

# **Microsoft<sup>®</sup>** | Virtualization

## **Disaster Recovery in a Geographically Dispersed Cross-Site Virtual Environment**

*Enabled by the EMC CLARiiON CX4 Platform, EMC RecoverPoint, and Microsoft Hyper-V on Windows Server 2008*

Applied Technology Blueprint

Published: May 2009

**Introduction:** This white paper describes a blueprint for an EMC<sup>®</sup> disaster recovery solution for Windows<sup>®</sup> virtualized environments. The solution is enabled by the EMC CLARiiON<sup>®</sup> CX4 platform, EMC RecoverPoint, and Hyper-V<sup>™</sup> on Microsoft<sup>®</sup> Windows Server<sup>®</sup> 2008 using Fibre Channel. The blueprint was designed by EMC and Microsoft to demonstrate automated failover capabilities in a virtualized Microsoft environment.

# Copyright

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. **MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.**

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2009 Microsoft Corporation. All rights reserved.

Microsoft, Windows, Windows Server, Windows Server 2008, Hyper-V, SharePoint, Excel, SQL Server, SQL Server 2008, SQL Server 2005, and SQL Server 2000 are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Copyright © 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as 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

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

Part number h6346

# Contents

Executive Summary.....	4
Introduction and Scope.....	4
Key Components.....	4
EMC CLARiiON CX4 Networked Storage.....	4
EMC RecoverPoint.....	6
Windows Server 2008 Hyper-V.....	7
System Center Virtual Machine Manager.....	8
Physical Architecture.....	9
System Integration.....	10
Configuration Considerations.....	11
Hyper-V Considerations.....	11
Disaster Recovery Solution Considerations.....	13
Configuration Testing.....	16
Test Bed Profile.....	16
Test Description.....	18
Results.....	19
Additional Information.....	20

## Executive Summary

This white paper describes a joint EMC and Microsoft applied technology blueprint of an automated disaster recovery solution for a virtualized Windows Server environment. The blueprint is intended to provide customers with an example disaster recovery solution configuration that can be used to consider, evaluate, and select the EMC and Microsoft disaster recovery solution components that best fit their requirements. This white paper is not intended to be an exhaustive study of specific architectures for every environment. To evaluate specific datacenter disaster recovery requirements, please contact an EMC or Microsoft sales representative.

## Introduction and Scope

For Information Technology (IT), mitigating the risks to critical data, systems, and applications in addition to computing infrastructures in the event of system outages or complete disasters presents an ongoing challenge from both a technological and a business perspective. Organizations are continually challenged with finding solutions that not only meet application and data requirements for capacity, performance and availability, but also have proven Return on Investment (ROI) and cost reduction capabilities. The business challenge is to create a cost effective, highly available, and protected virtual server infrastructure that ensures that applications meet business defined service level agreements (SLAs) for availability and disaster recovery preparedness. This white paper will discuss an option that meets these requirements.

The purpose of this blueprint is to demonstrate a disaster recovery solution using the EMC® CLARiiON® CX4 platform and integrate all the components required to run a complete, automated disaster recovery solution. The blueprint provides configuration guidelines for building similar solutions.

The solution described in this white paper is enabled by the EMC CLARiiON CX4 platform, EMC RecoverPoint, and Hyper-V™ on Microsoft® Windows Server® 2008 using Fibre Channel. It was tested jointly by EMC and Microsoft using a mix of fixed LUNs and virtual provisioning that is commonly found in a production environment.

## Key Components

The key components of this solution include:

- EMC CLARiiON CX4
- EMC RecoverPoint
- Windows Server 2008 Hyper-V
- System Center Virtual Machine Manager

### EMC CLARiiON CX4 Networked Storage

EMC CLARiiON CX4 series storage systems, powered by Intel® Xeon® processors, deliver industry-leading innovation in midrange storage. The unique combination of flexible, scalable, hardware design and advanced software capabilities enables CLARiiON CX4 to meet the growing and diverse needs of today's midsize and large enterprises. Through innovative technologies like Flash drives, UltraFlex™ technology, and Virtual

Provisioning™, customers can reduce costs and energy use while optimizing availability and performance.

The dynamic nature of virtualized environments requires a flexible storage solution in order to realize the full benefit of virtualization technology. Virtualization customers who leverage the industry leading EMC CLARiiON CX4 SAN for storage consolidation enjoy the most flexible and highest performing deployments available for virtualized data centers.

EMC CLARiiON CX4 delivers:

Features	Benefits
UltraFlex technology	CLARiiON's unique UltraFlex technology provides flexible consolidation for dynamic virtualized environments. Integrated Fibre Channel and iSCSI connectivity on all systems allows customers to easily leverage their existing LAN infrastructure through iSCSI and deploy Fibre Channel for applications with high throughput requirements. The modular design also provides easy, online expansion to accommodate additional connectivity requirements and emerging I/O technologies.
Five 9's availability	Fully redundant hardware architecture and unique software capabilities deliver 99.999 percent system level availability.
Enterprise Flash Drive technology	CLARiiON Flash Drive technology breaks the performance barrier of traditional disk technologies, providing 30 times the IOPS of a 15K FC drive and consistently delivering less than 1 millisecond response times.
Virtual Provisioning	Virtual Provisioning, also known as thin provisioning, allows users to easily provision storage to virtual servers, increase capacity utilization rates, and defer purchase of additional capacity. It also simplifies storage management by enabling just-in-time capacity allocation through monitoring, alerts, and reporting, and enables more efficient capacity planning.
Virtual LUN	Virtual LUN enables users to move virtual servers between drive types and RAID groups to meet performance and availability requirements without disruption as they change over time.

**Table 1. EMC CLARiiON CX4 features**

## EMC RecoverPoint

EMC RecoverPoint is an enterprise-scale solution designed to protect application data on heterogeneous SAN-attached servers and storage arrays. 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.

With RecoverPoint, administrators are able to implement a single, unified solution to protect and/or replicate data across heterogeneous storage with minimal to no data loss as compared to traditional host and array snapshots or disk-to-tape backup products. Administrators are able to simplify management and reduce costs, recover data at a local or remote site at any point in time, and ensure continuous replication to a remote site without impacting performance.

EMC RecoverPoint delivers:

Features	Benefits
Continuous data protection	Protection and recovery to any point in time.
Continuous remote replication	Bi-directional, heterogeneous block-level replication across WAN or FC networks.
Concurrent local and remote data protection	Protect and replicate data in local and remote site combinations for operational and disaster recovery.
Policy-based management	Leverage service-level policies that optimize RPO and RTO.
IP Bandwidth optimization	IP bandwidth management with prioritization, data compression, and integrated FC to IP conversion.
Block-level journaling of data changes	Enables full read/write access to a point-in-time image and wizard-based production recovery.
Data protection	Use RecoverPoint to protect against data corruption with flexible protection and recovery options.
Virtual infrastructure integration	Simplify and automate disaster recovery in Hyper-V environments with Windows failover cluster integration that works across heterogeneous storage arrays.

**Table 2. RecoverPoint features**

## Windows Server 2008 Hyper-V

Hyper-V is the hypervisor-based virtualization technology from Microsoft that is integrated into all Windows Server 2008 x64 Editions operating systems. As a virtualization solution, Hyper-V enables users to take maximum advantage of the server hardware by providing the capability to run multiple operating systems (on virtual machines) on a single physical server.

The availability of Hyper-V as a role in a mainstream Windows operating system provides several key advantages:

Features	Benefits
Built in technology	Hyper-V allows enterprises to easily leverage the benefits of virtualization without adopting a new technology.
Broad device driver support	The new 64-bit micro-kernelized hypervisor architecture leverages the broad device driver support in Windows Server 2008 parent partition to extend support to a broad array of servers, storage, and devices.
SMP support	Hyper-V supports Symmetric Multiprocessors (SMP) in virtual machines.
High availability	Windows Server 2008 clustering provides high availability to virtual machines to minimize unplanned downtime.
Easy VM migration	Quick/Live Migration capability to support business continuity during planned and unplanned downtime.
VSS support	Robust host-based backup of virtual machines by leveraging the existing Windows VSS-based infrastructure.
Easy extensibility	Easy extensibility using the standards-based Windows Management Instrumentation (WMI) interfaces and APIs.
Simplified integrated management	With its tight integration into Microsoft System Center family of products, customers have end-to-end physical and virtual infrastructure management capability for Hyper-V environments.

**Table 3. Hyper-V features**

## System Center Virtual Machine Manager

Microsoft System Center Virtual Machine Manager 2008 (SCVMM 2008) is enterprise-class management software that enables administrators to easily and effectively manage both the physical and virtual environments from a single management console and thus avoiding the complexity of using multiple consoles typically associated with managing an IT infrastructure. The key capabilities of SCVMM 2008 include:

Features	Benefits
Enterprise-class management suite	Manages both Hyper-V and VMware ESX virtualization environments.
Intelligent VM placement	Support for intelligent placement of virtual machines.
System Center Operations Manager 2007 integration	Integration with System Center Operations Manager 2007 to provide proactive management of both virtual and physical environments through a single console by leveraging PRO.
Native P2V/V2V migration	Native capability for Physical-to-Virtual migration and Virtual-to-Virtual migrations.
Failover integration	Integration with Failover Clustering to support high availability and Quick/Live Migration of virtual machines.
Automation	Easy automation capabilities leveraging Windows PowerShell.

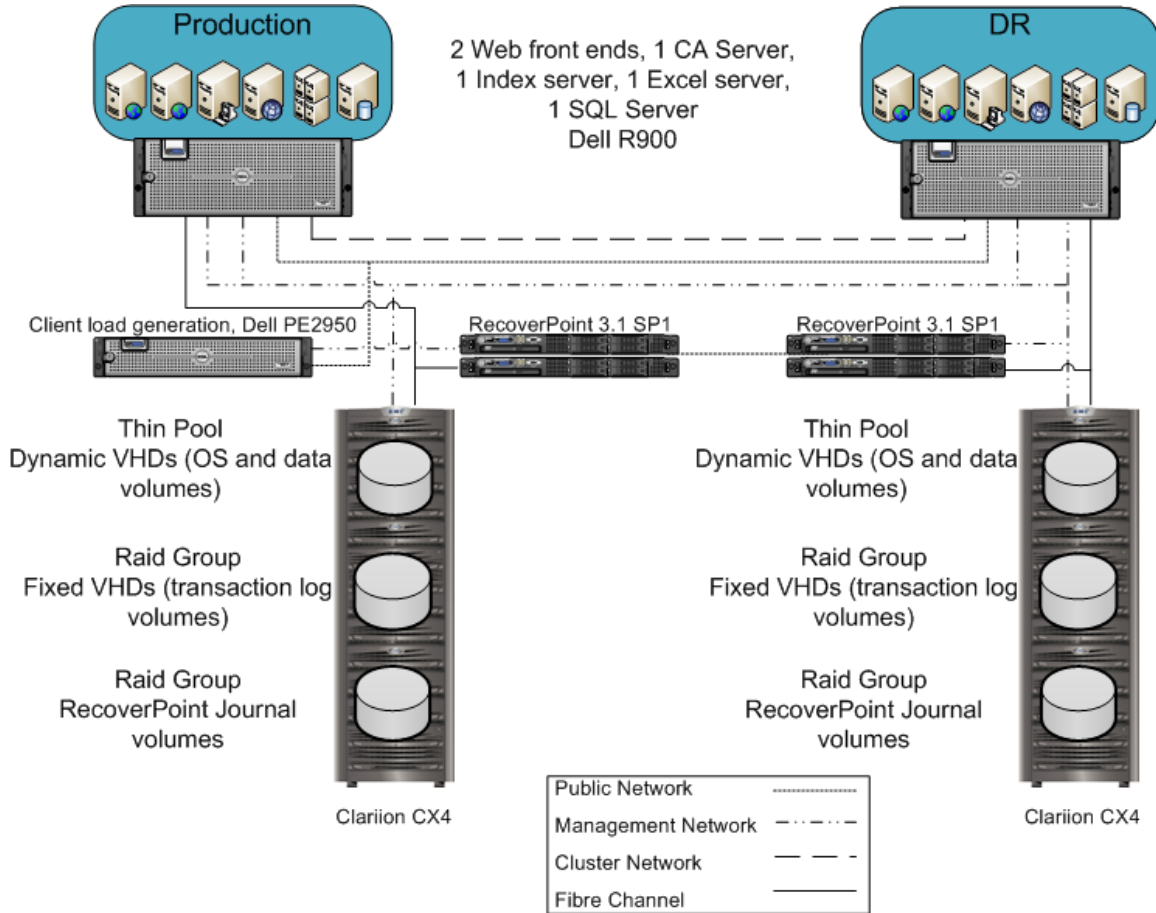
**Table 4. System Center Virtual Machine Manager features**

SCVMM 2008 can be configured in multiple configurations depending on the implementation requirements. A basic configuration will have SCVMM 2008 installed and running on a standalone server with local disks on the server as storage. Attaching a storage enclosure to the standalone server hosting SCVMM is recommended if the deployment requires a relatively large library server. The library server is a capability built into SCVMM for storing VHD templates, inactive VM files, ISO images, and so on.



# Physical Architecture

The following illustration depicts the overall physical architecture of the EMC disaster recovery solution. For validation purposes, it was configured as a virtualized SharePoint® Server farm, but the blueprint can be applied to any disaster recovery configuration.



**Figure 1. Solution architecture**

The disaster recovery solution is configured as a two node cluster with a file share witness. All of the SharePoint roles are running in separate child partitions on a single node of the cluster. Each node is connected to a separate array. The arrays were configured with both thin and fully provisioned LUNs. RecoverPoint is being used to replicate the data between the two arrays. RecoverPoint/CE is the integration point for Windows failover clustering and RecoverPoint. RecoverPoint/CE creates a new CE resource in the child partition cluster groups. This resource is what communicates with the RecoverPoint appliances to initiate failover at the storage level. The cluster disks are dependent on this resource and will not come online until the CE resource comes online. The CE resource comes online when the storage failover is complete.

A document loader is used on a third server for load generation. The document load is utilizing the public network connection to a network load balanced connection on the SharePoint Web front end.

The cluster network is dedicated to the cluster heartbeat. The following table describes the various components within the configuration.

Server Configuration	
Servers	Dell PowerEdge R900
CPU	4 Sockets (4 cores) 1.6 GHz
Memory	64GB
Network	4 ports
HBAs	Brocade 815 (2)
Storage Configuration	
Storage Fabric	Brocade DCX
	Brocade 7500
Storage Device	EMC CLARiiON CX4-960
	EMC CLARiiON CX4-480
Distance Replication	EMC RecoverPoint 3.1 SP1
Virtualization Components	
Operating System	Microsoft Windows Server 2008 Enterprise x64
Virtualization Software	Microsoft Windows Server 2008 Hyper-V
Management Software	
Virtualization Management	SCVMM 2008 Enterprise Edition
Physical Server Management	SCOM 2007 with PRO
Storage Management	EMC Navisphere® Management Suite
Storage Replication Management	EMC RecoverPoint Management Application
Stretch Cluster Application	EMC RecoverPoint/Cluster Enabler

Table 5. Solution components

## System Integration

After the infrastructure is set up, the system must be integrated to enable the disaster recovery solution. The system integration steps include:

- **RecoverPoint Integration:** RecoverPoint must be integrated with the CLARiiON CX4 platform through the CLARiiON-based splitter and with the Microsoft Failover Clusters with the RecoverPoint/Cluster Enabler. This requires integrating with the following:
  - **CLARiiON:** RecoverPoint replicates data by intercepting the application writes through the use of intelligent write splitting modules which can reside on the host or on a CLARiiON CX3 or CX4 array. The CLARiiON array-based splitter runs in each storage processor of a CLARiiON CX3 and CX4 array and will mirror all writes to a CLARiiON SAN or iSCSI volume, sending one copy to the original target and the other copy to the RecoverPoint appliance. The CLARiiON CX4 platform discussed in this paper is included in the integrated RecoverPoint write-splitting technology. When tested:
    - The CLARiiON-based splitter ran on each CLARiiON storage processor.

- RecoverPoint ran on the RecoverPoint appliances and used the CLARiiON-based splitters.
- CLARiiON and RecoverPoint were managed through their respective management applications.
- Integration Guidance: EMC has a RecoverPoint Installation guide (P/N 300-007-559) that details the RecoverPoint integration with CLARiiON. For EMC clients, the guide can be found at the secured EMC Powerlink® site at: <http://powerlink.emc.com/>.
- **Microsoft Failover Clusters:** RecoverPoint/Cluster Enabler integrates RecoverPoint continuous remote replication with Microsoft Failover Clusters on Windows Server 2008 to facilitate the failover processes for geographically dispersed cluster nodes with RecoverPoint replicated storage. EMC provides an integration guide for its clients:
  - Integration Guidance: EMC has a RecoverPoint/Cluster Enabler product guide (P/N 300-007-681) that details the integration process. For EMC clients, the guide can be found at the secured EMC Powerlink site at: <http://powerlink.emc.com/>.
  - **System Center Virtual Machine Manager 2008 integration:** For the System Center Virtual Machine Manager to work with this disaster recovery solution, a Microsoft hotfix is required to be installed on the SCVMM server. The required SCVMM hotfix can be downloaded from Microsoft at: <http://support.microsoft.com/kb/961983>.

## Configuration Considerations

When structuring a disaster recovery solution within a virtualized environment, there are many configuration options that should be considered to optimize the solution for a particular enterprise environment.

### Hyper-V Considerations

There are many considerations when configuring and tuning Hyper-V that can affect the performance of the environment and the disaster recovery solution including:

- **Hardware Considerations for Hyper-V Environments:** Each of the four following subsystems should be optimized to provide the best overall quality of service in the datacenter:
  - **CPU Subsystem:** The number of virtual processors to be assigned to a specific virtual machine depends on the requirements of the workload that runs on the guest operating system. Windows Server 2008 Hyper-V allows up to four virtual processors to be assigned to a specific virtual machine (VM). Support varies for different guest operating systems.

In Hyper-V, virtual processors in VMs do not have a one-to-one mapping with the logical processors on the server. The Hyper-V hypervisor handles the scheduling of virtual processors on any available logical processors. The general guidance is that when the virtual workload is highly processor intensive, using multiple virtual processors in a VM might be effective since more physical processors will be used. Since the use of multiple virtual processors also adds additional

overhead, careful planning is required in determining the virtual processor allocation to VMs.

- **Memory Subsystem:** Virtualization workloads are memory intensive. This is especially true when one of the goals for the virtualization deployment is consolidation. The memory need for a specific Hyper-V host server is driven by the amount of memory needed by each virtual machine, the amount of memory needed by the hypervisor to manage those VMs, and the memory needed to allow additional virtual machines to be migrated onto the system.
- **Hypervisor and Parent Partition:** The general guidance is to reserve 1GB of memory.
- **Virtual Machines:** The amount of memory required is the sum of all memory requirements for the running virtual machines on the server. For each virtual machine, the following formula should be utilized to determine the memory requirement:
  - If the VM is allocated  $\leq$  1GB memory: Amount of memory allocated to the VM + 32MB.
  - If the VM is allocated  $>$  1GB memory: Amount of memory allocated to the VM + 32MB + 8MB per GB allocated to the VM for each GB over 1GB.
  - For example: A VM allocated with 3GB of memory would require a reservation of 3GB + 48MB.
- **Migration:** If additional virtual machines happen to migrate (either planned or unplanned) to the server, memory, per the sizing guidance above, should be reserved for those virtual machines. Failure to do so can result in virtual machines being unable to start on the server due to the lack of available resources.
- **Disk Subsystem:** Hyper-V is designed to provide superior I/O capabilities through its shared device design. For a balanced system design, the storage subsystem needs to provide enough storage capacity while delivering QoS to the datacenter through acceptable IO request response times. In addition, the number of hosts connecting to the disk subsystem and the fabric type need to be factored in to pick the correct storage enclosure.
  - The capacity portion is a simple calculation based on RAID level choice and capacity needs for virtual machines and applications.
  - Performance sizing the disk subsystem is a critical aspect to ensure that solution has sufficient disk spindles and the bandwidth to handle the IO needs.
  - Another design criterion is the need for support of highly available virtual machines and virtual machine migration.
- **Network Subsystem:** The network subsystem needs to deliver QoS on I/O requests, connect the datacenter, and enable key hypervisor features. The decision points involve choosing enough NIC ports to handle the data traffic, management traffic, and cluster communication to meet current needs and enable scalability as the datacenter grows. In addition, if iSCSI storage arrays are utilized in the configuration, then iSCSI traffic drives networking

requirements as well. The guidance below is general best practices for each server in the datacenter.

- **Parent Partition / Virtualization Management:** A minimum of one port should be allocated. NOTE: This adapter can also be used as a fallback for the public cluster communication network if this Hyper-V host is part of a highly available cluster.
- **Virtual Machine Network:** A minimum of one port should be allocated for the virtual machine network. Additional adapters might be required based on the number of virtual machines and workloads running in the virtual machines.
- **Storage Network:** If iSCSI storage arrays will be utilized, a minimum of two ports should be allocated to ensure high availability.
- **Cluster Private Network:** One adapter should be allocated for the private cluster network if the Hyper-V host is part of a highly available cluster.
- **Performance Tuning:** Microsoft produced a performance tuning guide for Windows Server 2008 that includes a significant section on Hyper-V. The guide can be found at:  
[http://www.microsoft.com/whdc/system/sysperf/Perf\\_tun\\_srv.mspx](http://www.microsoft.com/whdc/system/sysperf/Perf_tun_srv.mspx).

### Disaster Recovery Solution Considerations

When configuring this disaster recovery solution, there are a number of configuration considerations that are important to the success of the solution as well as the optimization of its performance including:

- **Replication Bandwidth:** RecoverPoint reduces the amount of network bandwidth required while replicating data asynchronously over a WAN IP network. This bandwidth reduction is provided through a superior optimization analysis and a compression algorithm. EMC provides analysis of the I/O load during production and based on empirical studies, provides guidance on required network bandwidth, configuration, and the number of RecoverPoint appliances needed to continue to meet the business SLA and Recovery Point Objectives (RPO).

The data reduction algorithms consume CPU and memory resources if the RecoverPoint appliance is overburdened which can limit the incoming writes that can be processed per appliance. When optimized, a single RecoverPoint cluster solution can accommodate between two to eight nodes per site. The tradeoff between network savings due to data reduction and number of appliances needs to be carefully considered during the planning stage to ensure proper operation.

- **Planned/Unplanned Failover:** The distribution and replication algorithms of RecoverPoint ensures write-order fidelity on the replica LUNS. That is, every image that can be presented on the target side is crash-consistent and can be used for recovery purposes. In case of a planned graceful failover, the system administrator might choose to additionally create a bookmark within the RecoverPoint consistency group after the production system has shut down. This bookmark-based image can then be accessed by the disaster recovery host to ensure zero-data loss and application consistency. During an unplanned failover, typically the latest available image is presented to the disaster recovery host for recovery. A significant advantage with the RecoverPoint solution is that if the latest image is unusable or

corrupt, it is possible to roll back to a previous point in time that satisfies the RPO.

The RecoverPoint system also provides two wizard-based workflows:

- **Failover:** This workflow is used for promoting a disaster recovery target system to be produced and immediately start replication in the reverse direction. The same procedure can also be used for a failback to the original production side after any potential issues have been resolved.
- **Recover Production:** This workflow is used to roll the production system to a previous point in time. This workflow is typically used to recover from an application corruption.
- **Consistency Groups:** It is possible to use separate consistency groups for each application and the operating system SAN boot volumes based on the recovery requirement. A consistency group has one or more replication sets corresponding to each LUN that is part of that application. A consistency group is the smallest unit of recovery in RecoverPoint. In this solution, a consistency group was carved for each function such as Excel<sup>®</sup> Services, Index Services and Microsoft SQL Server<sup>®</sup>. Additionally, a consistency group was created for each of the Web Front End (WFE) virtual machines. This enables the ability to recover each of the applications independently.

RecoverPoint also includes a Group Set feature that can be used to define a set of consistency groups on which to automatically apply a parallel bookmark according to a specified frequency. The automatic parallel bookmark applied to each of the groups in the group set will include the name of the group set, followed by an ID number. The same ID is used across all the groups. Automatic parallel bookmarks are successfully applied for all of the consistency groups in the group set only when the source site is the same and transfer is active. An individual parallel bookmark can also be initiated from the command line or the GUI prior to a planned failover using the `apply_parallel_bookmark` feature by including the consistency groups as parameters. When using the group set feature or the `apply_parallel_bookmark` feature, administrators can pause on the image that corresponds to the parallel bookmark in each of the consistency groups.

The advantage of the group set is the ability to failover either part of the solution or the entire solution to the DR site. For instance, it might be necessary to failover only the Exchange environment on the DR site for a test. The other extreme is a site-wide disaster requiring the failover of all the consistency groups Fibre Channel versus iSCSI: Either option is acceptable depending on the performance requirements.

- **LUN Configuration:** When using RecoverPoint/CE and SCVMM, there are certain requirements when it comes to drive letters and mount points.
  - SCVMM requires that there is only one VM per LUN.
  - The setup of RecoverPoint/CE with SCVMM requires that all LUNs associated with a virtual machine are either drive letters or mount points associated with the LUN that the virtual machine operating system is installed on.
- **Active Directory:** Ensure that the disaster recovery site is part of the Active Directory<sup>®</sup> infrastructure. Any applications that utilize integrated security require



access to the Domain Controller. An additional site in Active Directory can be configured to enable the replication of Active Directory data. This also prevents user authentications in the production site from using the Domain Controller in the disaster recovery site.

- **Automation:** There are many things to consider when determining whether to automate failover or to do this manually:
  - **Availability of Key Personnel to Manually Failover:** While initial configuration complexity and running complexity can be a demand on resources, the execution of a failover plan is where automation and simplicity must be the focus. When a disaster is declared, key personnel might not be available in addition to the loss of servers, storage, networks, buildings, and so forth. If the complexity of the failover solution is such that skilled personnel with an intimate knowledge of all systems involved is required to restore, recover, and validate application and database services, the solution has a high probability of failure.
  - **Recovery Time Objective:** A server cluster provides high availability by making application software and data available on several servers linked together in a cluster configuration. If one server stops functioning, the failover mechanism automatically shifts the workload of the failed server to another server in the cluster. The failover process is designed to ensure the continuous availability of critical applications and data.

RecoverPoint/CE ensures the availability of user data during a failover. It automates the RecoverPoint processes of moving cluster resources between storage systems. By integrating with Microsoft Failover Clusters, RecoverPoint/CE enables the cluster to automatically manage resource and application failover which greatly improves the recovery time objective (RTO).

- **User Session State Consistency:** While clusters can be designed to handle failure, they are not fault tolerant with regard to user data. The cluster by itself does not guard against loss of a user's work. Typically, the application software handles the recovery of lost work. Therefore, the application software must be designed to recover the user's work, or it must be designed in such a way that the user session state can be maintained in the event of a failure.
- **Failover Complexity Reduction:** Regardless of failing over automatically or manually, the failover process must be made to be as simple as possible to ensure its success. For example, when using RecoverPoint without RecoverPoint/CE, users typically configure separate clusters at the primary and secondary sites. In this case, moving volumes between the clusters involves scripting the un-mounting of LUNs from the cluster at the primary location, making the RecoverPoint replica available to the secondary cluster, and then redirecting application clients to the new cluster. Without a certain degree of automation, this process might be challenging during an actual total site disaster requiring cross-site geographical failover.

## Configuration Testing

The disaster recovery solution uses the RecoverPoint/Cluster Enabler to enable remote site failover for Hyper-V in a virtualized Windows Server environment. To validate this disaster recovery solution, a SharePoint Server farm was virtualized in the lab. A document loader was used to produce a load on the SharePoint Server farm which was then failed over to the disaster recovery server to validate the solution.

### Test Bed Profile

To configure RecoverPoint/CE and to use SCVMM in the solution, every disk associated to a virtual machine operating system was required to be given a drive letter as previously mentioned in the disaster recovery considerations section of this white paper. In addition, all of the disks associated with each virtual machine were configured as mount points under the appropriate drive letter to conserve drive letters. To further reduce the number of drive letters needed, all of the virtual machines could have been placed under one drive letter. However, because of the way RecoverPoint/CE builds the LUNs dependencies, SCVMM would not have been able to manage the environment.

A total of six virtual machines are running in the SharePoint farm. This includes two virtual machines hosting the Web Front Ends (WFEs) that are also running the Query Service. There is one virtual machine hosting the SQL Server. Another virtual machine is running the Index Service. The Index Service is configured to crawl the data instead of using the WFEs for crawling. The other two virtual machines host Excel Services and the Central Administration site.

The table below details the virtual machine configuration.

Child Partition Name	vCPU	Memory (GB)	Network Adapters	Application/Service
CA Site	1	2	1	Central Administration Site
Excel Services	1	2	1	Excel Services
Index Services	4	6	2	Indexing Services
SQL Server	4	32	1	SQL Server 2008
WFE 1	4	4	2	Web Front End and Query Services
WFE 2	4	4	2	Web Front End and Query Services

**Table 6. Virtual machine configuration**

All of the virtual machines are configured with virtual hard drives (VHD) for backend storage and leverage the backup capabilities of Windows Server 2008 and VSS in conjunction with Hyper-V.

The LUNs were configured with a mixture of thin and fully provisioned LUNs to provide a balance of performance and economy as needed. For example, the database transaction log LUNs are on fully provisioned LUNs and are using fixed VHDs to provide better performance since transaction logs are very sensitive to performance. All of the remaining LUNs are thin provisioned and utilize dynamic VHDs to reduce overall storage requirements.

The thin LUNs for the content databases are sized to accommodate 100 GB of data, which is the recommended maximum for SharePoint content databases. Another 10 to



20 percent is allocated for NTFS overhead. A total of five content databases are used in the SharePoint farm.

The table below details the SharePoint Servers storage pool configuration.

Storage Group	Disks	Disk Type	Raid Type
Thin Pool 0	15	300GB 15k	RAID 5
RAID Group 100	4	300GB 15k	RAID 10
RAID Group 101	4	300GB 15k	RAID 10

**Table 7. SharePoint Server storage pool configuration**

The table below details the LUN storage configuration.

LUN Name	Allocated Size (GB)	Used Size (GB)	Storage Group	VHD Type
CA Site OS	50	19	Thin Pool 0	Dynamic
Excel Services OS	50	21	Thin Pool 0	Dynamic
Index OS	65	40	Thin Pool 0	Dynamic
Index Storage	115	7	Thin Pool 0	Dynamic
SQL Config Data/Log	35	4	Thin Pool 0	Dynamic
SQL Content Data 1	130	58	Thin Pool 0	Dynamic
SQL Content Data 2	130	58	Thin Pool 0	Dynamic
SQL Content Data 3	130	58	Thin Pool 0	Dynamic
SQL Content Data 4	130	58	Thin Pool 0	Dynamic
SQL Content Data 5	130	58	Thin Pool 0	Dynamic
SQL Content Log 1	12	12	RAID Group 100	Fixed
SQL Content Log 2	12	12	RAID Group 100	Fixed
SQL Content Log 3	12	12	RAID Group 100	Fixed
SQL Content Log 4	12	12	RAID Group 100	Fixed
SQL Content Log 5	12	12	RAID Group 100	Fixed
SQL Search Data 1	85	9	Thin Pool 0	Dynamic
SQL Search Data 2	85	8	Thin Pool 0	Dynamic
SQL Search Log 1	15	15	RAID Group 100	Fixed
SQL Server OS	90	40	Thin Pool 0	Dynamic
SQL TempDB Data 1	20	4	Thin Pool 0	Dynamic
SQL TempDB Data 2	20	4	Thin Pool 0	Dynamic
SQL TempDB Data	20	4	Thin Pool 0	Dynamic

3				
SQL TempDB Data	20	4	Thin Pool 0	Dynamic
SQL TempDB Log	15	15	RAID Group 100	Fixed
WFE 1 OS	50	30	RAID Group 100	Fixed
WFE 1 Query	65	7	RAID Group 100	Fixed
WFE 2 OS	50	30	RAID Group 100	Fixed
WFE 2 Query	65	7	RAID Group 100	Fixed
CA Site Journal	10	10	RAID Group 101	
Excel Services Journal	10	10	RAID Group 101	
Index Journal	15	15	RAID Group 101	
SQL Server Journal	200	200	RAID Group 101	
WFE 1 Journal	15	15	RAID Group 101	
WFE 2 Journal	15	15	RAID Group 101	
Totals	1890	883		

Table 8. LUN storage configuration

## Test Description

The test bed consists of a single SharePoint portal that is network load balanced between two virtual Web Servers on the front end. The SharePoint farm is tied to a single virtual SQL Server. All the virtual machines are located on the same physical server. The disaster recovery server was built on a second physical server with identical specifications.

A third physical server was used to host a document loader. The document loader generated a load on the SharePoint Server farm of 16 MB per second which is equivalent to 25 documents per second. The uploaded test documents were relatively small in size, the largest file being 471 KB in size.

Three scenarios were tested. The first two scenarios were planned failovers with and without a load. The third scenario was an unplanned failover with a load. The planned failovers were done by moving the cluster groups in the Cluster Management console. The unplanned failover was simulated by pulling the plug on the server. The failover time was then measured to determine the average time required for normal operations to begin on the disaster recovery server.

## Results

In all scenarios, the disaster recovery solution worked as expected. After the failover was simulated, the disaster recovery site took over operations from the production site. Because the various servers required different amounts of time to failover, the longest failover time was used for the total average failover time and is bolded for each of the tests in the table below.

The table below show the various failover times for each of the virtual machines.

Child Partition Name	Planned Failover with no load (mm:ss)	Planned Failover with load (mm:ss)	Unplanned Failover with load (mm:ss)
CA Site	1:15	1:15	1:10
Excel Services	1:15	1:15	1:10
Index Services	1:25	1:50	1:40
SQL Server	<b>2:15</b>	<b>4:20</b>	<b>1:40</b>
WFE 1	1:20	1:45	1:40
WFE 2	1:20	1:45	1:40

**Table 9. Failover results**

Most of the average failover times for each server were less than two minutes. The SQL Server consistently took the longest time to fail over. For example, the planned failover with and without load for SQL Server took the longest time of all the tests. This is because failover clustering saves the state of the virtual machines during failover which requires all of the memory contents to be destaged to disk. The SQL Server virtual machine has 32GB of memory allocated to it which is more than five times the next highest memory allocation. This required significantly more time for SQL Server to destage its memory contents to disk. The unplanned failover for SQL Server did not experience this delay since the virtual machine was not shut down gracefully.

What is important to take from the validation of this disaster recovery solution is that it works automatically and without ongoing management. Nothing else is required unless new virtual machines are added to the configuration. In this case, the RecoverPoint/CE configuration process must be run again to map the new configuration. This automation replaces the chaos and missteps that are likely to occur during a real disaster scenario.

## Additional Information

EMC Information Infrastructure Solutions for Microsoft virtualization helps organizations to realize the cost, service level, and agility benefits that the Microsoft virtualization technologies deliver. EMC Information Infrastructure Solutions for Microsoft Virtualization bring together EMC's leading infrastructure technology and services expertise with Microsoft Virtualization to help deliver dynamic IT infrastructures.

Additionally, EMC can help accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of a disaster recovery solution for a virtualized Microsoft environment. To learn more about this and other solutions, contact an EMC representative or see the following:

- EMC Information Infrastructure Solutions for Microsoft Virtualization <http://www.emc.com/solutions/business-need/virtualizing-information-infrastructure/information-infrastructure-solutions-for-microsoft-virtualization.htm>
- Powerlink page for EMC Customers (Password Required): [http://powerlink.emc.com/km/appmanager/km/secureDesktop?\\_nfpb=true&internalId=0b014066800d792a](http://powerlink.emc.com/km/appmanager/km/secureDesktop?_nfpb=true&internalId=0b014066800d792a)

### About EMC

EMC Corporation (NYSE: EMC) is the world's leading developer and provider of information infrastructure technology and solutions that enable organizations of all sizes to transform the way they compete and create value from their information.

Information about EMC's products and services can be found at <http://www.emc.com/>.

### For more information, see:

- EMC CLARiiON CX4 family product information: <http://www.emc.com/products/series/cx4-series.htm>
- EMC RecoverPoint product information: <http://www.emc.com/products/family/recoverpoint-family.htm>
- Microsoft Windows Server 2008 Hyper-V product information: <http://www.microsoft.com/windowsserver2008/en/us/hyperv-main.aspx>
- Microsoft System Center Virtual Machine Manager product information: <http://www.microsoft.com/systemcenter/virtualmachinemanager/en/us/default.aspx>