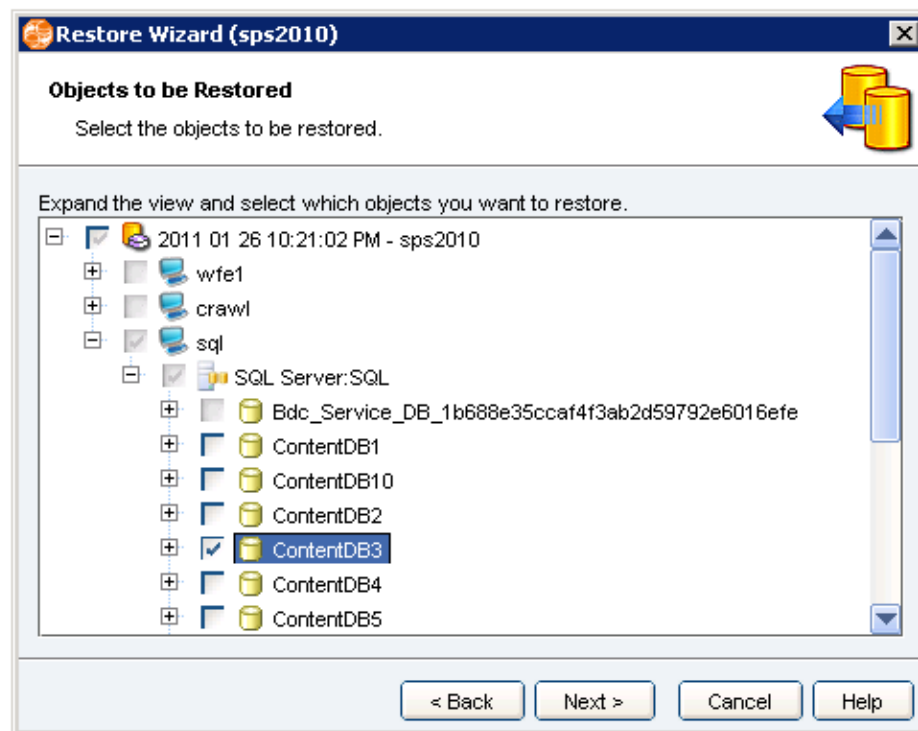


Figure 22. Restoring a single content database from the replica

While restoring the content database, the **Recovery** mode was selected as shown in Figure 23. The restore operation rolled back all uncommitted transactions. After the recovery process, the database was ready to be used.

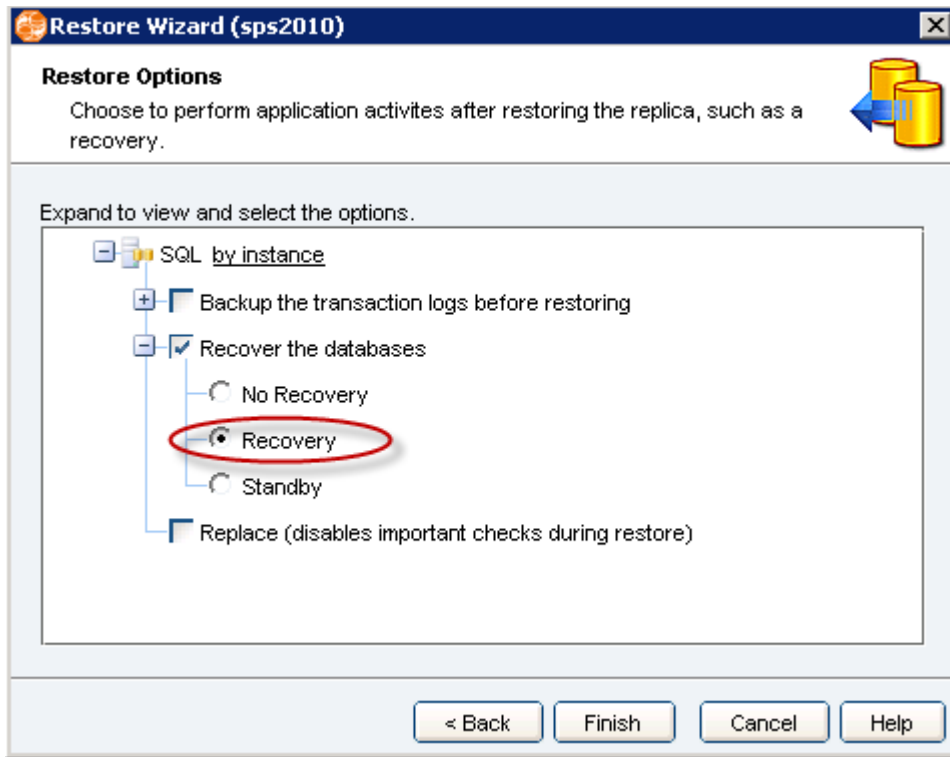


Figure 23. Recovery mode for content database restoring

For more detailed information about database recovery option, refer to *EMC Replication Manager Guide 5.3.1 – Product Guide*.

Test 7: Single item recovery with Metalogix Selective Restore Manager

In this test scenario, the following steps were followed to use Metalogix Selective Restore Manager to restore a SharePoint list item. The RTO for restoring a single item is less than 10 minutes:

1. Mount the content database replica through Replication Manager. It is not recommended to mount on a production host although Replication Manager provides an alternative path option to prevent collisions. It is recommended to have an extra mount host. In this solution, the Replication Manager server is the mount host.

To mount the content database replica, select Original path or Alternative path in the mount options. If the content database is enabled with SQL RBS with FileStream, it needs to be mounted with the original path in the mount option as shown in Figure 24. The mount job took 7 minutes to complete.




Figure 24. Original path under path options

2. Attach the content database to a non-production SQL instance after mount. Only seconds were needed to complete the database attachment.
3. Connect Selective Restore Manager to the SharePoint database, and select the attached content database as shown in Figure 25.

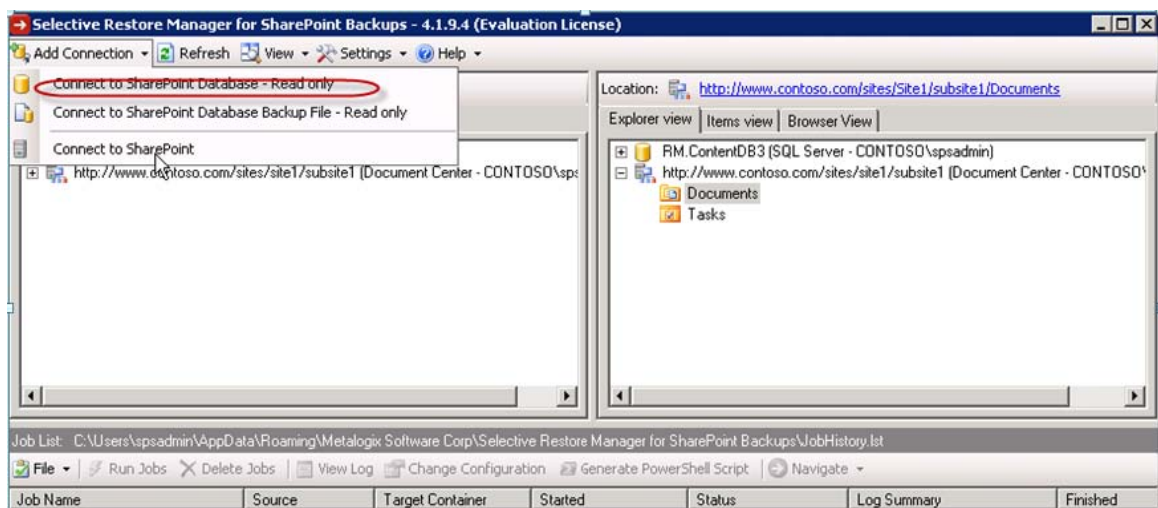


Figure 25. Metalogix Selective Restore Manager connection

4. Connect to the production SharePoint site. See Figure 26.

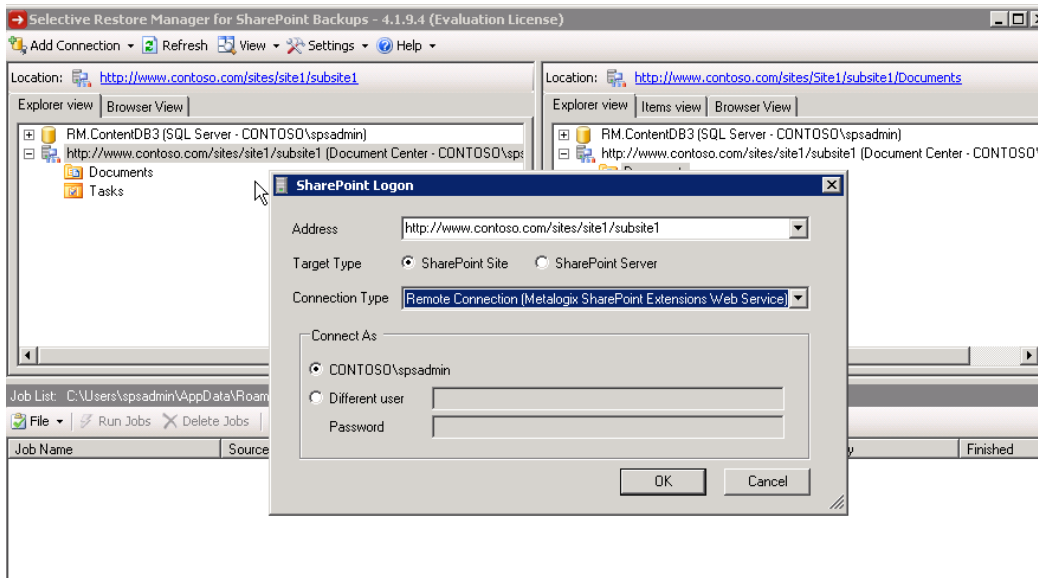


Figure 26. Connecting to the SharePoint site using Selective Restore Manager

5. Navigate to where the document is located at the SharePoint site in the explore view and restore the item in the item view. You can select **Save File and Metadata to Disk** from the list to restore the document to the disk. It took only 1 second to restore a 100 KB document to the production site. See Figure 27.

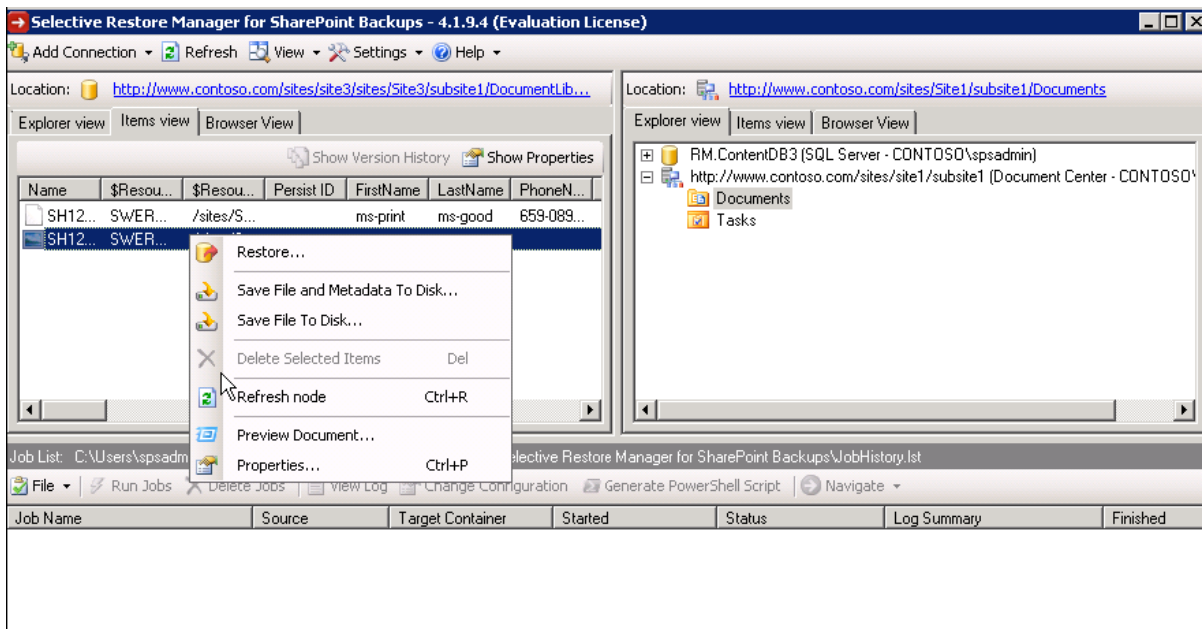


Figure 27. How to restore the document

Related metadata, versions, views, and permissions of that document can be preserved. For more features about Metalogix Restore Manager, refer to the online product description at <http://www.metalogix.net/products/selective-restore-manager-pro-for-sharepoint>.

Test 8: Remote site failover and failback using SRM 4.1 and RecoverPoint CRR

In this test scenario, all the SharePoint virtual machines (including one SQL server, two WFEs, one crawler server, and one application server) were failed over from the production site to the DR site with RecoverPoint CRR. The SharePoint farm was running with a full user load. The test shows the following results:

- 15 minutes 33 seconds were needed to fail over all the SharePoint virtual machines from the production site to the DR site.
- 16 minutes 8 seconds were spent running the test plan for failover.

Table 15 lists the detailed recovery time during the test and the operation of the failover plan. It shows that most of the time was spent in preparing storage of SQL virtual machines because it has 42 RDMS attached to the virtual machine. It took only 20 seconds for RecoverPoint to fail over to the disaster recovery site.

Table 15. Recovery time during testing and failover

Test type	Total time (minutes:seconds)	Time for shutting down the production virtual machines (minutes:seconds)	Time for preparing the storage (minutes:seconds)	Time for recovering normal priority virtual machines (minutes:seconds)	Maximum time lag (seconds)
Performing recovery plan for a fully operational farm serving user requests	15:33	00:33	11:45	3:15	3
Testing recovery plan for a fully operational farm serving user requests	16:08	00:00	11:36	4:32	3

As shown in Figure 28, the downtime of passed tests per second was caused by the failover of all the SharePoint virtual machines to the DR site. All the user connections to the SharePoint farm were down during the failover. After the entire SharePoint farm was recovered at the DR site by VMware SRM, the number of passed tests per second was back to normal. And the user requests to the SharePoint farm were back to normal.

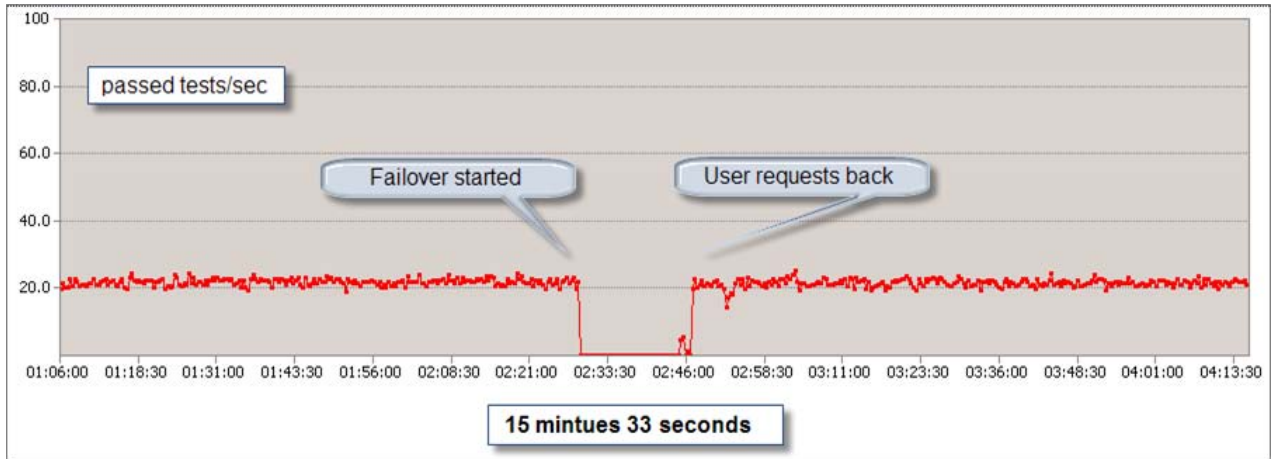


Figure 28. Passed tests per second during the failover

Table 16 lists the key performance matrix of RecoverPoint during the replication.

Table 16. Key performance matrix of RecoverPoint during replication

Consistency group	Initial synchronization duration	Average bandwidth reduction	Maximum time lag (seconds)	Capacity (TB)
SharePoint 2010	16 hours 26 minutes	3.35	3	4.79

Figure 29 shows the SharePoint 2010 replication performance with a full user load.




Figure 29. SharePoint 2010 replication performance with a full user load

At the same time, failback using the EMC RecoverPoint Site Recovery Manager Failback plug-in was also tested. It took only 4 minutes to fail back a virtual machine.

For the virtual machines with RDM devices, the production site and the DR site were switched for failback in the tests. The failback duration was similar to the failover duration.

Conclusion

Summary

This solution tested the performance of 1 TB of data with a BLOB store in a fully virtualized midsize SharePoint 2010 farm with the rapid replication of SharePoint data (in terabytes) by using EMC Replication Manager and EMC SnapView snap technology for local protection. This solution also integrated the VMware Site Recovery Manager with EMC RecoverPoint/SE for rapid, reliable, and cost-effective remote data protection.

Findings

The key findings of the proven solution testing are as follows:

- The SharePoint farm in the solution (including the SQL server) was virtualized depicting a midsize SharePoint environment with high availability.
- The sustained maximum user capacity was 13,080 at 10 percent user concurrency on the virtualized SharePoint farm.
- The average search and browse response time is improved after the EMC FAST Cache technology is enabled.
- Thirty percent of physical disk space was saved by using EMC LUN compression features on the SharePoint server BLOB store LUNs. With the SQL database datafile storage space, 91.2 percent of space was freed after enabling SQL RBS FileStream.
- A full-site disaster caused only 15 minutes 33 seconds of downtime while all SharePoint virtual machines were failed over from the production site to the disaster recovery site by leveraging EMC RecoverPoint/SE and VMware SRM.
- It took 27 minutes to replicate a 1.5 TB SharePoint 2010 farm by using EMC Replication Manager 5.3.1 with the EMC SnapView snap technology for local protection with no downtime to the application and negligible host resource impact.
- With EMC Replication Manager 5.3.1, it took only 6 minutes to restore a 100 GB content database from a replica.

Next steps

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

References

White papers

For white papers that provide similar solutions, see the following:

- *EMC CLARiiON, Celerra Unified, and VNX FAST Cache – A Detailed Review*
- *EMC CLARiiON SnapView Snapshots and Snap Sessions Knowledgebook*
- *EMC Data Compression – A Detailed Review*
- *EMC Virtual Infrastructure for Microsoft SharePoint Server 2010 Enabled by EMC CLARiiON and VMware vSphere 4 – A Detailed Review*
- *EMC Unisphere: A Unified Storage Management Solution*

Product documentations

For additional information on the products discussed in this white paper, see the following:

- *EMC Replication Manager 5.3.1 Product Guide*
- *EMC Replication Manager 5.3.1 Administrator's Guide*
- *EMC RecoverPoint Administrator's Guide*
- *EMC RecoverPoint Site Recovery Manager Failback Plug-in Technical Notes*
- *Site Recovery Manager Administration Guide –vCenter Site Recovery Manager 4.1*
- *EMC RecoverPoint Installation Guide*
- *VMware vCenter Site Recovery Manager Performance and Best Practices for Performance*