



EMC Business Continuity for Microsoft Office SharePoint Server 2007

Enabled by EMC CLARiiON CX4,
EMC RecoverPoint/Cluster Enabler,
and Microsoft Hyper-V

Reference Architecture

EMC Global Solutions



Microsoft | Virtualization

Copyright © 2010 EMC Corporation. All rights reserved.

Published May 2010

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

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

All performance data contained in this report was obtained in a rigorously controlled environment. Results obtained in other operating environments may vary significantly.

EMC Corporation does not warrant or represent that a user can or will achieve similar performance expressed in transactions per minute.

No warranty of system performance or price/performance is expressed or implied in this document. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

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

Part number: H7041

Table of Contents

Reference architecture overview	4
Document purpose	4
Solution purpose	4
The business challenge	4
The technology solution	5
Key components	6
Introduction.....	6
EMC CLARiiON CX4-240	6
EMC RecoverPoint.....	6
EMC RecoverPoint / Cluster Enabler.....	7
EMC PowerPath / Virtual Edition	8
Microsoft SharePoint Server 2007	8
Microsoft Windows Server 2008 R2 with Hyper-V	8
Microsoft Failover Clusters.....	8
Microsoft System Center Virtual Machine Manager 2008 R2.....	8
Physical architecture	9
Architecture diagram	9
Validated environment profile	10
Profile characteristics	10
Hardware	10
Software	11
Content mix	11
Testing methodology.....	12
Introduction.....	12
User response times	12
Maximum user capacity	12
Test results.....	13
RecoverPoint splitter performance testing	13
Distance testing.....	13
Distance testing with query server scaling.....	14
Migrations.....	15
Failover.....	16
Bandwidth reduction.....	17
Conclusion.....	18
Summary.....	18
Findings.....	18
Next steps	19

Reference architecture overview

Document purpose

Microsoft Office SharePoint Server is increasingly considered a business-critical platform and as such should have the highest levels of availability under all circumstances—for example, a site disaster.

This document describes the reference architecture for the EMC® Business Continuity for Microsoft Office SharePoint Server 2007 solution, enabled by EMC CLARiiON® CX4, EMC RecoverPoint with EMC RecoverPoint/Cluster Enabler, and Microsoft Windows 2008 R2 Hyper-V.

Solution purpose

The purpose of this solution is to demonstrate the value of EMC's RecoverPoint / Cluster Enabler (RecoverPoint/CE) product suite in providing fully automated disaster recovery (DR) for enterprise-class virtualized SharePoint farm environments.

Using this technology, it is now possible to stretch existing Microsoft Failover Clusters to provide geographically-separated high availability and DR. Both the failover and failback times demonstrated by the solution were under 10 minutes. This is significantly less than what can be achieved with native Microsoft tools.

With RecoverPoint's data reduction and compression features, SharePoint environments can now be replicated across longer distances, using less bandwidth, and customers can expect automated disaster recovery within minutes.

The business challenge

Disaster recovery for federated applications and environments becomes increasingly difficult to achieve with consistency as the applications and environments grow. As a federated application, SharePoint requires that all server roles, configurations, and data are consistent across the farm.

Challenge	Solution
Implementing and sustaining DR for Enterprise SharePoint environments is both complex and difficult.	The solution enables fully automated failover. Once configured, failover becomes automatic and planned failback is quick, simple, and causes minimal disruption.
Adjusting to dynamic workloads as data volume grows and user workloads change.	Server virtualization allows for simplified configuration and provisioning and rapid modification of the SharePoint farm when and where necessary.
Maintaining SharePoint search consistency in a DR scenario.	RecoverPoint/CE's support of full farm failover enables the SharePoint farm to resume with crash consistency. It minimizes lengthy re-indexing and degraded search capabilities during failover or failback processes.

Enterprise SharePoint environments can stretch to tens of servers, with differing roles, such as Index, SQL servers, and application servers. Writing DR plans to meet failover and recovery SLAs becomes an arduous task and failovers become unreliable or fail. Typically, with legacy solutions, critical working components of a SharePoint farm, such as a valid search content index, have to be rebuilt after a site failure.

The technology solution

This solution describes a virtualized Microsoft Office SharePoint Server 2007 enterprise farm environment, protected with remote disaster recovery and fully automated failover, enabled by EMC technology.

The environment consists of a six-node Hyper-V Windows Failover Cluster, with three active nodes (production site) and three passive nodes (DR site). The cluster contains the entire host infrastructure required to operate an Office SharePoint Server 2007 farm—for example, domain controllers, application servers, Web front ends (WFEs), and SQL servers.

SharePoint Server 2007 uses Microsoft SQL Server 2008 as its data store.

Microsoft Windows 2008 R2 Enterprise with Hyper-V provides the virtualization platform to maximize hardware utilization and improve SharePoint performance and availability. Hyper-V enables virtual machine (VM) high availability through Microsoft Windows Failover Clustering (WFC) and provides both Live and Quick Migration features. In synchronous replication mode, the solution supported Live Migration of VMs between sites with minimal disruption to the availability of the VM.

EMC CLARiiON CX4-240 arrays provide consolidated, managed, and highly-available storage for both production and DR sites. In addition, they provide a built-in mechanism (CLARiiON splitter) for RecoverPoint to provide continuous remote replication (CRR) of production data to the DR site.

In the solution, RecoverPoint CRR was tested in both available modes, as follows:

- Synchronous and asynchronous replication over a Fibre Channel (FC) inter-site link—distances up to 300 km or latencies up to 3 ms were achieved.
- Asynchronous replication over an IP WAN link—distances up to 2,500 km or latencies up to 25 ms were achieved.

Fully integrating with Microsoft's Windows Failover methodologies, RecoverPoint/CE ensures that a complete site failover can happen with minimal downtime and zero user intervention. Planned failover of a VM from site to site is now possible from Microsoft's Failover Cluster Manager. The operator does not need to be trained in RecoverPoint or CLARiiON technologies to achieve this.

Integration of this solution in an existing CLARiiON CX4-240 SharePoint environment requires minimal downtime (minutes to convert cluster groups) and a minimal footprint per site: two server rack units (U) for RecoverPoint appliances (RPAs), eight 4/8 GB FC ports, four 1 GbE network ports, and RecoverPoint/CE software installed on all cluster nodes.

Key components

Introduction

This section identifies and briefly describes the key solution components, as follows:

- EMC CLARiiON CX4-240
 - EMC RecoverPoint
 - EMC RecoverPoint/Cluster Enabler (RecoverPoint/CE)
 - EMC PowerPath[®]/Virtual Edition (PP/VE)
 - Microsoft SharePoint Server 2007
 - Microsoft Windows Server 2008 R2 with Hyper-V
 - Microsoft Failover Clusters
 - Microsoft System Center Virtual Machine Manager 2008 R2 (SCVMM)
-

EMC CLARiiON CX4-240

EMC CLARiiON CX4 model 240 provides a powerful networked storage system that scales seamlessly (up to 231 TB of capacity) so more applications can be consolidated. The CLARiiON CX4-240 combines CLARiiON's proven five 9s (99.999 percent) availability with innovative technologies such as Flash drives, Virtual Provisioning[™], a 64-bit operating system, and multi-core processors.

The CLARiiON CX4-240 provides UltraFlex[™] technology with dual-protocol (iSCSI and FC), online-expandable connectivity options, and the ability to integrate future technologies. With the CLARiiON CX4-240, customers can exploit CX4 advanced functionality to protect data from unexpected events and to decrease energy use.

EMC RecoverPoint

EMC RecoverPoint is a comprehensive replication and data protection solution that provides continuous remote replication, continuous data protection, and concurrent local and remote data protection for SAN volumes residing in a CLARiiON array. RecoverPoint runs on an out-of-band appliance and combines industry-leading continuous data protection technology with a bandwidth-efficient, no-data-loss continuous replication technology, enabling it to protect data both locally and remotely.

RecoverPoint CRR can operate in two modes, as follows:

- Synchronous replication over an FC inter-site link. Remote data is kept in sync with the production site, enabling zero RPO and minute-based RTO.
- Asynchronous replication over an IP WAN link. Remote data is asynchronously sent to the remote site. RecoverPoint's bandwidth reduction feature allows for replication over longer distances and/or slower WAN links. Metropolitan to intercontinental distances can be achieved.

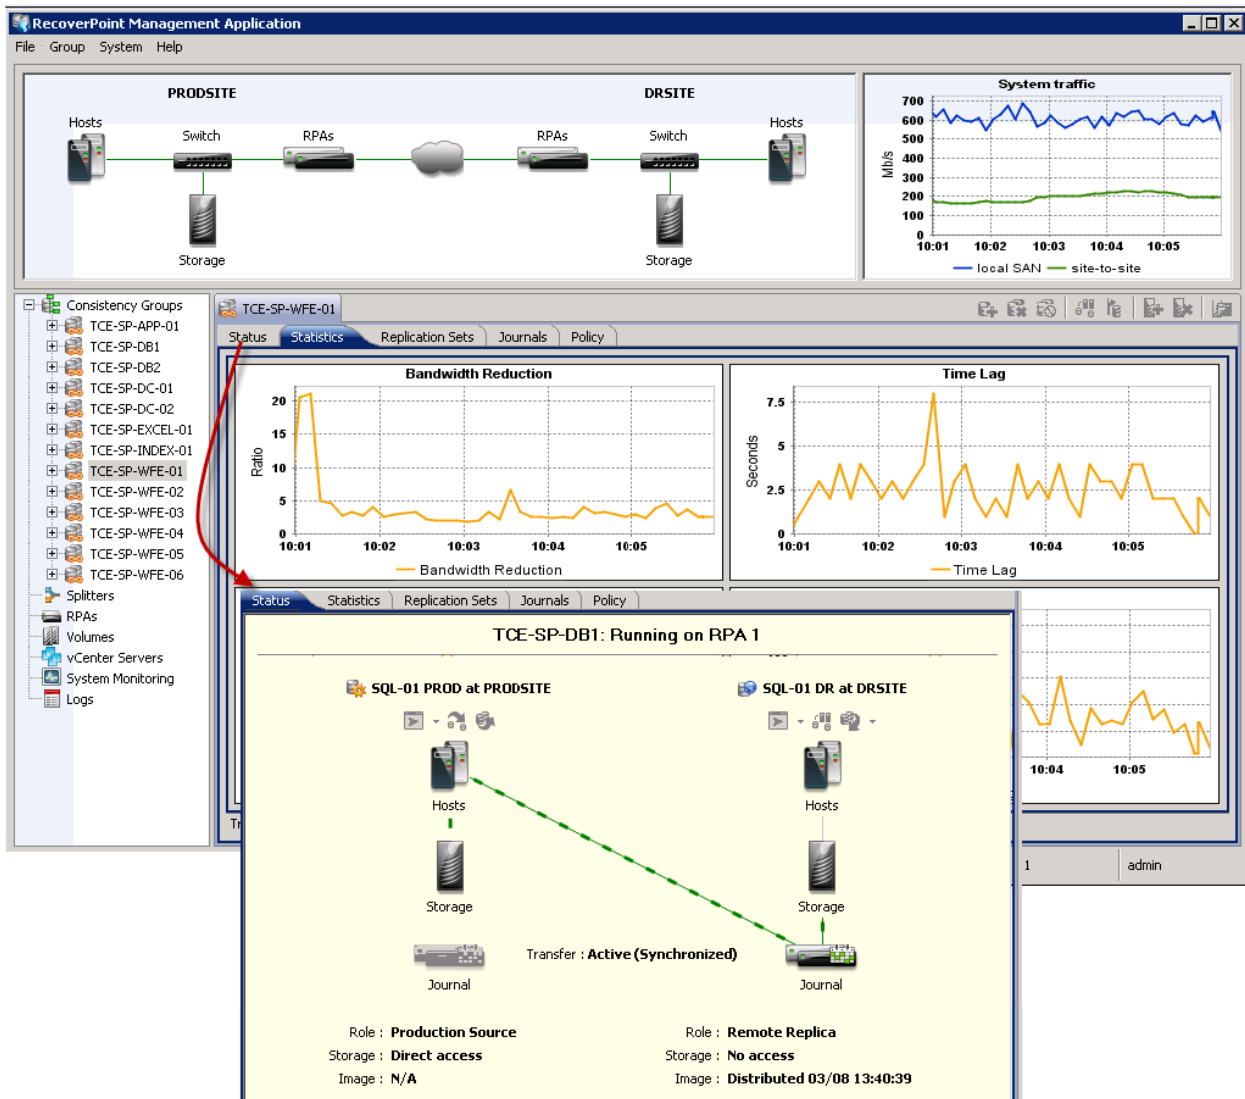
EMC CLARiiON CX4 arrays include the array-based RecoverPoint splitter function, which simplifies deployments (integration) and installation, since no additional fabric components are required. The CLARiiON array-based splitter, which is built into FLARE[®], also provides protection for FC and/or iSCSI environments.

EMC RecoverPoint / Cluster Enabler

RecoverPoint/CE is a software package that provides tight integration with Microsoft Failover Clusters to enable continuous remote replication between geographically-dispersed cluster nodes. RecoverPoint/CE works seamlessly with applications designed to take advantage of Failover Clusters, such as Microsoft Windows Server 2008 and Hyper-V environments.

RecoverPoint/CE enables existing Microsoft Failover Cluster customers to extend the protection of their current solution to include site disaster recovery, and enables existing RecoverPoint customers to reduce recovery time objectives (RTO) by allowing Microsoft Failover Clusters to automate resource and application failover between sites.

The following image shows the RecoverPoint Management Application GUI.



**EMC
PowerPath /
Virtual Edition**

EMC PowerPath/Virtual Edition (PP/VE) delivers EMC PowerPath multipathing features to optimize the VMware vSphere™ and Microsoft Windows Server 2008 Hyper-V virtual environments. With PP/VE, standardized path management can be achieved across heterogeneous physical and virtual environments.

PP/VE enables automated optimal server, storage, and path utilization in a dynamic virtual environment. This eliminates the need to manually load-balance hundreds or thousands of VMs and I/O-intensive applications in hyper-consolidated environments.

**Microsoft
SharePoint
Server 2007**

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

**Microsoft
Windows
Server 2008 R2
with Hyper-V**

Hyper-V is a hypervisor-based virtualization technology from Microsoft that makes it easier than ever to take advantage of the cost savings of virtualization through Windows Server 2008 R2.

Hyper-V enables customers to make the best use of server hardware investments by consolidating multiple server roles as separate VMs running on a single, physical machine. It can also efficiently run multiple different operating systems (Windows, Linux, and others) in parallel, on a single server, and fully use the power of x64 computing.

**Microsoft
Failover
Clusters**

Microsoft Failover Clusters is the clustering extension to Windows Server 2008 R2. It enables up to eight servers, running the same Windows operating system in the same domain, to be connected to a shared storage system.

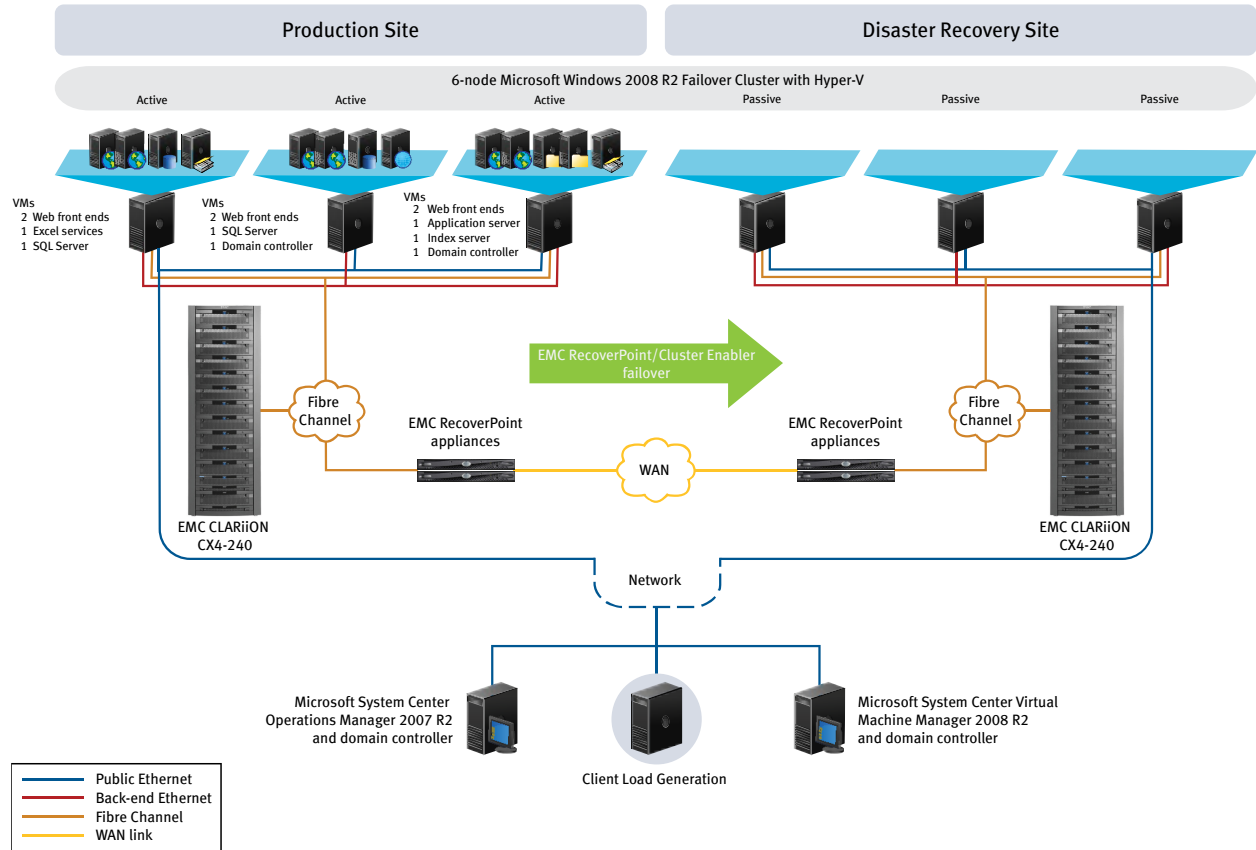
**Microsoft
System Center
Virtual Machine
Manager 2008
R2**

Microsoft System Center Virtual Machine Manager (SCVMM) 2008 R2 is a unified, multi-vendor management solution for the virtualized data center that helps to enable centralized administration of physical and virtual assets, increase server utilization, and provide dynamic resource optimization of virtual IT infrastructure.

Physical architecture

Architecture diagram

The following diagram depicts the solution's overall physical architecture.



CL4465

Validated environment profile

Profile characteristics

EMC validated the solution with the following environment profile.

Profile characteristic	Value
SharePoint farm user data	1.5 TB
Concurrency	1%
Site collections	1
Sites per site collection	15
Microsoft SQL Server 2008 (VMs)	2
Hyper-V cluster (physical)	6 nodes (3 production / 3 DR)
Web front ends (VMs)	6 (also running Query role)
Excel Services (VM) (hosting Central Admin)	1
Index server (VM)	1
Application server (VMs)	1
Domain controllers (VMs)	2
System Center Virtual Machine Manager (SCVMM) (physical)	1
RecoverPoint consistency groups (CGs)	1 CG per VM

Hardware

The following table lists the hardware used to validate the solution in this document.

Equipment	Quantity	Configuration
Storage array (Production site)	1	CLARiiON CX4-240 with: <ul style="list-style-type: none">8 FC ports per SPFLARE 2944 * 300 GB / 15k rpm FC drives
Storage array (DR site)	1	CLARiiON CX4-240 with: <ul style="list-style-type: none">8 FC ports per SPFLARE 2952 * 300 GB / 15k rpm FC drives
Distance emulation device	1	Anue H-series GEM and FC dual-blade Network Emulator software v3.3.5
Network switch	4	48-port trunkable network switches (2 production / 2 DR)
FC switch	3	48-port FC 4 GB switch (2 production / 1 DR)
Hyper-V server	6	16-core, 48 GB RAM
Infrastructure server	2	8-core, 16 GB RAM
RecoverPoint appliance	4	Gen-3

Software

The following table lists the software used to validate the solution in this document.

Software	Version
Windows Server 2008 Enterprise Edition R2	RTM
Microsoft Hyper-V	2008 R2
Microsoft SQL Server 2008	64-bit Enterprise Edition SP1
Microsoft Office SharePoint Server 2007	SP2
Microsoft SCVMM 2008 R2	RTM
PowerPath (with VE capabilities)	5.3 SP1
Visual Studio Test Suite 2008	SP1
KnowledgeLake Doc loader	1.1
EMC RecoverPoint	3.2 SP2 Patch 2
EMC RecoverPoint/Cluster Enabler	4.0.1

Content mix

The following table lists the content mix used for the solution.

File type	Average size (KB)
.doc	251
.docx	102
.xls	820
.xlsx	20
.ppt	2,021
.pptx	189
.jpg	93
.gif	75
.vsd	471
.mpp	235

Testing methodology

Introduction

The SharePoint user load was simulated using three key activities: browse, search, and modify. The testing methodologies used adhere to Microsoft recommendations and guidelines for SharePoint performance testing.

User response times

The following table details the acceptable ([Microsoft guidelines](#)) user response times for SharePoint operations.

In SharePoint performance testing it is vital that these response times are adhered to. If the response times are exceeded then the test is deemed to have failed.

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

Maximum user capacity

The maximum user capacity is the number of user requests that the SharePoint farm is able to service concurrently while maintaining the response times noted in the previous table.

The maximum user capacity is derived from the following formula:

$$\# = \text{seconds per hour} / \text{user RPH} / \text{Concurrency\%} * \text{RPS}$$

Examples:

$$3,600 / 60 / 1\% * 41.6 = 249,600 \text{ (supported user capacity for 1\% concurrency)}$$

$$3,600 / 60 / 10\% * 41.6 = 24,960 \text{ (supported user capacity for 10\% concurrency)}$$

Notes

All users were run against a Microsoft heavy user profile that performed 60 requests per hour. Also zero percent think time was applied to all tests—that is, typical user decision time was eliminated when browsing, searching, or modifying in Office SharePoint Server. Each user request was completed from start to finish without user pause, therefore creating a continuous workload on the system.

Test results

RecoverPoint splitter performance testing

One objective of the testing performed in this solution was to determine the performance impact of RecoverPoint replication on a very active SharePoint farm.

A suite of tests was performed before and after RecoverPoint replication was enabled on the CLARiiON in order to determine the impact of the CLARiiON-based RecoverPoint splitter on both the CLARiiON and the SharePoint application.

The following table details key performance counters with and without RecoverPoint replication enabled. The user profile was 80% Browse, 10% Search, and 10% Modify.

Replication enabled	User concurrency	Passed tests per second	Max user capacity	CLARiiON SP utilization	Average user response time
No	10%	41.4	24,840	12.9%	Browse = 1.23 sec Search = 1.09 sec Modify = 1.15 sec
Yes	10%	41.6	24,960	13.8%	Browse = 1.21 sec Search = 1.14 sec Modify = 1.16 sec

As noted in the table, there is a negligible performance impact on the SharePoint farm in both RPS and user response times. CLARiiON SP utilization was marginally higher with the splitter enabled.

Neither CLARiiON CX4-240 was stressed during any of the tests, and LUN utilization was within acceptable parameters.

Distance testing

Various distance and bandwidth combinations were tested during full SharePoint user load in order to determine the minimal inter-site replication link requirements to support this solution.

The following table shows those distances achieved between the production and DR sites. As shown in the table, the WAN throughput requirements for SharePoint are significant. This was due more to systematic farm operations (for example, timer jobs and index propagations) than to user load.

RecoverPoint/CE supports synchronous (FC) and asynchronous (IP) replication between sites so customers can use their existing link technology to implement this solution.

Note

Latency to distance conversions were derived using the standard formula of 1 ms per 100 km round-trip distance.

Distance scenario	Achieved results		
	Round-trip distance (km)	Latency (ms)	Bandwidth (Mb/s)
Asynchronous replication (IP)			
Baseline (same site)	0	0	300
City-to-city, shorter distance	400	4	500
State-to-state / inter-country (Europe)	1,600	16	900
Synchronous replication (FC)			
Baseline (same site)	0	0	1,000
Metro CWDM / DWDM	100	1	1,000

Distance testing with query server scaling

As part of the test observations, it was found that the query servers were responsible for utilizing most of the available bandwidth. This is because the Index server propagation to the query servers (which occurs after incremental crawls) causes very large burst writes to the production disks of the query servers. These writes need to be replicated to the DR site and therefore require a significant amount of bandwidth.

As a result of this finding the number of query servers was scaled down and tests were performed to determine the bandwidth requirements of SharePoint farms with differing numbers of query servers.

The following table shows the distances achieved between the production and DR farms when the query server roles are reduced on the WFEs.

Number of query servers	Round-trip distance (km)	Latency (ms)	Bandwidth (Mb/s)
Asynchronous query server scaling			
6	1,600	16	900
4	1,800	18	900
2	2,500	25	900
Synchronous query server scaling			
6	100	1	1,000
4	200	2	1,000
2	300	3	1,000

Migrations

The following table shows the time that it takes for VMs to perform Hyper-V Live and Quick migrations, both inter- and intra-site. The tests were performed with both replication modes: synchronous at 1,000 Mb/s and 1 ms latency, and asynchronous at 300 Mb/s and 0.5 ms latency.

VM role	Failover type	Synchronous		Asynchronous	
		VM downtime (seconds) *	Cluster migration time (seconds) **	VM downtime (seconds) *	Cluster migration time (seconds) **
Live migration					
Domain controller	Local to remote	45	35	105	63
Web front end	Local to remote	210	38	210	60
SQL Database	Local to remote	105	27	210	52
Web front end	Local to local	30	1	30	1
Quick migration					
Domain controller	Remote to local	105	26	75	26
Web front end	Remote to local	210	26	105	26
SQL Database	Remote to local	750	27	570	55
Web front end	Local to local	135	1	90	1

* VM downtime: The amount of time that the VM is unresponsive to application requests.

** Cluster migration time: The amount of time required by the cluster to fail over the cluster group (a subset of the overall VM downtime). This is the time from an offline request to a VM cluster group on the source node to the time that the VM online action completes on the target node.

Notes

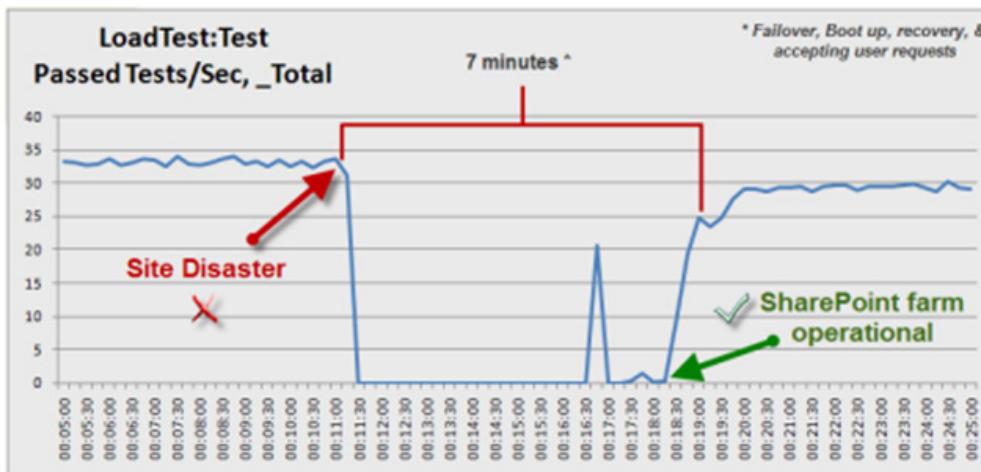
Live migration: VMs with large memory configurations take longer to migrate than VMs with smaller memory configurations. This is because active memory is copied over the network to the receiving cluster node prior to migration. Live migration requires high bandwidth, low latency networks in order to function.

Quick migration: This migration type is essentially a suspend-to-disk and then a resume operation. VM memory must be committed to disk before migration. This causes a significant number of burst writes, which need to be replicated quickly, and so requires high bandwidth WAN links between sites.

Failover

To demonstrate a fully automated SharePoint farm failover, one of the tests executed was an entire production site disaster, through a simulated crash.

The following image shows an example where a site outage occurred in an asynchronous replication configuration with a WAN link of 300 Mb/s and with 0.5 ms latency.



Test result highlights

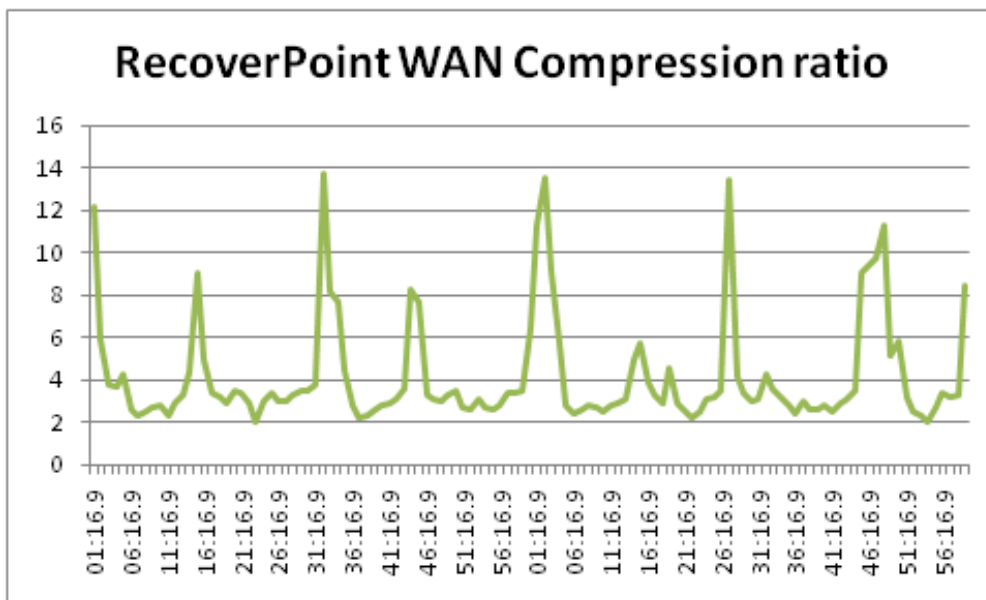
Total SharePoint application downtime due to the outage was just seven minutes. No user intervention was required for the application to recover on the DR site.

Using standard Microsoft failover clustering methodologies (Preferred Owner model) ensured that, on failover, the SharePoint virtual environment was evenly balanced across cluster nodes on the DR site, again with no user action required.

Bandwidth reduction

In asynchronous continuous remote replication mode, over IP WANs, RecoverPoint has the ability to do both data reduction and compression, ensuring maximum write I/O replication performance from available WAN throughput. Only changed I/O blocks are replicated, and RecoverPoint compresses the data before sending it over the WAN.

Typical bandwidth compression ratios depend on the nature of the data but are normally expressed in factors of 1 to 10 times. For this solution, a seed set of generic documents (for example, .xlsx, .jpg, and .docx) was used, with a compression setting of 1 (the highest compression level). An average compression ratio of 4.2 times was achieved. This is illustrated by the following graph.



Conclusion

Summary

This reference architecture depicts a validated virtualized Microsoft Office SharePoint Server 2007 farm that is enabled by Hyper-V, Microsoft Windows failover clustering and RecoverPoint with Cluster Enabler technology. EMC's CLARiiON CX4-240 was key in providing both consolidated storage and splitter replication technology.

Designing and implementing an automated disaster recovery solution for an enterprise Office SharePoint Server 2007 farm using native Microsoft technology is a daunting, if not impossible, task. With native Microsoft technology, recovering from a disaster is a labor-intensive task that requires many hours of reconfiguration and possibly many days of index crawls.

RecoverPoint with Cluster Enabler provides best-in-class disaster recovery for active SharePoint farms. You can now expect automated disaster recovery, within minutes (RTO), with minimal (asynchronous) to zero (synchronous) SharePoint data loss (RPO), even in the event of a full site disaster.

This solution also saves on licensing costs (no VM OS licenses on the DR site) and on time for updating and maintaining a mirror DR farm (software / configuration updates).

Time and attention need to be invested to understand the current and future requirements of a customer's Office SharePoint Server 2007 farm. Microsoft SharePoint Server is increasingly considered as a business-critical platform and as such should have the highest levels of availability under all circumstances—for example, a site disaster.

Findings

The key findings of solution testing were as follows:

- A full site disaster caused only seven minutes of downtime for users connecting to the SharePoint farm (for both synchronous and asynchronous replication).
- Planned failback of the SharePoint had a cumulative downtime of only two minutes.
- The impact of the RecoverPoint splitter on the CLARiiON's SP utilization was minimal at ~1 percent under full SharePoint user load.
- Enabling the CLARiiON RecoverPoint splitter had negligible impact on application performance.
- Bandwidth reduction by RecoverPoint was significant (4.2 times).
- Quick Migration of VMs with large amounts of memory (for example, SQL server) requires sufficient bandwidth and low enough latency to manage the large amounts of data being committed to disk in burst writes and requiring quick replication.
- In the solution, query servers had the busiest disks in the SharePoint farm and were therefore the highest users of replication bandwidth. Reducing the number of query servers allowed for longer distances between the production and DR sites.

The following table demonstrates the advantages of using EMC technology over native tools.

Description	Native tools	EMC technology	Notes
Entire farm crash consistency on failover	No	Yes	RecoverPoint's Consistency Group feature ensures write order fidelity across all disks in the event of a crash.
Automated full farm failover	No	Yes	RecoverPoint/CE allows for Microsoft clusters to stretch beyond the local data center.
Farm failover time	Hours / Days	7 Minutes	RecoverPoint minimizes downtime by using Microsoft's WFC to fail over the entire farm to a DR site where the disks have a crash consistent copy ready at all times.
Search Index consistency on failover	No	Yes	RecoverPoint write order consistency ensures that the Index server's content is consistent with its Shared Services Provided (SSP) search database.
Ability to incrementally re-synchronize data in the event of link loss	No	Yes	RecoverPoint always monitors disk I/O writes. If a WAN link used for replication goes down and then recovers, only the missing writes need to be transmitted as opposed to a full resync.

Next steps

EMC can help to accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of a backup/disaster recovery solution for a virtualized Microsoft SharePoint 2007 environment.

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