

EMC PROTECTION FOR MICROSOFT EXCHANGE SERVER 2010

EMC VNX5700, EMC RecoverPoint, EMC Replication Manager,
EMC Data Protection Advisor, VMware vSphere, and
VMware vCenter Site Recovery Manager

- Fully automated recovery for Microsoft Exchange Server 2010
- Unlimited snapshots to protect from logical database corruption
- In-depth monitoring and reporting

EMC Solutions Group

Abstract

The EMC® Protection for Microsoft Exchange Server 2010 solution is an Exchange business continuity solution designed for enterprises with two or more active data centers at different geographic locations. The solution offers high availability at every location and near-instantaneous recovery from a disaster at any location.

The solution leverages EMC VNX5700 storage with the VNX™ Total Protection Pack (EMC RecoverPoint, EMC Replication Manager, and EMC Data Protection Advisor) and VMware® vSphere™ 5 with vCenter™ Site Recovery Manager (SRM) 5.

December 2011



Copyright © 2011 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.

VMware, ESX, 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 H8891

Table of contents

Executive summary	6
Business case.....	6
Solution overview	6
Introduction.....	7
Purpose	7
Scope	7
Audience	7
Technology overview	8
Overview.....	8
EMC VNX5700 storage system	8
EMC RecoverPoint.....	9
RecoverPoint appliances (RPA)	9
RecoverPoint splitter	10
RecoverPoint journals.....	10
RecoverPoint consistency groups	10
EMC Replication Manager	11
EMC Data Protection Advisor.....	12
EMC PowerPath/VE	13
VMware vSphere	13
VMware vCenter Site Recovery Manager.....	13
EMC RecoverPoint Storage Replication Adapter for VMware vCenter SRM.....	14
Overall solution architecture and design	15
Clarification of terminology	15
Design attributes	15
Virtualization	16
Active/active site configuration.....	16
Data protection	17
Exchange Server 2010 Design.....	17
Exchange storage design	19
Building block design approach.....	20
Mailbox server virtual machine building block details	20
About storage pools.....	21
VNX storage pool design	22
vSphere server design.....	23
Exchange local protection	25

RecoverPoint design	26
RecoverPoint consistency group design	26
RecoverPoint storage requirements.....	27
RecoverPoint journal design.....	27
Considerations when sizing journal volumes	28
Sizing journal volumes.....	28
Consistency group policy configuration for management by SRM	30
RecoverPoint group sets.....	30
Group Sets configuration	30
Replication Manager configuration and integration with RecoverPoint.....	32
Replication Manager process overview	32
VSS API and KVSS utility	32
Exchange tasks that can be automated by RM.....	32
How this solution uses RM	33
RM application set configuration.....	33
Special RM environment variable required	34
About replication options.....	34
Replication option configuration	34
vCenter Site Recovery Manager configuration	36
Prerequisites.....	36
SRM requirements for vSphere.....	36
SRM configuration steps	36
Step 1 – Install and configure SRM with the RecoverPoint SRA adapter	37
Step 2 – Configure the connection between the protected and recovery sites	37
Step 3 – Configure RecoverPoint array managers	37
Step 4 – Configure resource inventory mappings	38
Step 5 – Configure protection groups.....	39
Step 6 – Create the recovery plan.....	40
Step 7 – Customize virtual machine recovery options, including IP customization	41
Exchange configuration to support data center failover	42
Analysis and reporting with EMC Data Protection Advisor (DPA)	43
DPA integration with RecoverPoint	43
DPA/RecoverPoint sample reports.....	44
DPA integration with VMware vSphere	48
Registering the plug-in	48
Viewing the plug-in reports.....	48
DPA/vSphere sample report.....	48
DPA integration with VNX/CLARiiON	49

DPA/VNX sample reports	49
DPA integration with Microsoft Exchange Server.....	50
Solution validation and test results.....	51
RecoverPoint journal usage under Exchange load	51
About Microsoft LoadGen.....	51
RecoverPoint journal space usage during testing	51
SRM recovery plan testing	53
Recovery plan testing.....	53
Exchange failover with SRM	55
Changing the DAG IP address.....	56
Validating client access	57
Exchange failback with SRM	58
Resynchronization options.....	58
Failback steps.....	59
Exchange database recovery with Replication Manager	60
Restore options.....	60
Consistency group restore options	61
Conclusion	62
Summary	62
Findings.....	62
References	63
White papers	63
Product documentation.....	63
Other documentation.....	63

Executive summary

Business case

In the enterprise, email, calendaring, and unified messaging have become mission-critical applications, with the leading platform being Microsoft Exchange. IT resources face the conflicting challenges of reducing administration costs while delivering greater protection for both the application and the user data. Because any unplanned downtime can have a dramatic impact on the continuity of business, IT customers are demanding improved levels of service, including continuous application availability. What's more, enterprises with data centers at multiple geographic locations want the ability to weather a natural or man-made disaster at any data center site without any disruption to business.

Solution overview

The EMC® Protection for Microsoft Exchange Server 2010 solution is an Exchange business continuity solution designed for enterprises that require application protection both locally, within a data center, and also across multiple data centers within a metropolitan area or across the globe. The solution offers high availability at every location and automated recovery from a disaster at any location.

The solution leverages:

- EMC VNX5700 storage with the VNX™ Total Protection Pack, which includes:
 - EMC RecoverPoint/SE
 - EMC Replication Manager
 - EMC Data Protection Advisor for replication analysis
- VMware® vSphere™ 5 with vCenter™ Site Recovery Manager (SRM) 5

Introduction

Purpose

The purpose of this white paper is to demonstrate how EMC and VMware technologies can be leveraged to provide high availability and disaster recovery protection for Microsoft Exchange Server 2010 enterprise deployments at multiple, active, geographically distributed data centers.

Scope

The scope of this paper corresponds to the scope of the solution validation (build, test, and document) activities performed by EMC engineers in an EMC laboratory.

What was built and tested is described and, where possible, recommendations and guidelines are provided for professionals to design an identical or similar solution for a customer.

The concepts, instructions, procedures, recommendations, and guidelines presented in this document are by no means exhaustive.

Audience

The target audience for this white paper is business executives, IT directors, and infrastructure administrators who are responsible for their company's Exchange landscape.

The target audience also includes professional services groups, system integration partners, and other EMC teams tasked with deploying Exchange systems in a customer environment.

A high-level understanding of Exchange solutions and Exchange landscapes is beneficial. Familiarity with virtualization concepts is also beneficial.

Technology overview

Overview

This section provides an overview of the primary technologies used in this solution. The tight integration of these products and technologies make possible all of the benefits of the Exchange business continuity solution described in this paper.

- EMC VNX5700 storage system
- EMC RecoverPoint
- EMC Replication Manager
- EMC Data Protection Advisor
- EMC PowerPath®/VE
- VMware vSphere
- VMware vCenter Site Recovery Manager (vCenter SRM)

EMC VNX5700 storage system

The EMC VNX5700 storage system is a member of the VNX series next-generation storage platform, powered by Intel processors. The VNX series is powered by Intel® Xeon® Processors, for intelligent storage that automatically and efficiently scales in performance, while ensuring data integrity and security. The VNX series is designed to deliver maximum performance and scalability for mid-tier enterprises, enabling them to grow, share, and cost-effectively manage multiprotocol file and block systems. VNX arrays incorporate the RecoverPoint splitter, which supports unified file and block replication for local data protection and disaster recovery.

EMC Unisphere™ is the central management platform for the EMC VNX series, providing a single combined view of file and block systems, with all features and functions available through a common interface. Unisphere is optimized for virtual applications and provides industry-leading VMware integration, automatically discovering virtual machines and ESXi® servers and providing end-to-end, virtual-to-physical mapping.

For additional VNX specifications, refer to the *EMC VNX Series Unified Storage Systems Specification Sheet* (<http://www.emc.com>).

EMC RecoverPoint

The EMC RecoverPoint family provides a cost-effective solution for protection of critical production data at both local and remote sites, and provides customers with a centralized management interface that allows the recovery of data to practically any point in time.

EMC RecoverPoint is an enterprise-class data protection, replication, and disaster recovery 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 (CDP) technology with unique bandwidth reduction technology, no-data-loss replica creation, and updating to protect data both locally and remotely. RecoverPoint provides the following replication options for both physical and VMware virtualized environments:

- Continuous Data Protection (CDP)—CDP continuously captures and stores data modifications locally, enabling local recovery from any point in time, with no data loss. Both synchronous and asynchronous replications are supported.
- Continuous Remote Replication (CRR)—CRR supports synchronous and asynchronous replication between remote sites over FC and a WAN. Synchronous replication is supported when the remote sites are connected through FC and provides a zero recovery point objective (RPO). Asynchronous replication provides crash-consistent protection and recovery to specific points in time, with a small RPO.
- Concurrent Local and Remote Replication (CLR)—CLR is a combination of CRR and CDP and provides concurrent local and remote data protection.

In RecoverPoint, CDP is normally used for operational recovery, while CRR is normally used for disaster recovery. The solution described in this white paper uses a RecoverPoint CLR configuration.

With RecoverPoint, data recovery can be performed locally and/or remotely by rewinding the target volumes back to a selected point in time through the use of earlier versions of data saved in the journal. RecoverPoint uses either FC or an IP network to send data over a WAN. This solution uses an IP network for replication of data between sites.

EMC RecoverPoint/SE, used in this solution, is the entry-level offering that allows replication and Continuous Data Protection for the EMC VNX series arrays.

RecoverPoint appliances (RPA)

The RecoverPoint appliance (RPA) is a 1U server that runs RecoverPoint software and includes four 4 gigabit FC connections and two 1 gigabit Ethernet connections. For local-only replication, two RPAs are installed at the protected site. For remote replication (or both local and remote replication), four RPAs are installed, two at the protected site and two at the recovery site.

RecoverPoint splitter

RecoverPoint uses lightweight splitting technology on the application server, in the fabric, or in the array to mirror application writes to the RecoverPoint cluster. CLARiiON and VNX arrays have integrated RecoverPoint splitters that operate in each storage processor, ensuring that the RecoverPoint application receives a copy of each write. The array-based splitter is the most effective write splitter for VMware replication, enabling replication of VMware Virtual Machine File System (VMFS) and raw device mapping (RDM) volumes without the cost or complexity of additional hardware. The splitter supports both FC and iSCSI volumes presented by the CLARiiON/VNX arrays to any host, including an ESX server. For this solution, array-based write splitters run inside the storage processors on the VNX5700 array.

RecoverPoint journals

RecoverPoint journals store time-stamped application writes for later recovery to selected points in time. Three journals are provisioned for local and remote replication—a production journal at the production site, and a history journal at both the production and recovery sites. For synchronous replication, every write is retained in the history journal for recovery to any point in time. For asynchronous replication, several writes are grouped before delivery to the history journal—this supports recovery to significant points in time. Bookmarked points in time can be created automatically or manually to enable recovery to specific application or system events. The production journals, one per consistency group, record the system delta-marking information. Production journals do not contain snapshots for point-in-time recovery.

RecoverPoint consistency groups

RecoverPoint uses replication sets and consistency groups to ensure the consistency and write-order fidelity of point-in-time images. A replication set defines an association between a production volume and the local and/or remote volumes to which it is replicating. A consistency group logically groups replication sets that must be consistent with one another. The consistency group ensures that updates to the production volumes are written to the replicas in consistent write order and that the replicas can always be used to continue working from or to restore the production source. RecoverPoint replication is policy-driven. A replication policy, based on a particular business need, can be uniquely specified for each consistency group. This policy governs the replication parameters for the consistency group—for example, the RPO and recovery time objective (RTO) for the consistency group, and its deduplication, data compression, and bandwidth reduction settings.

EMC Replication Manager

EMC Replication Manager manages EMC point-in-time replication technologies through a centralized management console. Replication Manager coordinates the entire data replication process—from discovery and configuration to the management of multiple, application-consistent, disk-based replicas. Replication Manager is used to safeguard and protect your business-critical applications, such as Oracle, Microsoft Exchange Server, Microsoft SQL Server, and VMware-based virtual machines.

EMC Replication Manager, when used with EMC RecoverPoint, provides solutions that eliminate backup windows without impacting production. This integration helps you to create application-consistent copies of business-critical data—at specific points in time—locally with RecoverPoint CDP, remotely with RecoverPoint CRR, or both locally and remotely with RecoverPoint CLR.

With Replication Manager, RecoverPoint replicas can be restored to any significant point in time that falls within the protection window. Additionally, Replication Manager can create more point-in-time copies of the same production volumes using other EMC replication technologies. For example, in addition to creating specific point-in-time copies of production volumes using CDP/CRR/CLR, Replication Manager can also create clones or snaps of the same volumes using EMC SnapView™ technology, which is available on VNX and CLARiiON storage systems.

Replication Manager leverages Microsoft Volume Shadow Copy Service (VSS) to provide point-in-time recovery and roll-forward recovery for Microsoft Exchange Server 2010 by using Copy and Full backup modes. Both modes back up the databases and transaction logs, but only Full mode truncates the logs after a successful replica validation. Because these snapshots are transportable, they can be used for repurposing. For example, if your server attaches to a SAN, you can mask the shadow copy from (make it invisible to) the production server but unmask it from (make it visible to) another server so that the copy can be reused for backup or mailbox-level recovery.

For more information about EMC RecoverPoint and EMC Replication Manager integration points, review the white paper entitled *EMC Replication Manager and EMC RecoverPoint* at <http://www.emc.com>.

For more information about EMC Replication Manager integration with Microsoft Exchange Server, review the *Replication Manager Product Guide* and *Replication Manager Support for Exchange 2010 Technical Notes* available on EMC Powerlink at <http://powerlink.emc.com>.

EMC Data Protection Advisor

EMC Data Protection Advisor (DPA) provides the visibility necessary for your organization to understand all of the details involved in managing large and complex data protection environments. DPA software automatically monitors, analyzes, and provides alerts on many aspects of a company's IT infrastructure. In this solution, DPA is used to analyze and monitor the replication environment.

DPA is a sophisticated reporting and analytics platform that provides customers with full visibility into the effectiveness of their data protection strategy. DPA does this by monitoring all of the technologies that a customer uses to protect data, including backup software, storage arrays, file servers, and tape libraries.

DPA's sophisticated reporting engine provides highly customizable reports to highlight problems within the environment and enables customers to perform capacity management, service level reporting, chargeback, change management, and troubleshooting.

DPA's Predictive Analysis Engine provides customers with early warning of problems that might occur and generates alerts that enable customers to resolve problems sooner to reduce business impact.

DPA is designed to help users identify and fix performance and capacity management (PCM) issues within large enterprises. DPA consists of a number of loosely-coupled processes and data stores for holding configuration and gathered data. In addition, DPA has a number of interfaces to import and export information, which allow for integration with other aspects of a large enterprise and to provide a significant part of an overall distributed systems management strategy.

EMC PowerPath/VE EMC PowerPath/VE provides intelligent, high-performance path management with path failover and load balancing optimized for EMC and selected third-party storage systems. PowerPath/VE supports multiple paths between a vSphere host and an external storage device. Having multiple paths enables the vSphere host to access a storage device even if a specific path is unavailable. Multiple paths can also share the I/O traffic to a storage device. PowerPath/VE is particularly beneficial in highly available environments because it can prevent operational interruptions and downtime. The PowerPath/VE path failover capability avoids host failure by maintaining uninterrupted application support on the host in the event of a path failure (if another path is available).

VMware vSphere VMware vSphere uses the power of virtualization to transform data centers into simplified cloud computing infrastructures and enables IT organizations to deliver flexible and reliable IT services. vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the data center. As a cloud operating system, vSphere manages large collections of infrastructure (such as CPUs, storage, and networking) as a seamless and dynamic operating environment and also manages the complexity of a data center.

VMware vCenter Site Recovery Manager VMware vCenter Site Recovery Manager (SRM) is a disaster recovery framework that integrates with EMC RecoverPoint to automate recovery of virtual machines and their storage so that recovery becomes as simple as pressing a single button. SRM is an extension to VMware vCenter that enables integration with array-based replication, discovery and management of replicated datastores, and automated migration of inventory from one vCenter instance to another.

SRM itself does not replicate any data. For this, SRM leverages an external replication solution such as RecoverPoint. SRM servers coordinate the operations of the replicated storage arrays and vCenter servers at the production and recovery sites so that, as virtual machines at the production site shut down, virtual machines at the recovery site start up and assume responsibility for providing the same services, using the data replicated from the production site.

Migration of protected inventory and services from one site to the other is controlled by a recovery plan that specifies the order in which virtual machines shut down and start up, the compute resources that are allocated, and the networks they can access. New SRM version 5 can now leverage native VMware vSphere replication or third-party, storage-based replication (RecoverPoint in this solution) to provide centralized management of recovery plans, enable non-disruptive testing, and automate site recovery and migration processes. One of the key new features in SRM version 5 is its automated failback capabilities, which provide automatic failback and re-protection of virtual machines by automatically reversing replication to the original site.

**EMC RecoverPoint
Storage
Replication
Adapter for
VMware vCenter
SRM**

VMware vCenter SRM integrates with the underlying replication product, RecoverPoint, through a RecoverPoint Storage Replication Adapter (SRA). The adapter enables SRM to identify which virtual machines are being replicated and to coordinate the execution of recovery plans with the replication layer. SRM requires separate vCenter Server instances at both the production and failover site. SRM instances are deployed at both sites and integrate directly with their local vCenter server instances.

New features in the RecoverPoint SRA adapter include fully automated failback (without the need for a separate plug-in). For more information, refer to *EMC RecoverPoint Storage Replication Adapter for VMware Site Recovery Manager Version 2.0 Release Notes*.

The RecoverPoint SRA supports the discovery of arrays attached to RecoverPoint and the discovery of consistency groups enabled for management by SRM. SRA also supports SRM functions, such as failover and failover testing, by mapping SRM recovery plans to the appropriate RecoverPoint actions.

Overall solution architecture and design

This solution demonstrates how EMC and VMware technologies can be leveraged to provide high availability and disaster recovery protection for Microsoft Exchange Server 2010 enterprise deployments at multiple, geographically distributed sites.

The solution, as validated, demonstrates the replication and activation of an entire Exchange Server 2010 Database Availability Group (DAG), along with its associated Exchange virtual machines, across a WAN between two (simulated) active data centers.

Clarification of terminology

In an active/active site configuration, such as the one this solution represents, both sites (or all sites, if there are more than two) are active production sites, that are locally protected and also globally protected by every other site. In other words, each site can be referred to as both a protected site and a recovery site. With this in mind, the subsequent sections refer to a pair of protected/recovery sites, but each site can reverse roles with the other, depending on where the disaster event occurs. Further, in a configuration with more than two active sites, every site is both a protected site and a potential recovery site for any peer site affected by a disaster.

Note: Even though this solution demonstrates the protection of an active/active site configuration, the solution can also be used to protect active/passive site configurations.

Design attributes

Table 1 presents attributes of the solution design, including specific protection requirements.

Table 1. Solution design attributes

Attribute	Description
Number of sites	2 sites
Number of users per site	<ul style="list-style-type: none"> 10,000 users per site 20,000 users supported at each site during DR event
Exchange 2010 user profile	150 messages sent/received/day (0.15 IOPS)
Mailbox size	2 GB
Database read/write ratio	3:2 in Mailbox Resiliency configuration
High availability design	<ul style="list-style-type: none"> 1 Exchange DAG per site 4 Mailbox servers per DAG 2 database copies per DAG
Mailbox server design	<ul style="list-style-type: none"> 5,000 users per server (2,500 active users / 2,500 passive users) 10 databases per server 500 users per database 6 vCPUs, 32 GB RAM per Mailbox server virtual machine

Attribute	Description
Storage design	<ul style="list-style-type: none"> • 1 VNX5700 storage array per site • 2 TB NL-SAS (7.2k rpm) disks for Exchange database and logs (RDM/P) • 2 TB NL-SAS (7.2k rpm) disks for virtual machine OS (VMFS) • 600 GB SAS (10k rpm) disks for RecoverPoint journals
vSphere host design	<ul style="list-style-type: none"> • 4 Intel-based servers with 4 six-core Xeon X7460 processors @ 2.66 GHz (24 CPUs) and 128 GB RAM. 2 servers per site, each supports 4 Exchange Mailbox Role virtual machines. • 4 Intel based servers with 4 six-core Xeon X7350 Processors @ 2.93 GHz (16 CPUs) and 64 GB RAM. 2 servers per site, each supports 4 Exchange Client Access and Hub Transport Role virtual machines.
Simulated RTO and RPO	<ul style="list-style-type: none"> • Local and remote RPO: 5 minutes • Local RTO: 30 minutes • Remote RTO: 2 hours
Connectivity between sites	1 Gb/s Ethernet
Storage connectivity to physical host	2 x 8 Gb/s FC ports
Data protection	<ul style="list-style-type: none"> • Daily full backups (hardware VSS) • 24 application-consistent snapshots per day (24 hours) • Unlimited crash-consistent snapshots per day (24 hours)*
<p>*Note: The number of crash-consistent snapshots depends on the size of the RecoverPoint journal, the application change rate, and policy settings. Refer to RecoverPoint design on page 26.</p>	

Virtualization

This solution virtualizes Exchange 2010 on VMware vSphere 5 and leverages EMC, Microsoft, and VMware best practices and guidelines for deploying Exchange in virtual configurations.

For more information on Exchange Server 2010 deployments for high availability and site resiliency, refer to <http://technet.microsoft.com/en-us/library/dd638121.aspx>.

For more details on Exchange Server 2010 deployments on VMware vSphere, refer to *Microsoft Exchange 2010 on VMware Best Practices Guide* (<http://www.vmware.com>).

Active/active site configuration

Figure 1 illustrates the Exchange Server 2010 environment deployed for this solution. The environment represents two active sites with four vSphere hosts at each site.

Each site provides enough performance and capacity to support up to 20,000 active users (10,000 users during normal operations and another 10,000 users during a disaster recovery event).

For this solution, the Active Directory and DNS services at each site are provided outside of the virtual infrastructure. The Active Directory is configured to meet all of the requirements for this Exchange deployment. Redundant Active Directory servers, as well as DNS servers, are provisioned at both sites.

Data protection

Data protection for both sites is provided by RecoverPoint, with Replication Manager initiating application-consistent (VSS) copies for Exchange databases and logs. In the event of a site-specific disaster, or if you want to move the servers or test disaster recovery, vCenter SRM is used to coordinate an entire Exchange server failover from one site to another site and back again.

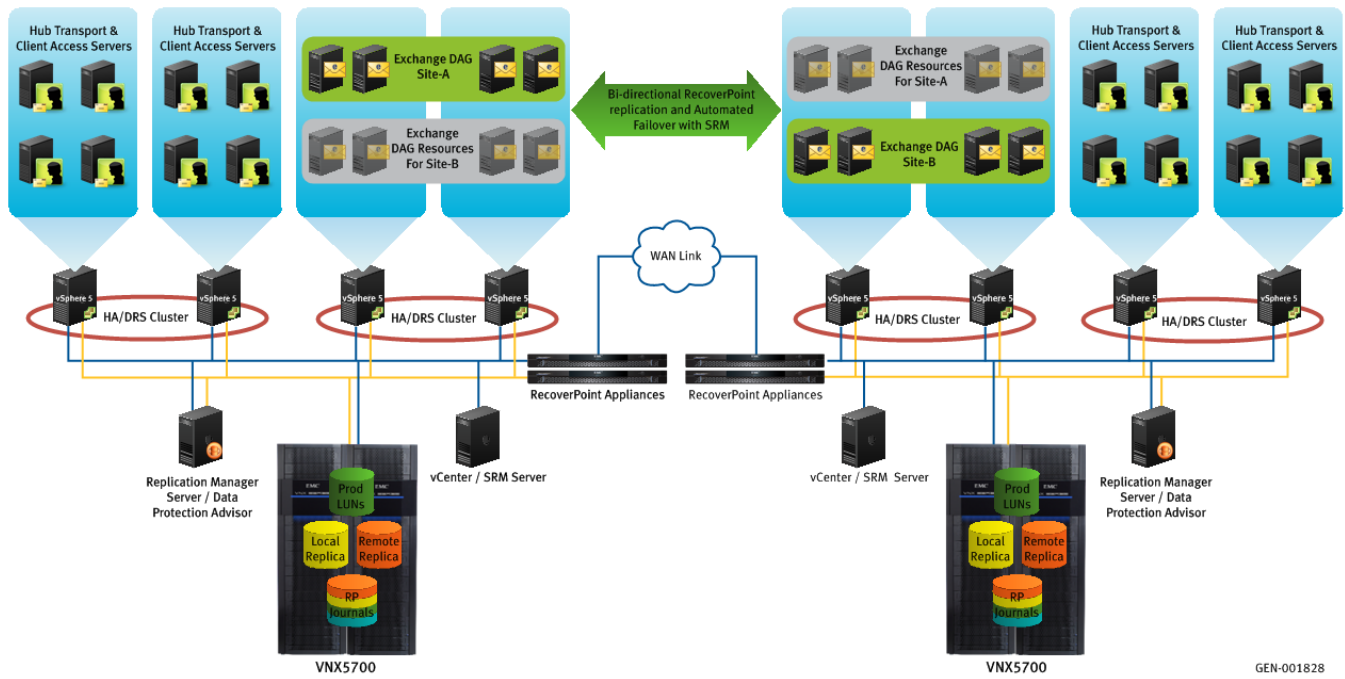


Figure 1. Solution architecture diagram

Exchange Server 2010 Design

One DAG with four Mailbox servers is deployed at each site. A DAG is a set of mailbox servers (up to 16) that replicate to each other, providing database-level protection from outages. DAG uses Exchange Continuous Replication (log shipping) and a subset of Windows failover clustering technologies to provide continuous high availability (HA) and site resiliency. It replaces earlier Exchange and Windows failover cluster-based technologies used for HA in Exchange 2007, such as single copy clusters (SCC), continuous cluster replication (CCR), and standby continuous replication (SCR). DAG members can reside in the same data center or they can be configured across multiple sites in cross-site, “stretched” configurations.

An Exchange Management Console screenshot, presented in Figure 2, shows details about the Exchange DAGs configured for this solution. An Exchange 2010 DAG is deployed at each site with four member servers and two database copies. Each Exchange Mailbox server is designed to support 5,000 users with a profile of 150 messages/user/day at 0.15 IOPS. During normal operation, each server provides service to 2,500 active users and can service an additional 2,500 users during a local switchover.

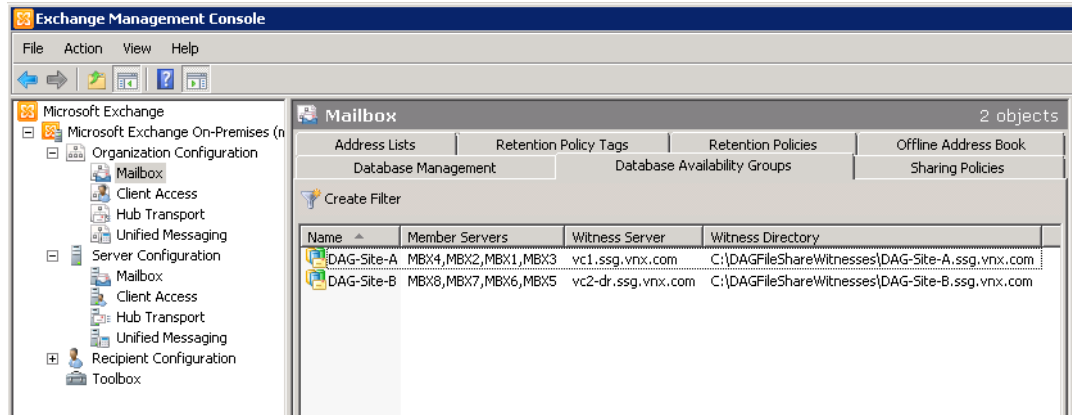


Figure 2. Exchange Management Console GUI: Exchange Server 2010 Database Availability Group configuration

Each Mailbox server is configured with 10 databases, five active and five passive. All databases are balanced and distributed between Mailbox servers within the DAG and between vSphere hosts to eliminate a single point of failure. This design eliminates downtime during vSphere host or Mailbox server virtual machine maintenance.

	Site A				Site B			
	ESXi Host 1		ESXi Host 2		ESXi Host 3		ESXi Host 4	
	Exchange 2010 DAG				Exchange 2010 DAG			
	MBX1	MBX2	MBX3	MBX4	MBX5	MBX6	MBX7	MBX8
DB1-5	Active		Passive		Active		Passive	
DB6-10	Passive		Active		Passive		Active	
DB11-15		Active		Passive		Active		Passive
DB16-20		Passive		Active		Passive		Active
DB21-25					Active		Passive	
DB26-30					Passive		Active	
DB31-35						Active		Passive
DB36-40						Passive		Active

Figure 3. Exchange Server 2010 DAG design: Database distribution between Mailbox servers and vSphere hosts

Figure 4 presents a screenshot from the Exchange Management Console showing the database configuration for one Mailbox server. It shows that ten databases are deployed on Mailbox server MBX1, where five of these databases are active and the other five are passive.

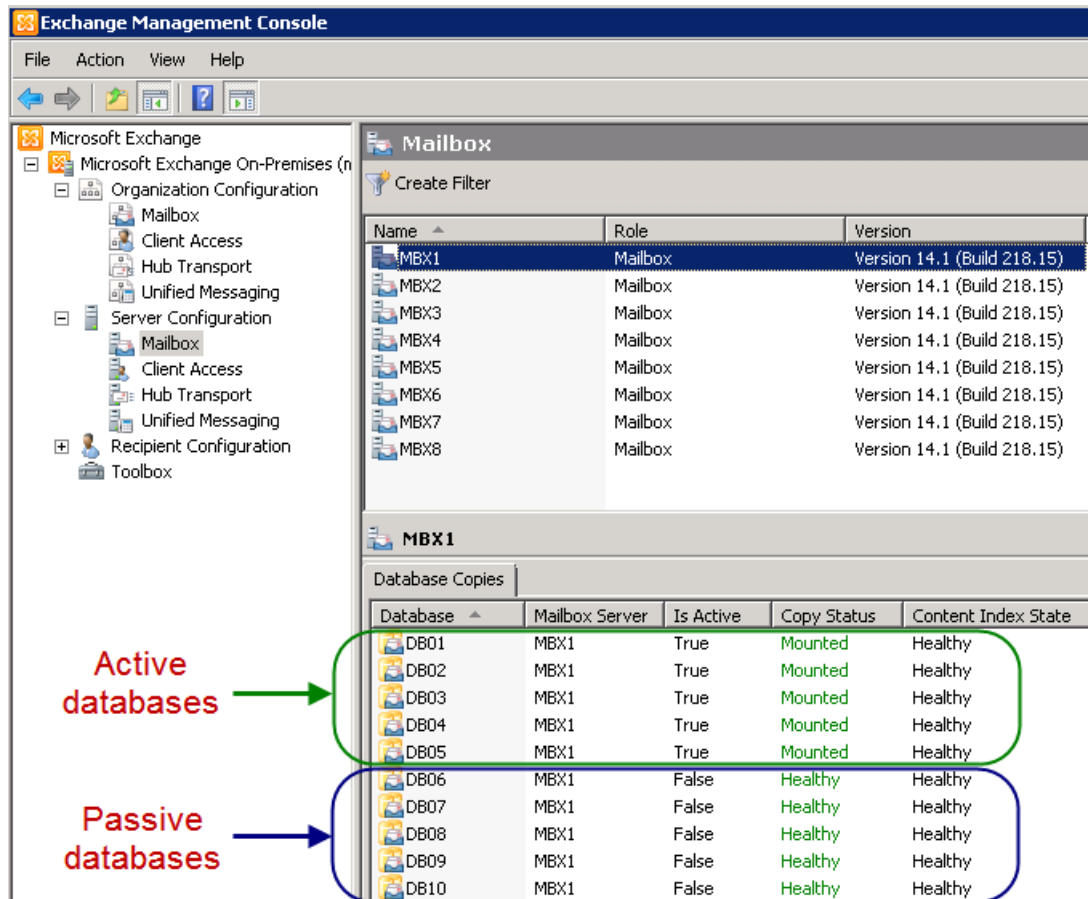


Figure 4. Exchange Management Console GUI: Mailbox server configuration

Exchange storage design

To support the Exchange Server 2010 user profile and mailbox size requirements for this solution (see Table 1 for details), each site uses an EMC VNX5700 storage array with 2 TB 7.2k rpm NL-SAS drives to house databases and logs.

To simplify design and deployment, a Mailbox server building block for each Mailbox server is designed and configured with the appropriate number of drives to support the required performance and capacity.

Building block design approach

A Mailbox server building block represents the amount of storage (I/O and capacity) and server compute (CPU, memory) resources required to support a specific number of Exchange Server 2010 users. The amount of required resources is derived from a specific user profile type (messages sent/received, IOPS per user, and mailbox size). Using the building block approach simplifies the design and implementation of Exchange Server 2010.

Once the initial building block is designed, it can be easily reproduced to support the required number of users in your enterprise. By using this approach, EMC customers can now create their own building blocks that are based on their company's specific Exchange environment requirements. This approach is very helpful when future growth is expected because it makes Exchange environment expansion simple and straightforward. EMC best practices involving the building block approach for Exchange Server design have proven to be very successful in many customer implementations.

For more details on how to use the EMC building-block methodology for Exchange 2010 deployments, see the *Microsoft Exchange 2010: Storage Best Practices and Design Guidance for EMC Storage* (<http://www.emc.com>).

Mailbox server virtual machine building block details

For this solution, the Exchange Mailbox server virtual machine design is based on a building block for 5,000 users. The storage, CPU, and memory resources for this building block are derived from the design attributes presented in [Table 1](#) on page 15. For additional Microsoft guidelines considered for this solution's design, see <http://technet.microsoft.com/en-us/library/ee712771.aspx>.

[Table 2](#) presents a summary of the Exchange Mailbox server virtual machine building block for 5,000 users with a 150-message profile and a 2 GB mailbox quota.

Table 2. Mailbox server building block summary

Users per Exchange Mailbox server virtual machine	Disks per Exchange Mailbox server virtual machine	vCPUs per Exchange Mailbox server virtual machine	Memory per Exchange Mailbox server virtual machine
5,000	20 (16+4) 16 x 2 TB NL-SAS disks in RAID 1/0 for databases 4 x 2 TB NL-SAS disks in RAID 1/0 for logs	6 CPUs	32 GB

Table 3 presents storage design details for a Mailbox server building block with 5,000 users.

Table 3. Storage design details for a 5,000-user Mailbox server building block

Attribute	Details
Number of users	5,000 (in a switchover condition)
User profile	150 messages sent/received (0.15 IOPS)
Mailbox size	2 GB
Disk type	2 TB NL-SAS, 7.2k rpm
Storage configuration/RAID type	20 disks per building block for Exchange data: <ul style="list-style-type: none"> • 16 disks in a RAID 1/0 storage pool for databases • 4 disks in a RAID 1/0 RAID group for logs
Databases per server/users per database	10 databases per server, 500 users per database
LUNs/LUN sizes	21 LUNs per Mailbox server virtual machine: <ul style="list-style-type: none"> • 10 database LUNs (1.6 TB) (RDM/P) • 10 log LUNs (100 GB) (RDM/P) • 1 LUN for virtual machine OS (200 GB) (VMFS)

About storage pools

When deploying storage for Exchange Server 2010 in a virtualized environment on a VNX system, you can choose to use either storage pools or traditional RAID groups.

The use of storage pools simplifies storage provisioning. Traditional storage provisioning with only RAID groups restricts the number of disks you can have in a group to 16. Storage pools, on the other hand, enable you to manage potentially hundreds of disks at a time. Such pool-based provisioning provides benefits similar to metaLUN striping across many drives but, unlike metaLUNs, storage pools require minimal planning and management effort. Storage pools support the same RAID protection levels as RAID groups do: RAID 5, RAID 6, and RAID 1/0.

Using homogeneous thick storage pools (pools with the same disk type) for deploying Exchange storage on VNX systems can help to simplify design and LUN provisioning.

There are three main models that can be used when designing storage pools to be used for Exchange 2010 deployed in physical or virtual environments.

- One storage pool per Mailbox server—In this model, a single storage pool is deployed for each Exchange Mailbox server virtual machine. This option provides more granularity and easy deployment with a building block methodology, and it also provides database copy isolation if the server is deployed in a DAG with multiple copies. This deployment model also works well when Exchange Server 2010 is deployed without mailbox resiliency (no DAG).

- One storage pool per vSphere host—This model works well when Exchange is deployed in a virtual environment with multiple DAG copies. This model can help reduce the number of storage pools deployed in the environment. It is important to be sure that copies of the same database replicated to different virtual machines are *not* hosted by the same ESX server. With careful design, this model can provide database copy isolation and balanced performance.
- One storage pool per database copy—This model provides database copy isolation and, in many cases, can minimize the number of pools that must be deployed. However, this model requires careful distribution of LUNs to vSphere hosts and virtual machines. Database copies from each pool are distributed to multiple Mailbox server virtual machines.

VNX storage pool design

For this solution, one homogeneous RAID 1/0 storage pool with 40 NL-SAS disks is deployed for each vSphere host hosting Exchange Mailbox Server Role virtual machines. This configuration supports databases for two Mailbox Server Role virtual machines on each vSphere host. It provides just over 36.6 TB of usable storage and meets the IOPS and total mailbox capacity requirements for two building blocks, each with 5,000 very heavy users and a 2 GB mailbox limit.

To support log capacity requirements, four 2 TB NL-SAS drives are deployed in a RAID 1/0 (2+2) RAID group configuration for each vSphere host with two Exchange Mailbox Server Role virtual machines.

From each database storage pool, 20 x 1.6 TB database LUNs are configured to accommodate a 1 TB database and index. The LUNs are distributed among Exchange Mailbox server virtual machines on the same vSphere host, and the LUNs are configured as RDM/P volumes.

The same approach is used for log LUNs. 20 x 100 GB LUNs are created from each RAID group. The LUNs are distributed among Mailbox server virtual machines on the same vSphere host, and the LUNs are configured as RDM/P volumes.

Note: Logs and databases in this solution are separated on different disks to satisfy the protection requirements presented in [Table 3](#).

At each site, 200 GB virtual machine operating system (OS) volumes are configured on the same VNX5700 storage array with 2 TB NL-SAS drives configured as RAID1 (2+2) RAID groups. Each OS LUN is formatted as a VMFS volume on the vSphere host.

Figure 5 illustrates storage configuration used for the Mailbox server virtual machines in this solution.

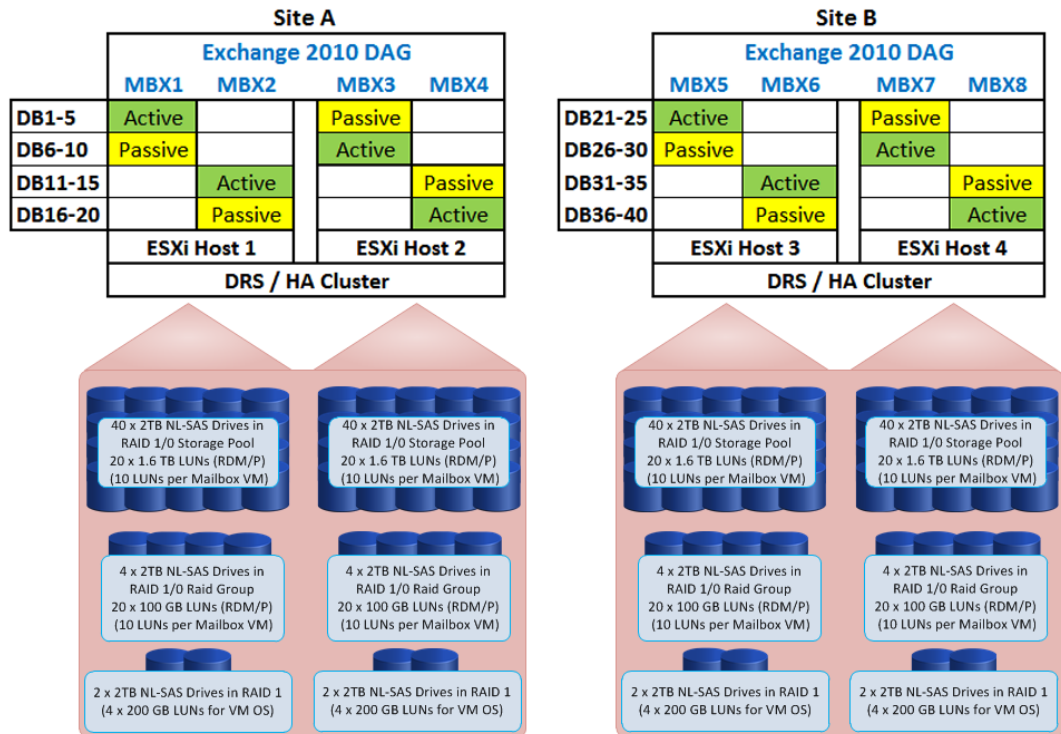


Figure 5. Storage design for Exchange Server 2010 Mailbox servers

vSphere server design

To satisfy the requirement to continue to support all 20,000 Exchange users in the event of a disaster at either site, four vSphere servers are deployed at each site. Two vSphere servers with 32 CPUs and 128 GB of RAM capacity support eight Mailbox Server Role virtual machines (with four virtual machines per vSphere server). Another two vSphere servers with 16 CPUs and 128 GB of RAM support eight Exchange Client Access and Hub Transport Role virtual machines (again, with four virtual machines per vSphere server).

Table 4 provides additional details about the host/role associations at the two sites (referred to in the table as Site A and Site B).

Note: Grey rows indicate server / compute resources active only during site failover.

Table 4. vSphere host and Exchange Server role associations

Site	vSphere host	Exchange Server Role / hostname	CPUs	Memory (GB)
A	ESX1 (24 CPUs, 128 GB RAM)	Mailbox Role - MBX1	6	32
		Mailbox Role - MBX2	6	32
		Mailbox Role - MBX5 from Site B	6	32
		Mailbox Role - MBX6 from Site B	6	32
	ESX2 (24 CPUs, 128 GB RAM)	Mailbox Role - MBX3	6	32
		Mailbox Role - MBX4	6	32
		Mailbox Role - MBX7 from Site B	6	32
		Mailbox Role - MBX8 from Site B	6	32
A	ESX3 (16 CPUs, 64 GB RAM)	Client Access and Hub Transport - CAS-HUB1	4	16
		Client Access and Hub Transport - CAS-HUB2	4	16
		Client Access and Hub Transport - CAS-HUB3	4	16
		Client Access and Hub Transport - CAS-HUB4	4	16
	ESX4 (16 CPUs, 64 GB RAM)	Client Access and Hub Transport - CAS-HUB5	4	16
		Client Access and Hub Transport - CAS-HUB6	4	16
		Client Access and Hub Transport - CAS-HUB7	4	16
		Client Access and Hub Transport - CAS-HUB8	4	16
B	ESX5 (24 CPUs, 128 GB RAM)	Mailbox Role - MBX5	6	32
		Mailbox Role - MBX6	6	32
		Mailbox Role - MBX1 from Site A	6	32
		Mailbox Role - MBX2 from Site A	6	32
	ESX6 (24 CPUs, 128 GB RAM)	Mailbox Role - MBX7	6	32
		Mailbox Role - MBX8	6	32
		Mailbox Role - MBX3 from Site A	6	32
		Mailbox Role - MBX4 from Site A	6	32
B	ESX7 (16 CPUs, 64 GB RAM)	Client Access and Hub Transport - CAS-HUB9	4	16
		Client Access and Hub Transport - CAS-HUB10	4	16
		Client Access and Hub Transport - CAS-HUB11	4	16
		Client Access and Hub Transport - CAS-HUB12	4	16
	ESX8 (16 CPUs, 64 GB RAM)	Client Access and Hub Transport - CAS-HUB13	4	16
		Client Access and Hub Transport - CAS-HUB14	4	16
		Client Access and Hub Transport - CAS-HUB15	4	16
		Client Access and Hub Transport - CAS-HUB16	4	16

Exchange local protection

For local protection and high availability, an Exchange 2010 DAG is deployed at each site. In most cases, this is sufficient to protect Exchange during Mailbox server failure or maintenance. However, DAG does not provide protection from logical database corruption, because such corruption can be replicated to other copies. To protect from logical corruption, VSS backups or a point-in-time recovery solution can be implemented.

Note: An Exchange Server 2010 lagged copy can also provide protection from database corruption. But lagged copies cannot provide point-in-time recovery. After a lagged copy is used for recovery, it must be re-created (seeded) again. For additional information about Exchange database corruption types and lagged copy limitations, visit <http://technet.microsoft.com/en-us/library/dd335158.aspx>.

In this solution, EMC RecoverPoint CLR provides protection from logical database corruption. RecoverPoint uses bookmarks (“named snapshots”) to provide point-in-time recovery. Each bookmark is stored in a RecoverPoint journal configured at both sites (CDP and CRR copies) and data can be recovered from either site. EMC offers alternatives to RecoverPoint for protecting against logical database corruption, including EMC VNX SnapView snapshots and clones. The alternative that is best for you depends on your specific requirements.

RecoverPoint design

In this solution EMC RecoverPoint provides protection from Exchange logical database corruption, which can potentially replicate to other Exchange database copies within the DAG. RecoverPoint also provides near-instantaneous point-in-time recovery with minimal data loss. In this solution, RecoverPoint is configured as follows:

- The replication method is Concurrent Local and Remote Replication (CLR). This combines both CDP and CRR technologies and provides point-in-time recovery from either CDP or CRR copies.
- Each Mailbox server virtual machine is placed into its own consistency group with all replicated database and log LUNs, including the LUN containing the guest OS. This configuration ensures full consistency across all volumes used by the Mailbox Server Role virtual machines.
- Three journals are set up for each consistency group. Two journals reside at the protected site to support the protected volumes and their local replicas, and one journal resides at the recovery site to support the remote replica.

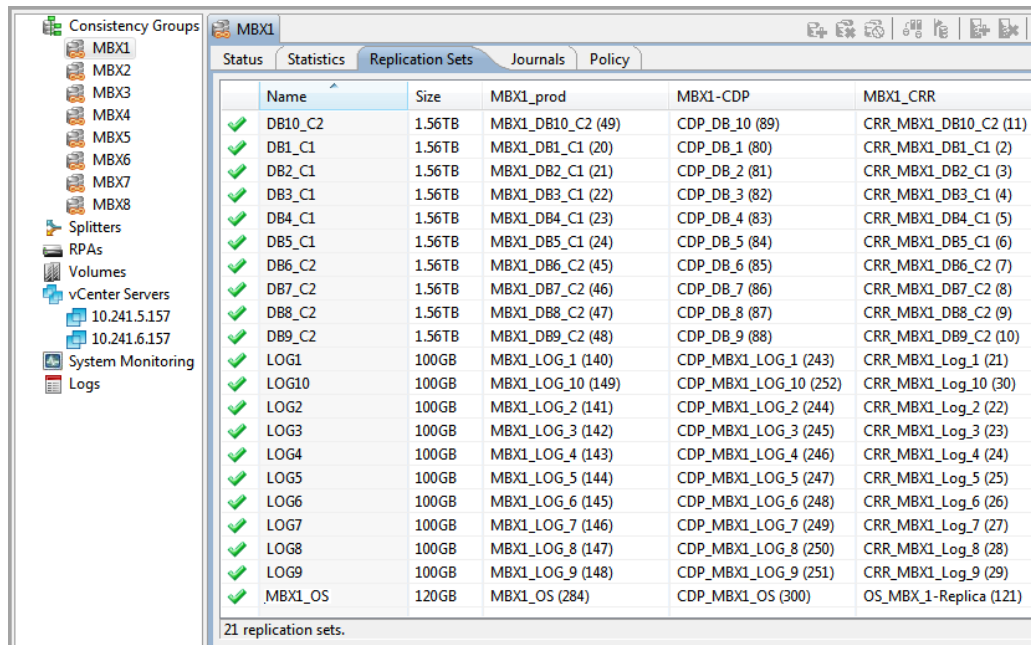
Note: RecoverPoint CRR alone can provide protection from logical database corruption, but CRR can replicate only a single DAG copy. With CRR, replicated Exchange databases can be mounted at the recovery site on any Exchange DAG member server, but the databases then require seeding to provide local native high availability for Exchange at the recovery site.

RecoverPoint consistency group design

RecoverPoint consistency groups are created for each replicated Mailbox server virtual machine. Separating replicated virtual machines into multiple consistency groups makes it possible to plan for different disaster recovery scenarios. For example, each Exchange Mailbox server virtual machine can be failed over separately from the other virtual machines, or the entire protected environment can be failed over in a single process.

Note: This approach does not prevent the recovery of a single database. Refer to [Exchange database recovery with Replication Manager](#) on page 60.

Figure 6 shows the details for the consistency group configuration with replication sets for one replicated Exchange Mailbox server virtual machine.



Name	Size	MBX1_prod	MBX1-CDP	MBX1_CRR
DB10_C2	1.56TB	MBX1_DB10_C2 (49)	CDP_DB_10 (89)	CRR_MBX1_DB10_C2 (11)
DB1_C1	1.56TB	MBX1_DB1_C1 (20)	CDP_DB_1 (80)	CRR_MBX1_DB1_C1 (2)
DB2_C1	1.56TB	MBX1_DB2_C1 (21)	CDP_DB_2 (81)	CRR_MBX1_DB2_C1 (3)
DB3_C1	1.56TB	MBX1_DB3_C1 (22)	CDP_DB_3 (82)	CRR_MBX1_DB3_C1 (4)
DB4_C1	1.56TB	MBX1_DB4_C1 (23)	CDP_DB_4 (83)	CRR_MBX1_DB4_C1 (5)
DB5_C1	1.56TB	MBX1_DB5_C1 (24)	CDP_DB_5 (84)	CRR_MBX1_DB5_C1 (6)
DB6_C2	1.56TB	MBX1_DB6_C2 (45)	CDP_DB_6 (85)	CRR_MBX1_DB6_C2 (7)
DB7_C2	1.56TB	MBX1_DB7_C2 (46)	CDP_DB_7 (86)	CRR_MBX1_DB7_C2 (8)
DB8_C2	1.56TB	MBX1_DB8_C2 (47)	CDP_DB_8 (87)	CRR_MBX1_DB8_C2 (9)
DB9_C2	1.56TB	MBX1_DB9_C2 (48)	CDP_DB_9 (88)	CRR_MBX1_DB9_C2 (10)
LOG1	100GB	MBX1_LOG_1 (140)	CDP_MBX1_LOG_1 (243)	CRR_MBX1_Log_1 (21)
LOG10	100GB	MBX1_LOG_10 (149)	CDP_MBX1_LOG_10 (252)	CRR_MBX1_Log_10 (30)
LOG2	100GB	MBX1_LOG_2 (141)	CDP_MBX1_LOG_2 (244)	CRR_MBX1_Log_2 (22)
LOG3	100GB	MBX1_LOG_3 (142)	CDP_MBX1_LOG_3 (245)	CRR_MBX1_Log_3 (23)
LOG4	100GB	MBX1_LOG_4 (143)	CDP_MBX1_LOG_4 (246)	CRR_MBX1_Log_4 (24)
LOG5	100GB	MBX1_LOG_5 (144)	CDP_MBX1_LOG_5 (247)	CRR_MBX1_Log_5 (25)
LOG6	100GB	MBX1_LOG_6 (145)	CDP_MBX1_LOG_6 (248)	CRR_MBX1_Log_6 (26)
LOG7	100GB	MBX1_LOG_7 (146)	CDP_MBX1_LOG_7 (249)	CRR_MBX1_Log_7 (27)
LOG8	100GB	MBX1_LOG_8 (147)	CDP_MBX1_LOG_8 (250)	CRR_MBX1_Log_8 (28)
LOG9	100GB	MBX1_LOG_9 (148)	CDP_MBX1_LOG_9 (251)	CRR_MBX1_Log_9 (29)
MBX1_OS	120GB	MBX1_OS (284)	CDP_MBX1_OS (300)	OS_MBX_1-Replica (121)

21 replication sets.

Figure 6. RecoverPoint consistency group configuration

RecoverPoint storage requirements

In addition to the requirement of having replication volumes at each site, RecoverPoint requires a very small volume, known as the repository volume, which must be configured on the SAN-attached storage at each site for each RPA cluster. The repository volume stores configuration information about the RPAs and consistency groups, which enables a properly functioning RPA to seamlessly assume the replication activities from a failing RPA in the same cluster.

RecoverPoint journal design

RecoverPoint journal volumes store consistent point-in-time snapshots (or bookmarks) to enable point-in-time recovery of Exchange data. These bookmarks allow DVR-like rollbacks of the data to different points in time. For each consistency group in the RecoverPoint CLR configuration, three journals are set up, two at the local site to support the protected volumes and their local replicas, and one at the remote site to support the remote replica. The production journals, one for each consistency group, contain system delta-marking information. Production journals do not contain snapshots for point-in-time recovery.

Considerations when sizing journal volumes

The two most important considerations when sizing storage for journal volumes are:

- Performance
- Protection window

For optimal journal performance, ensure that you choose the appropriate RAID and drive types. The solution, as validated, uses 600 GB 10k rpm SAS drives in a RAID 1/0 configuration for journal volumes, which are deployed in RAID 1/0 RAID groups.

The protection window depends on the change rate of the application and RPO requirements (how far back Exchange administrators want to go to for point-in-time recovery). Both of these factors determine the required journal size.

EMC recommends sizing your journals to meet your RPO requirements. Determining the journal size requires you to calculate the expected peak change rate in your environment. Twenty percent of the journal must be reserved for the target side log and five percent for internal system needs.

Sizing journal volumes

Use the following formula to calculate journal volume size based on the expected change rate and required rollback time period:

$$\text{Journal size} = \frac{(\text{new data writes per second}) \times (\text{required rollback time in seconds})}{(1 - \text{target side log size})} \times 1.05$$

For example, to support a 3-day (72-hour) rollback requirement (259,200 seconds), with 5 Mbps of new data writes to the replication volumes in a consistency group, the calculation is as follows:

$$\text{Journal size} = 5 \text{ Mbps} \times 259,200 \text{ seconds} / 0.8 \times 1.05 = 1,701,000 \text{ Mb} (\sim 213 \text{ GB})$$

In this calculation, 259,200 seconds represents a 72-hour rollback window, and 0.8 represents 20 percent for reserved journal space. As a general rule, EMC recommends making the journal size at least 20 percent of size of the data being replicated, when change rates are not available.

For full details about configuring and sizing RecoverPoint journal volumes, see the *EMC RecoverPoint Administrator's Guide* available on <http://powerlink.emc.com>.

For this RecoverPoint CLR configuration, 500 GB journal space is allocated for each local and remote copy of each consistency group. (In this solution, each consistency group contains all LUNs for one Mailbox server.) The 500 GB journal is composed of five 100 GB journal volumes. The deployment of smaller volumes enables journal space to be increased, seamlessly, by smaller increments. For best performance, it is a best practice to create and grow journals by using volumes that are identical in size, because RecoverPoint re-stripes the volumes during expansion.

Production journals, which do not contain snapshots for point-in-time recovery, are allocated 20 GB of storage. These journals are composed of four 5 GB journal volumes.

Figure 7 shows details of the journal volumes deployed in this solution as viewed with the RecoverPoint Management Console.

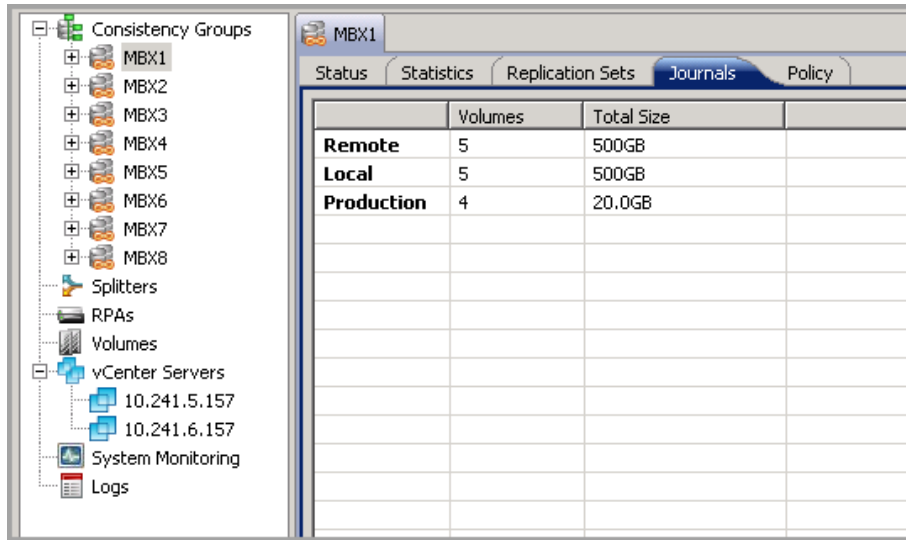


Figure 7. RecoverPoint journal volume configuration

Note: EMC strongly recommends adhering to the RecoverPoint guidelines presented in the *EMC RecoverPoint Administrator's Guide* to ensure that your journal volumes meet your specific RPO requirements.

Table 5 provides a summary of the RecoverPoint storage configuration for this solution. For validation of journal space under Exchange load, see [RecoverPoint journal usage under Exchange load](#) on page 51.

Note: Journal volume sizes (see Table 5) were configured to accommodate multiple test scenarios and might differ from guidelines listed in the *EMC RecoverPoint Administrator's Guide*. Refer to the section [RecoverPoint journal usage under Exchange load](#) for information about journal usage during Exchange workload validation testing.

Table 5. RecoverPoint storage configuration for this solution

Volume type	Site A and Site B	Comments
Repository	3 GB	1 volume per storage array/RPA cluster
Journal	40 x 100 GB 32 x 5 GB	<ul style="list-style-type: none"> Journal LUNs configured from RAID 1/0 groups. 500 GB journal per consistency group: 5 x 100 GB volumes for each local and remote copy. 20 GB journal per consistency group: 4 x 5 GB volumes for each production copy.
Replication	80 x 1.6 TB for database volumes 80 x 100 GB for log volumes	Replication volume sizes must be equal or larger than production volumes.

Consistency group policy configuration for management by SRM

When RecoverPoint is deployed with VMware vCenter SRM, RecoverPoint must be configured to enable automatic failover capability. This configuration is accomplished by using the “Group is managed by SRM, RecoverPoint can only monitor” policy setting in the RecoverPoint Management Console, as shown in [Figure 8](#).

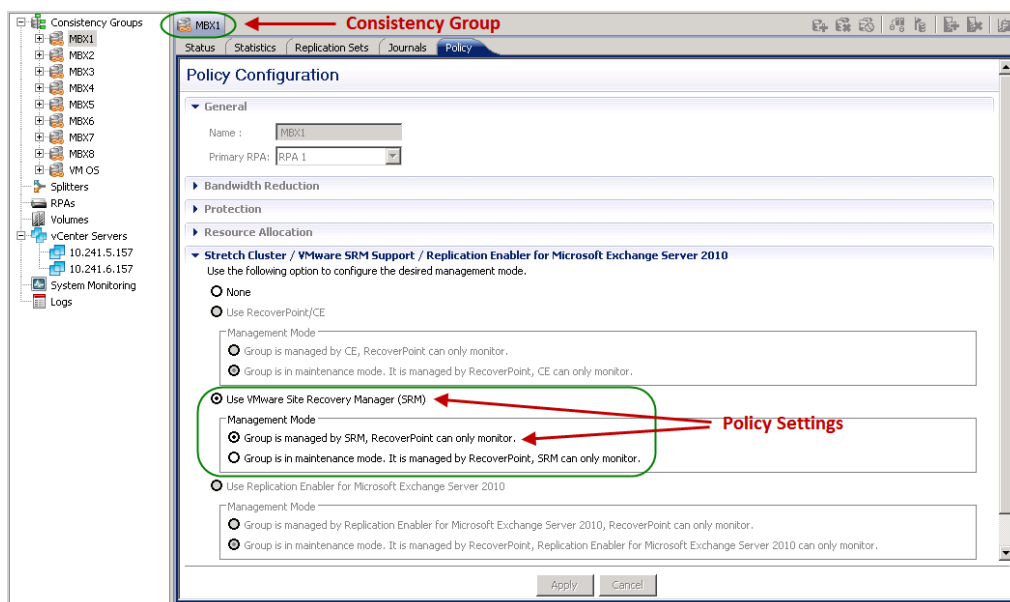


Figure 8. Consistency group policy configuration for management by SRM

For additional details on configuring RecoverPoint, consult the *EMC RecoverPoint Release Notes* and *EMC RecoverPoint Administrator's Guide*.

RecoverPoint group sets

A RecoverPoint group set is a set of consistency groups to which the system applies parallel bookmarks at a user-defined frequency. A group set allows you to automatically bookmark a set of consistency groups so that the bookmark represents the same recovery point in each consistency group in the group set. This allows you to define consistent recovery points for consistency groups distributed across different RPAs. Group sets are useful for consistency groups that depend on one another or that must work together as a single unit.

Group Sets configuration

For this solution, two group sets are configured, where each group set is made up of four consistency groups, each representing Mailbox servers in a DAG at each site. This configuration enables easy recovery, because all bookmarks for every consistency group within a group set have the same RPO and timestamp with a frequency of 15 minutes.

Figure 9 shows the group set information configured for this solution.

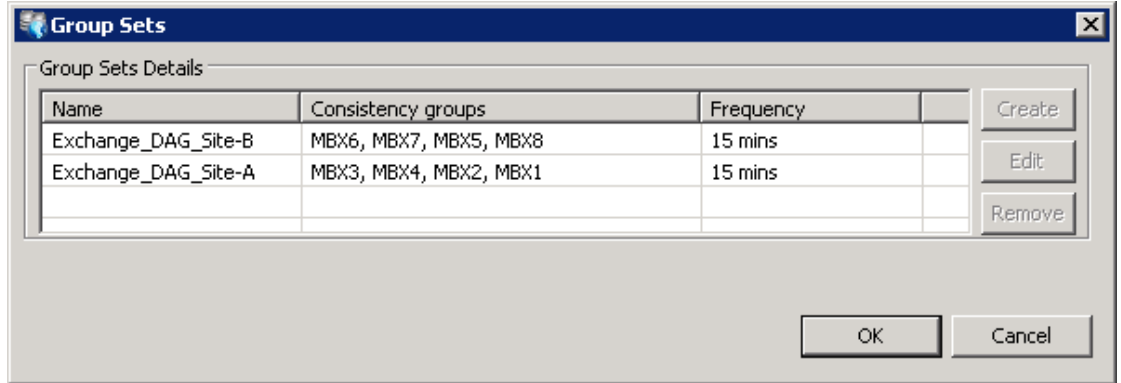


Figure 9. RecoverPoint group sets

Replication Manager configuration and integration with RecoverPoint

This solution integrates EMC RecoverPoint with EMC Replication Manager. Replication Manager's built-in application intelligence and automation enables you to leverage RecoverPoint's replication technology to schedule and create VSS copies of Exchange, both locally and remotely.

Replication Manager process overview

Replication Manager freezes the production Exchange database, requests a VSS copy, flushes Exchange Server's cache to disk, and creates a VSS bookmark for RecoverPoint and Replication Manager. Exchange is then thawed, all within the required VSS window of time. Once thawed, production resumes. Meanwhile the RecoverPoint VSS bookmark is captured in the local CDP journal as well as the remote CRR journal.

VSS API and KVSS utility

Microsoft Exchange Server supports the VSS API that enables RecoverPoint to integrate with Exchange. RecoverPoint provides a utility called KVSS that invokes the Microsoft VSS framework to prepare a VSS image. For full details on using KVSS utility for protecting Exchange, refer to *Replicating Microsoft Exchange Server with EMC RecoverPoint – Technical Notes*. In this solution, Replication Manager provides an additional level of automation by coordinating application-consistent snapshots.

Exchange tasks that can be automated by RM

Replication Manager (RM) can automate Exchange 2010 tasks such as:

- Create Exchange database application sets
- Create and manage multiple-site application-consistent replicas of Exchange databases
- Run a full Exchange job to create a VSS bookmark, run the Microsoft Exchange utility eseutil, and then truncate the Exchange logs
- Run an Exchange copy job that creates a VSS bookmark
- Start an on-demand mount and dismount of an Exchange database
- Restore Exchange databases from either local or remote bookmarks to the production Exchange server in seconds

It is very easy to use Replication Manager's GUI interface to automate the restore of Exchange to an application-consistent point in time. Replication Manager handles the restore locally by using the RecoverPoint CDP replica or, if necessary, RM can restore back to production by using the RecoverPoint CRR Exchange replica.

How this solution uses RM

In this solution, Replication Manager automates RecoverPoint application-consistent replicas and performs Exchange log truncation. To satisfy the solution requirements presented in [Table 1](#), the design supports 24 application-consistent (VSS) snapshots in a 24-hour cycle.

Note: In addition to maintaining hourly application-consistent replicas (snapshots) initiated by Replication Manager, the RecoverPoint journal maintains many more crash-consistent snapshots from which Exchange databases can be easily recovered. The number of crash-consistent snapshots in the journal depends on the size of the journal and the application change rate. Refer to [RecoverPoint journal usage under Exchange load](#), which describes testing of journal usage under an Exchange workload.

RM application set configuration

For this solution, an RM application set is created for each Exchange Mailbox server virtual machine. The application set includes all ten databases configured on the server. This design aligns with the RecoverPoint consistency group design, where each consistency group contains all of the Exchange databases and logs for a specific Mailbox server.

[Figure 10](#) shows the RM application set properties for Mailbox server MBX1.

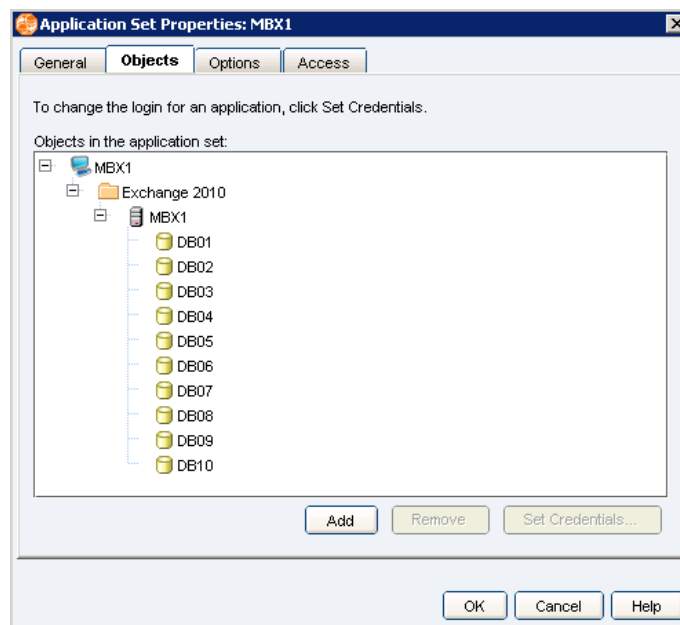


Figure 10. RM application set properties for one Exchange Mailbox server

Special RM environment variable required

For this solution, each RecoverPoint consistency group is configured with an additional replication set for the Mailbox server OS LUN that resides on the VMware VMFS datastore volume. This creates a potential conflict between Replication Manager and RecoverPoint, because the Microsoft VSS framework does not support VMFS data stores. To resolve this conflict, the following environment variable is created on each Exchange Mailbox server virtual machine with Replication Manager Agent installed:

```
ERM_RP_EXTRA_DEVS_ALLOWED=1
```

With this setting in place, RecoverPoint continues to provide crash-consistent replicas of all LUNs attached to the Exchange Mailbox server virtual machine, including the OS VMFS volume, and Replication Manager continues to provide application-consistent replicas of all Exchange RDM volumes.

About replication options

In an Exchange/VSS environment, Exchange administrators can select the Online Full or Online Copy replication option. For the Online Full option, Replication Manager replicates the databases and transaction logs, verifies the consistency of the databases and logs (optional for Exchange 2010), and then truncates the logs. For the Online Copy option, Replication Manager replicates the databases and transaction logs in the same way it does during an Online Full option; however, it does not truncate the logs.

Replication option configuration

In this solution, Replication Manager jobs are configured with two Replica Type options to protect the Exchange Server 2010 databases, Online Full and Online Copy. Replication Manager uses the VSS framework to initiate both of these replica types.

Online Full jobs are set up for daily backups, and Online Copy jobs are set up to run at hourly intervals during business hours to provide application-consistent recovery.

Figure 11 shows the Replication Manager Online Full job configuration settings to perform daily backup of Mailbox server MBX1. Advanced replication settings enable Replication Manager to verify database states before performing this job.

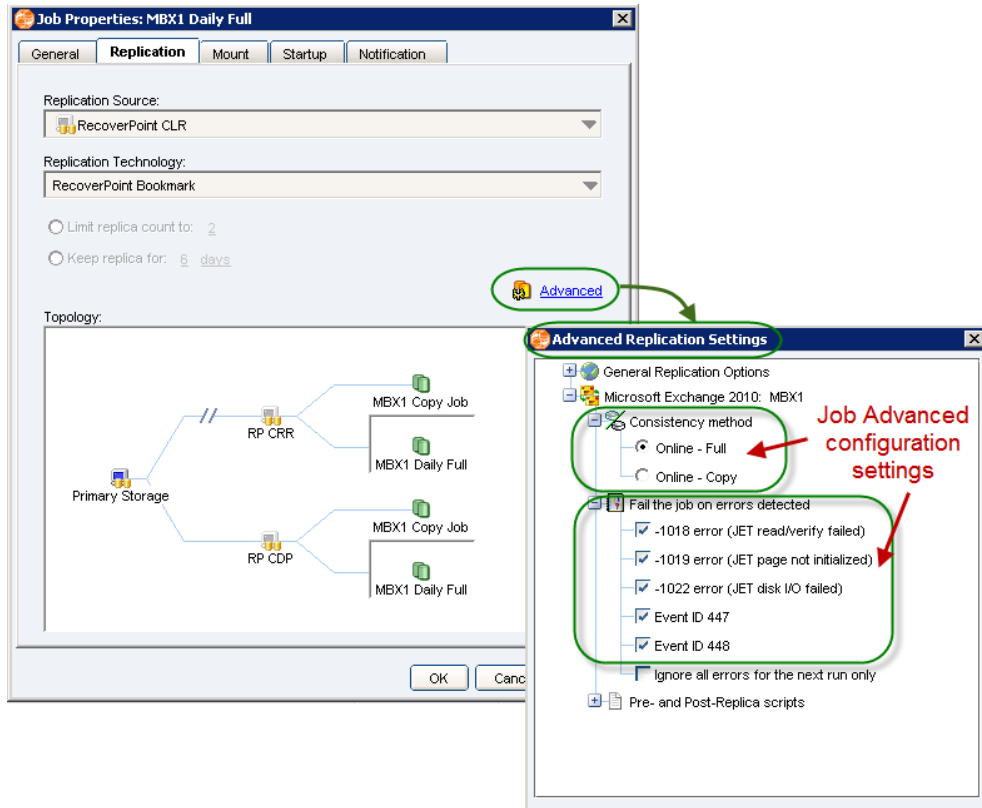


Figure 11. Replication Manager job advanced replication settings

vCenter Site Recovery Manager configuration

This section describes how vCenter Site Recovery Manager (SRM) version 5 is configured for automatic failover of the protected site to the recovery site.

Prerequisites

When designing a disaster recovery solution, it is essential that all software components are configured and working correctly. Careful planning and coordination of IT personnel responsible for these activities is also essential for a successful recovery plan execution for a planned or unplanned failover. This ensures the shortest RTO and compliance with SLAs.

SRM requirements for vSphere

SRM has several requirements for the vSphere configuration at each site:

- Each site must include a vCenter server containing at least one vSphere data center.
- The recovery site must support array-based replication with the protected site, and the recovery site must have hardware and network resources that can support the same virtual machines and workloads as the protected site.
- At least one virtual machine must be located on a data store that is replicated by RecoverPoint at the protected site.
- The protected and recovery sites must be connected by a reliable IP network. Storage arrays may have additional network requirements.
- The recovery site must have access to the same public and private networks as the protected site, although the recovery site does not necessarily require access to the same range of network addresses.
- VMware tools must be installed on all virtual machines.

SRM configuration steps

Installing and configuring VMware vCenter Site Recovery Manager (SRM) for automated site recovery using RecoverPoint involves these steps:

1. Install and configure SRM with the RecoverPoint SRA adapter.
2. Configure the connection between the protected and recovery sites.
3. Configure RecoverPoint array managers.
4. Configure resource inventory mappings.
5. Configure protection groups.
6. Create the recovery plan.
7. Customize virtual machine recovery options, including IP customization.

Most of these steps are executed from the SRM interface in vCenter, where intelligent wizards support quick and easy configuration. For full details of these steps, consult the *VMware vCenter Site Recovery Manager Administration Guide 5.0* (<http://www.vmware.com>).

Step 1 – Install and configure SRM with the RecoverPoint SRA adapter

Installing and configuring SRM involves the following tasks:

1. Configure SRM databases at both sites. These databases store the recovery plans, inventory information, and so on.
2. Install the SRM server at both sites.
3. Install the RecoverPoint Storage Replication Adapter on the SRM server at both sites. This adapter enables SRM and RecoverPoint integration.
4. Install the SRM client plug-in on one or more vSphere clients at both sites.

RecoverPoint Storage Replication Adapter is included in the SRM installation package which is available for download from the VMware website. Installation and configuration instructions are provided in the product release notes, which are included in the installation package.

Step 2 – Configure the connection between the protected and recovery sites

With the SRM database and SRM server installed at each site, you then need to configure the connection from the protected site to the recovery site. This is accomplished by specifying the remote SRM server IP details in the **Connect to Remote Site** wizard at the protected site.

Step 3 – Configure RecoverPoint array managers

For SRM to integrate with RecoverPoint, RecoverPoint array managers must be configured at both the protected and recovery sites. Use the SRM configuration wizard to configure properties for RecoverPoint SRA. The wizard discovers the replicated storage devices at the protected and recovery sites and identifies the VMFS data stores that they support. When finished, the wizard presents a list of replicated data store groups. [Figure 12](#) shows properties for RecoverPoint SRA manager at the protected site. The connection parameters for RecoverPoint at the protected site are specified. The same parameters must be specified for SRM at the recovery site.

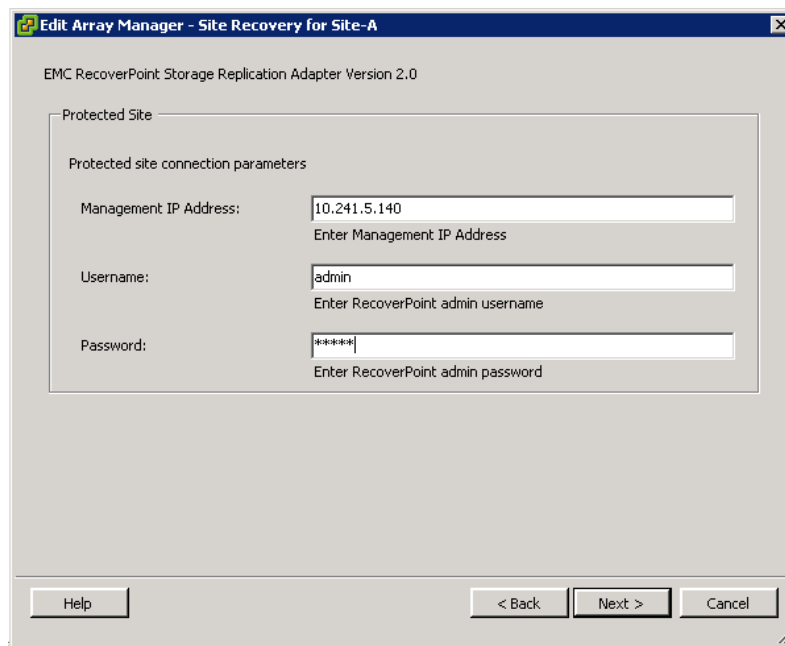


Figure 12. SRM RecoverPoint Array Manager configuration

After an array manager is configured for each site, array pairs must be enabled for use with SRM. This can be done from either site.



Figure 13. Enabled RecoverPoint array manager

After the array managers are enabled, the “local” and “remote” devices for each enabled array pair are displayed.

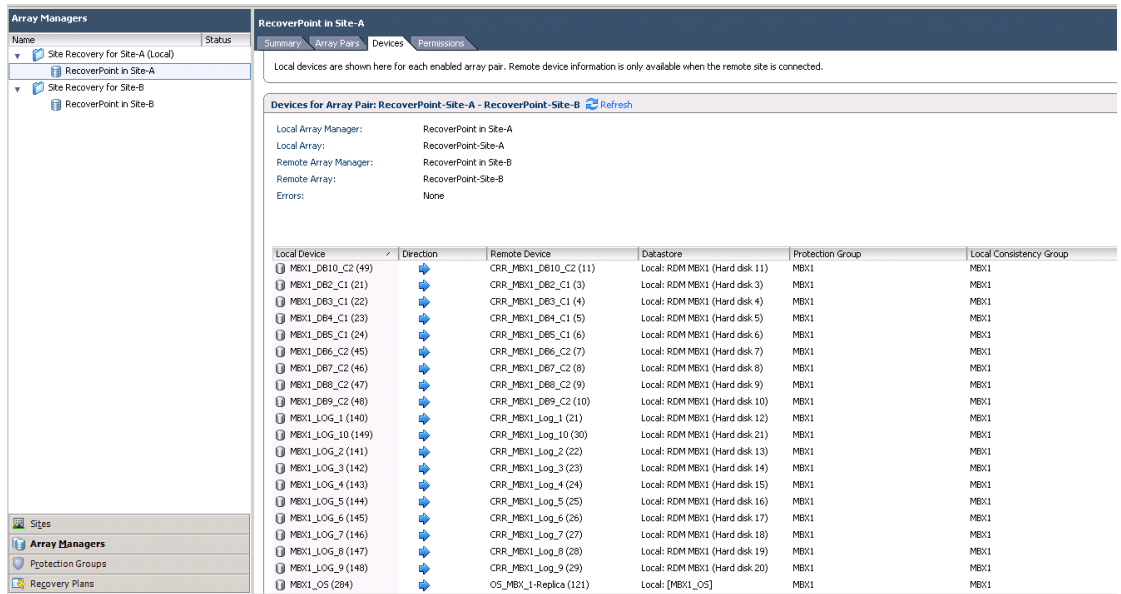


Figure 14. Replicated devices shown in RecoverPoint SRM array manager

Step 4 – Configure resource inventory mappings

Once SRM and array managers are configured, mappings between site resources must be defined.

Step 5 – Configure protection groups

A protection group is a collection of virtual machines that all use the same datastore group (the same set of replicated LUNs) and that all fail over together. To create a protection group, select the data store groups to protect, and specify a data store group at the recovery site where SRM can create placeholders for members of the protection group. Use the **Create Protection Group** wizard to do this. The wizard automatically detects the data stores currently protected by RecoverPoint and allows you to select the ones to include in the protection group.

Eight protection groups were created for this solution, one for each Exchange Mailbox server virtual machine. [Figure 15](#) shows the protection groups created for the solution.

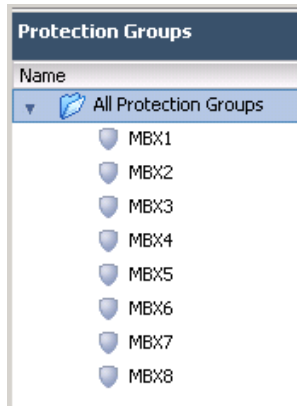


Figure 15. SRM protection groups created for this solution

After a protection group is created, shadow virtual machines are automatically created in the recovery site vCenter inventory. These act as placeholders for the virtual machines that can be failed over with SRM. The shadow virtual machines cannot be started independently from SRM and are removed if their protection group is deleted. [Figure 16](#) illustrates this.

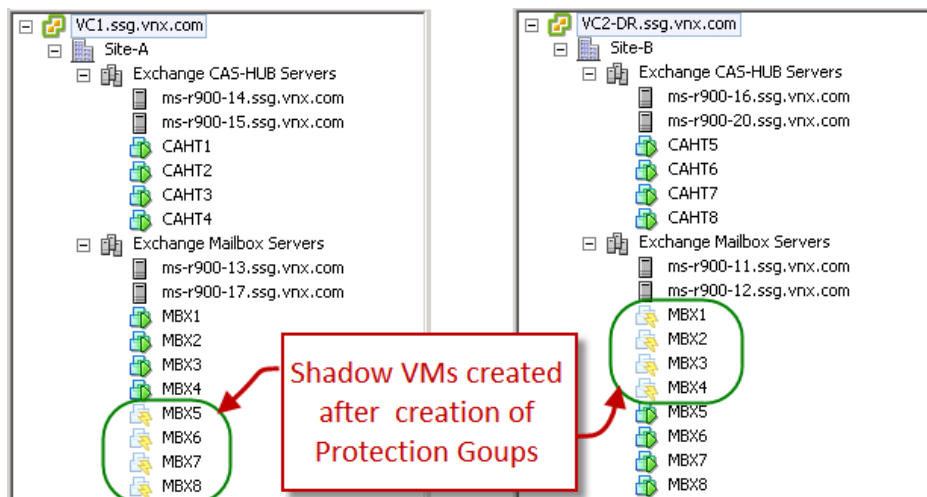


Figure 16. Shadow virtual machines are created following the creation of protection groups

Step 6 – Create the recovery plan

After protection groups are created, the next step is to create recovery plans. You can create multiple recovery plans to accommodate many different recovery scenarios. As shown in Figure 17, ten recovery plans are created for this solution, two to fail over the entire environment from each site, and the other eight to fail over individual Exchange Mailbox server virtual machines independently. The figure also shows the SRM option for creating a recovery plan that contains individual protection groups from one of the sites.

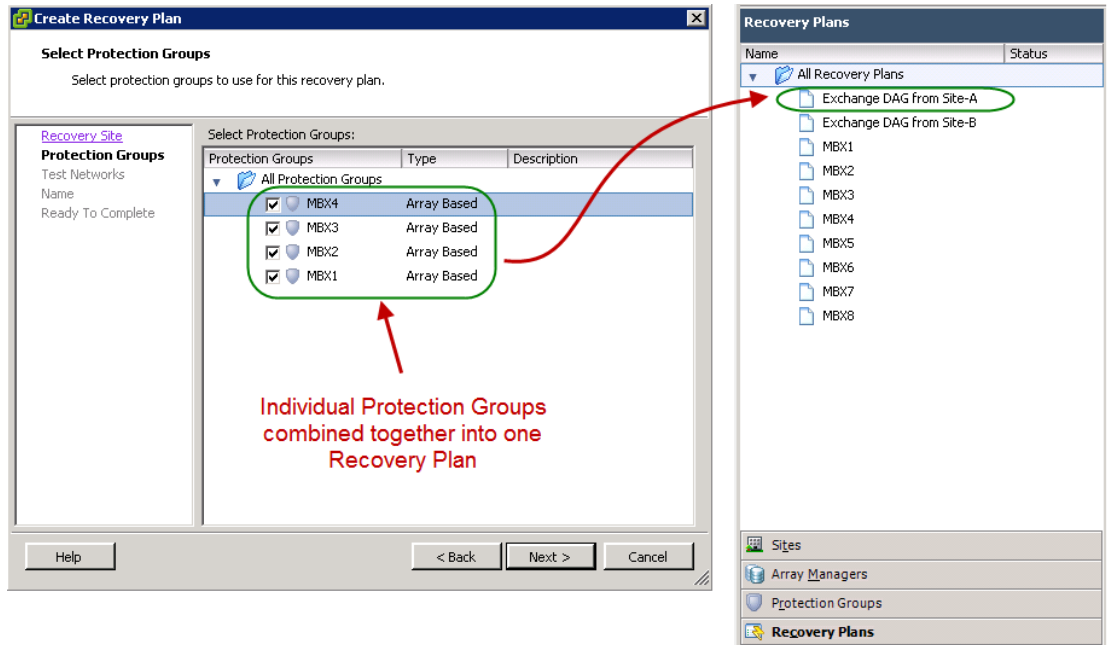


Figure 17. SRM recovery plan configuration options

Step 7 – Customize virtual machine recovery options, including IP customization

The recovery options for individual virtual machines can be customized to satisfy specific requirements. The available customization options constitute a powerful tool for ensuring that complex recovery plans can be implemented with ease. For example, if there are multiple virtual machines in a protection group, their startup sequence can be customized by assigning them different recovery priorities. Options include the ability to change virtual machine IP addresses so that they point to specific subnets at the recovery site. Options can also be set to run specific pre- or post-processing scripts before or after the execution of a recovery plan.

To ensure that the entire Exchange DAG at the protected site successfully fails over to the recovery site, the IP address of each Mailbox server virtual machine at the recovery site is specified. The original IP addresses of the protected virtual machines are also specified to ensure successful and seamless failback. Figure 18 shows the virtual machine customization wizard for one Exchange Mailbox server virtual machine where the IP addresses and DNS information are specified.

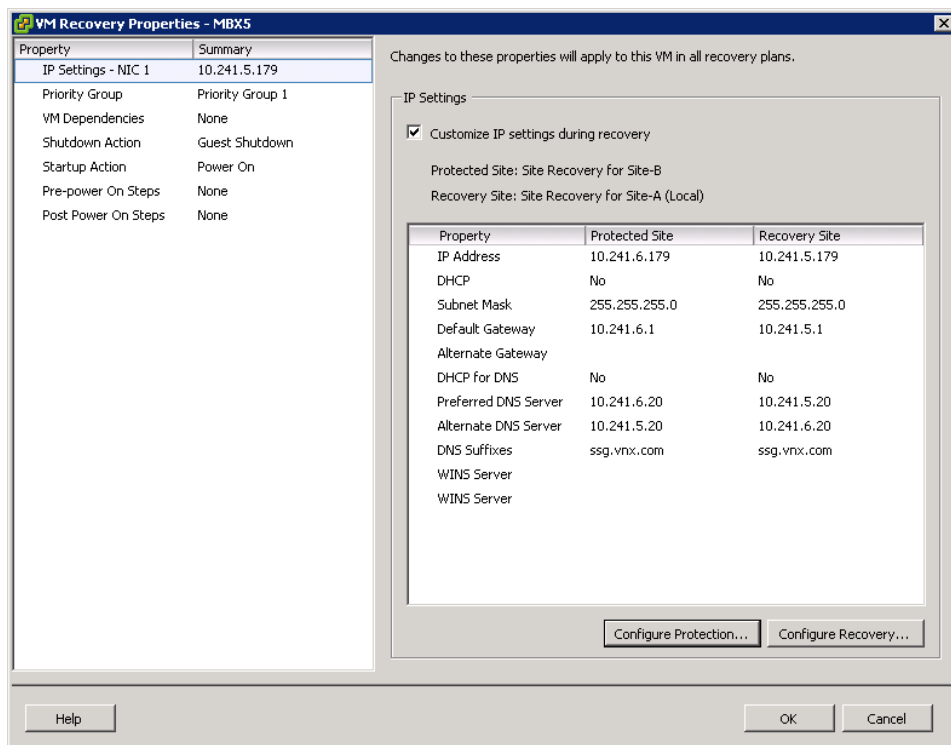


Figure 18. Exchange virtual machine IP customization properties for SRM recovery plans

Exchange configuration to support data center failover

Microsoft Exchange Server is tightly coupled with the domain controller and uses the Active Directory service to provide authorization and authentication services for the Exchange servers and users. Disaster recovery for Exchange Server requires that the Active Directory service is available at the both the protected site and the recovery site.

EMC recommends ensuring that a domain controller running Active Directory is configured at the remote site before beginning RecoverPoint replication. This facilitates better communication between the Active Directory server and the Exchange server during the Exchange installation process. In addition, in the event of a total source-side failure during remote replication, this eliminates the potential service delay that can result from the need to construct a new Active Directory server at the target site.

To achieve a successful failover, it is very important to ensure that the Exchange 2010 DAGs and all underlying cluster and Active Directory components are configured correctly. DAG property values are stored in both the Active Directory and the cluster database. However, some properties are stored only in the cluster database.

In this solution, an entire DAG (with underlying Windows Failover Cluster components) is moved from one active production site to another active production site where Exchange services are currently configured for users serviced by that site.

For complete instructions on configuring Exchange 2010 DAGs and all underlying cluster and Active Directory components, refer to the articles *Managing High Availability and Site Resilience* at <http://technet.microsoft.com/en-us/library/dd638215.aspx> and *Configure Database Availability Group Properties* at <http://technet.microsoft.com/en-us/library/dd297985.aspx>.

Analysis and reporting with EMC Data Protection Advisor (DPA)

EMC Data Protection Advisor (DPA) is a sophisticated reporting and analytics platform that provides customers with full visibility into the effectiveness of their data protection strategy. DPA does this by monitoring all of the technologies that a customer uses to protect data, including backup software, storage arrays, file servers, and tape libraries.

DPA's sophisticated reporting engine provides highly customizable reports to highlight problems within the environment and enables customers to perform capacity management, service level reporting, chargeback, change management, and troubleshooting.

DPA's Predictive Analysis Engine provides customers with early warning of problems that might occur and generates alerts that enable customers to resolve problems sooner to reduce business impact.

DPA is designed to help users identify and fix performance and capacity management (PCM) issues within large enterprises. DPA consists of a number of loosely-coupled processes and data stores for holding configuration and gathered data. In addition, DPA has a number of interfaces to import and export information, which allow for integration with other aspects of a large enterprise and to provide a significant part of an overall distributed systems management strategy.

DPA integration with RecoverPoint

DPA gathers configuration, status, and performance information from a RecoverPoint system in your environment.

The Collector component is a process that gathers data from an application or device. A Collector is automatically installed on the Data Protection Advisor server when you install the Server component. You can install the Collector component on other devices in the environment to provide enhanced data collection.

Data Protection Advisor stores three types of data:

- Configuration data, which includes information on the configuration of Data Protection Advisor itself—This information is maintained by the controller process and is generally static and small in size.
- Gathered data, which is the data gathered by Collectors—The size of this data can become quite large due to the dynamic nature of the collected data. As a result, it is maintained separately from the configuration data.
- Recoverability analysis data, which is gathered by the recoverability engine—Recoverability analysis is only available for a DPA Server installed on a Windows or Solaris platform. It is not available for DPA servers installed on Red Hat or SUSE Enterprise Linux.

The default names for the data stores are config, datamine, and illuminator.

Monitoring of RecoverPoint requires additional licensing beyond what is included in the standard DPA license. For more information on how to obtain licenses for RecoverPoint reporting, contact an EMC representative.

RecoverPoint components are displayed in the navigation tree in a hierarchy under **Storage > Replication**, with the managing host at the top level. RecoverPoint appliances (RPAs) and splitters are divided by site, if there is more than one site in your RecoverPoint system. Consistency group copies are mapped in the hierarchy to their consistency groups. To view reports for RecoverPoint, right-click the host, RPA, or RPA component node in the navigation tree.

DPA/RecoverPoint sample reports

Figure 19 through Figure 25 present some sample DPA reports based on the monitoring of a RecoverPoint system.

Figure 19 shows the RecoverPoint consistency group journal lag for an RPA over a 30-day period.




Figure 19. DPA/RecoverPoint sample report: RecoverPoint consistency group journal lag

Figure 20 shows the SAN throughput and the WAN throughput at both sites for a single nightly backup.

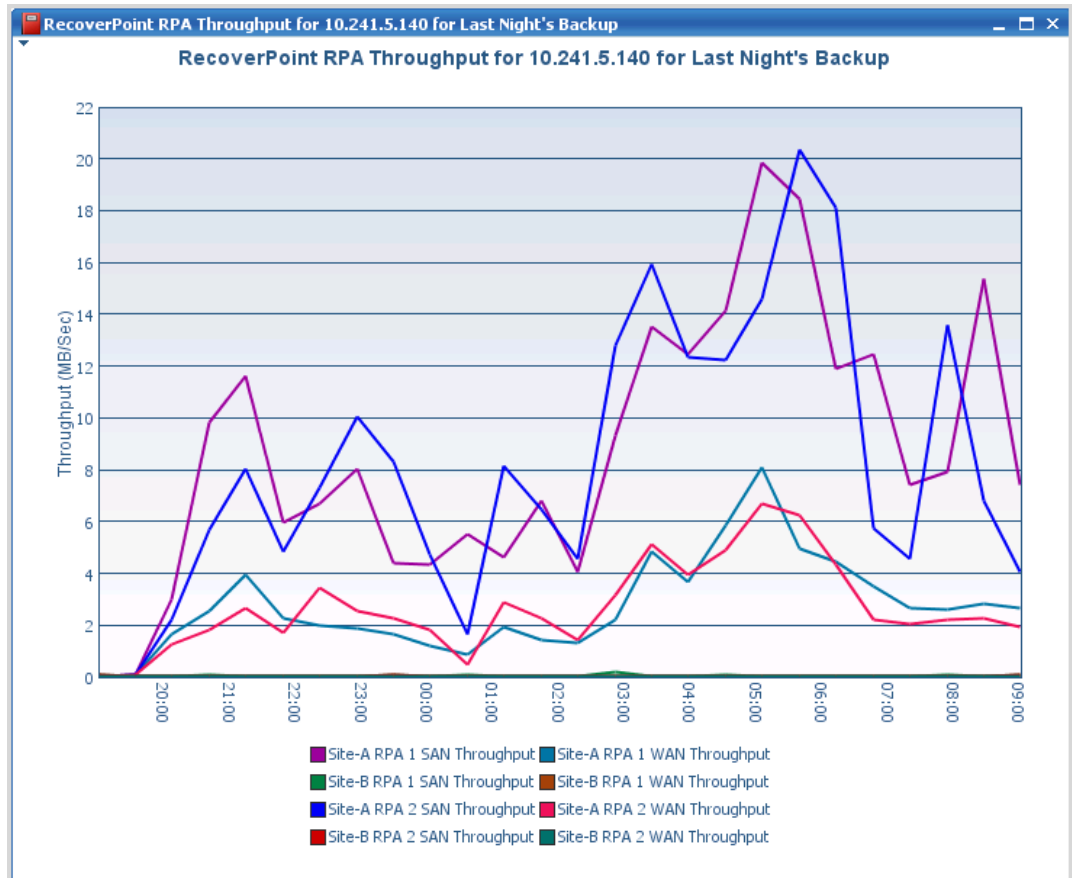


Figure 20. DPA/RecoverPoint sample report: RecoverPoint RPA throughput

Figure 21 shows the RecoverPoint consistency group copy status, including Group Name, Storage Access, Journal Usage, and so on.

RecoverPoint Consistency Group Copy Status for 10.241.5.141 at 10/20/11 02:06														
Server	Consistency Group	Name	Enabled	Active RPA	Data Transfer	Journal State	Storage Access	Current Protection Window (week)	Predicted Protection Window (week)	Link Mode	Journal Mode	Journal Usage (GB)	Latest Journal Image	Journal Lag (GB)
10.241.5.141	MBK1	MBK1-CDP	<input checked="" type="checkbox"/>	RPA 1	Active	Distributing	No Access	2	5	Asynchronous	Normal	255.0	10/20/11 01:50	0.0
10.241.5.141	MBK1	MBK1_CRR	<input checked="" type="checkbox"/>	RPA 1	Active	Distributing	No Access	2	11	Asynchronous	Normal	240.0	10/20/11 01:50	0.0
10.241.5.141	MBK1	MBK1_prod	<input checked="" type="checkbox"/>	RPA 1			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK2	MBK2-CDP	<input checked="" type="checkbox"/>	RPA 1	Active	Distributing	No Access	2	3	Asynchronous	Normal	131.0	10/20/11 01:50	0.0
10.241.5.141	MBK2	MBK2-CRR	<input checked="" type="checkbox"/>	RPA 1	Active		Logged Access	2	3	Asynchronous	Normal	127.0	10/20/11 01:50	5.6
10.241.5.141	MBK2	MBK2-Prod	<input checked="" type="checkbox"/>	RPA 1			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK3	MBK3-CDP	<input checked="" type="checkbox"/>	RPA 2	Active	Distributing	No Access	1	5	Asynchronous	Normal	346.0	10/20/11 01:12	0.0
10.241.5.141	MBK3	MBK3-CRR	<input checked="" type="checkbox"/>	RPA 2	Active		Logged Access	1	1	Asynchronous	Normal	67.0	10/20/11 01:12	7.2
10.241.5.141	MBK3	MBK3-Prod	<input checked="" type="checkbox"/>	RPA 2			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK4	MBK4-CDP	<input checked="" type="checkbox"/>	RPA 2	Active	Distributing	No Access	2	3	Asynchronous	Normal	133.0	10/20/11 01:12	0.0
10.241.5.141	MBK4	MBK4-CRR	<input checked="" type="checkbox"/>	RPA 2	Active		Logged Access	2	2	Asynchronous	Normal	127.0	10/20/11 01:12	5.4
10.241.5.141	MBK4	MBK4-Prod	<input checked="" type="checkbox"/>	RPA 2			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK5	MBK5-CRR	<input checked="" type="checkbox"/>	RPA 1	Active	Distributing	No Access	0	75	Asynchronous	Normal	1.0	10/20/11 01:26	0.0
10.241.5.141	MBK5	MBK5-Prod	<input checked="" type="checkbox"/>	RPA 1			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK6	MBK6-CRR	<input checked="" type="checkbox"/>	RPA 1	Paused	Distributing	No Access			Asynchronous	Normal	0.0		0.0
10.241.5.141	MBK6	MBK6-Prod	<input checked="" type="checkbox"/>	RPA 1			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK7	MBK7-CRR	<input checked="" type="checkbox"/>	RPA 2	Active	Distributing	No Access	0	51	Asynchronous	Normal	1.0	10/20/11 01:57	0.0
10.241.5.141	MBK7	MBK7-Prod	<input checked="" type="checkbox"/>	RPA 2			Direct Access (marking data)			Asynchronous				
10.241.5.141	MBK8	MBK8-CRR	<input checked="" type="checkbox"/>	RPA 2	Paused	Distributing	No Access			Asynchronous	Normal	0.0		0.0
10.241.5.141	MBK8	MBK8-Prod	<input checked="" type="checkbox"/>	RPA 2			Direct Access (marking data)			Asynchronous				

Figure 21. DPA/RecoverPoint sample report: RecoverPoint consistency group copy status

Figure 22 shows an overview of RecoverPoint events, including warnings, errors, and informational messages.

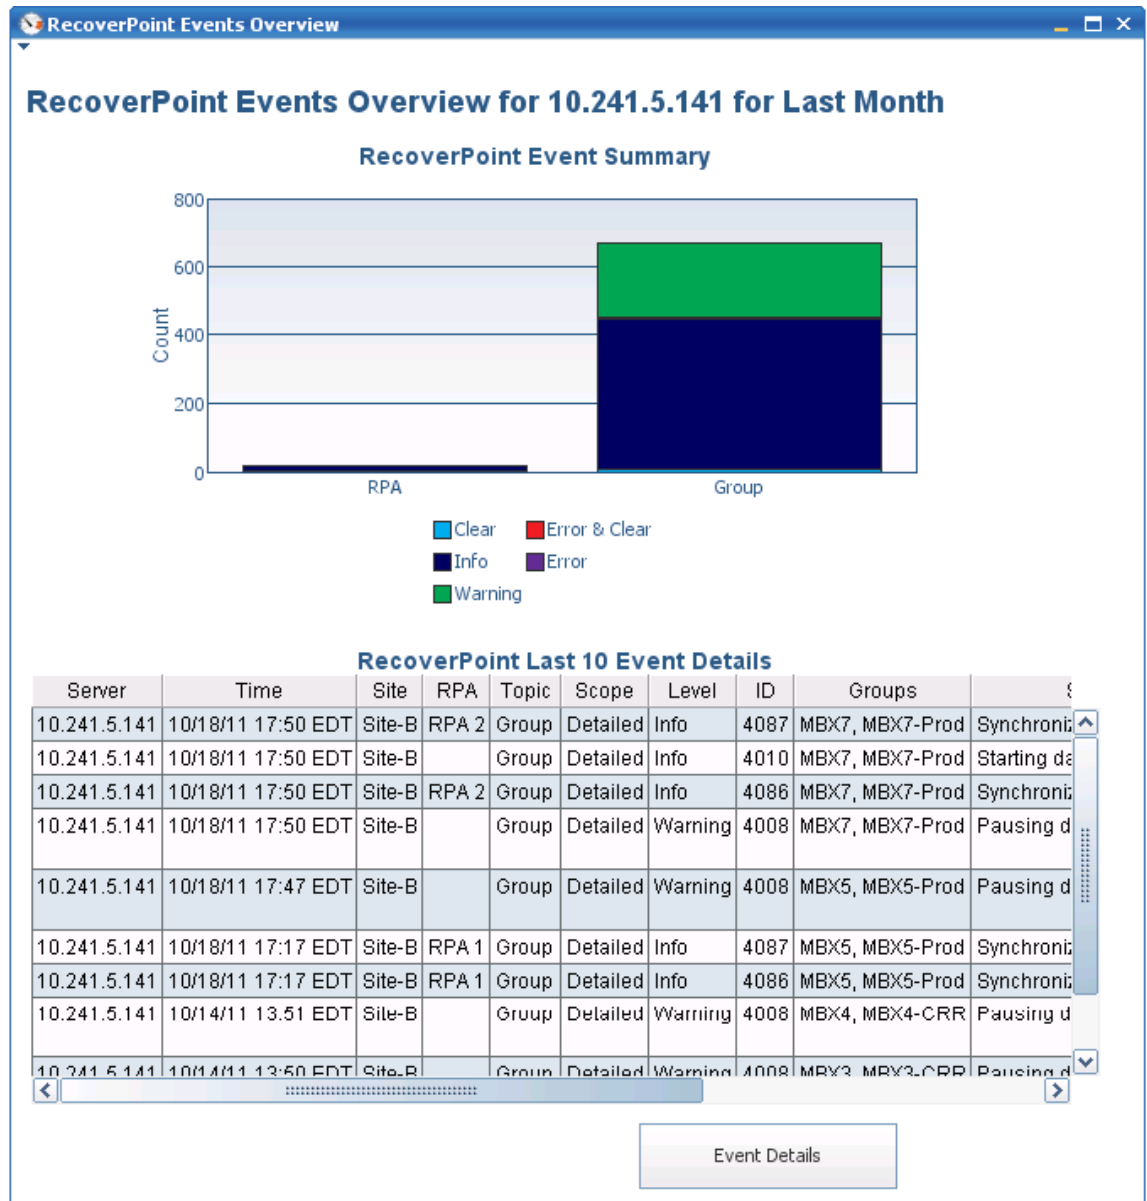


Figure 22. DPA/RecoverPoint sample report: RecoverPoint events overview

DPA integration with VMware vSphere

DPA can be integrated with VMware vSphere with a VMware plug-in. This plug-in allows you to view reports from a vSphere client that provides information on virtualization hosts and virtual machines.

Registering the plug-in

After you configure a virtualization node in DPA, register the VMware plug-in from the DPA GUI to view reports from the infrastructure client.

Viewing the plug-in reports

After registering the plug-in, tabs are available from the vSphere client registered with the plug-in. The tabs display reports that differ according to whether a virtualization host or virtual machine is selected in the VMware hierarchy.

DPA/vSphere sample report

Figure 23 presents a sample report generated by DPA to monitor resource utilization in a vSphere environment.




Figure 23. DPA/vSphere sample report: Resource utilization overview

DPA integration with VNX/CLARiiON

DPA can integrate with EMC Symmetrix VNX series and EMC CLARiiON storage arrays. VNX and CLARiiON storage arrays replicated with EMC RecoverPoint require additional configuration to enable complete recoverability analysis. For configuration instructions, refer to the DPA Help system.

DPA/VNX sample reports

Figure 24 and Figure 25 show sample DPA reports based on the monitoring of an EMC VNX series storage array.

Figure 24 shows an overview of changes made to the storage, for example, changes made to LUN sizes.

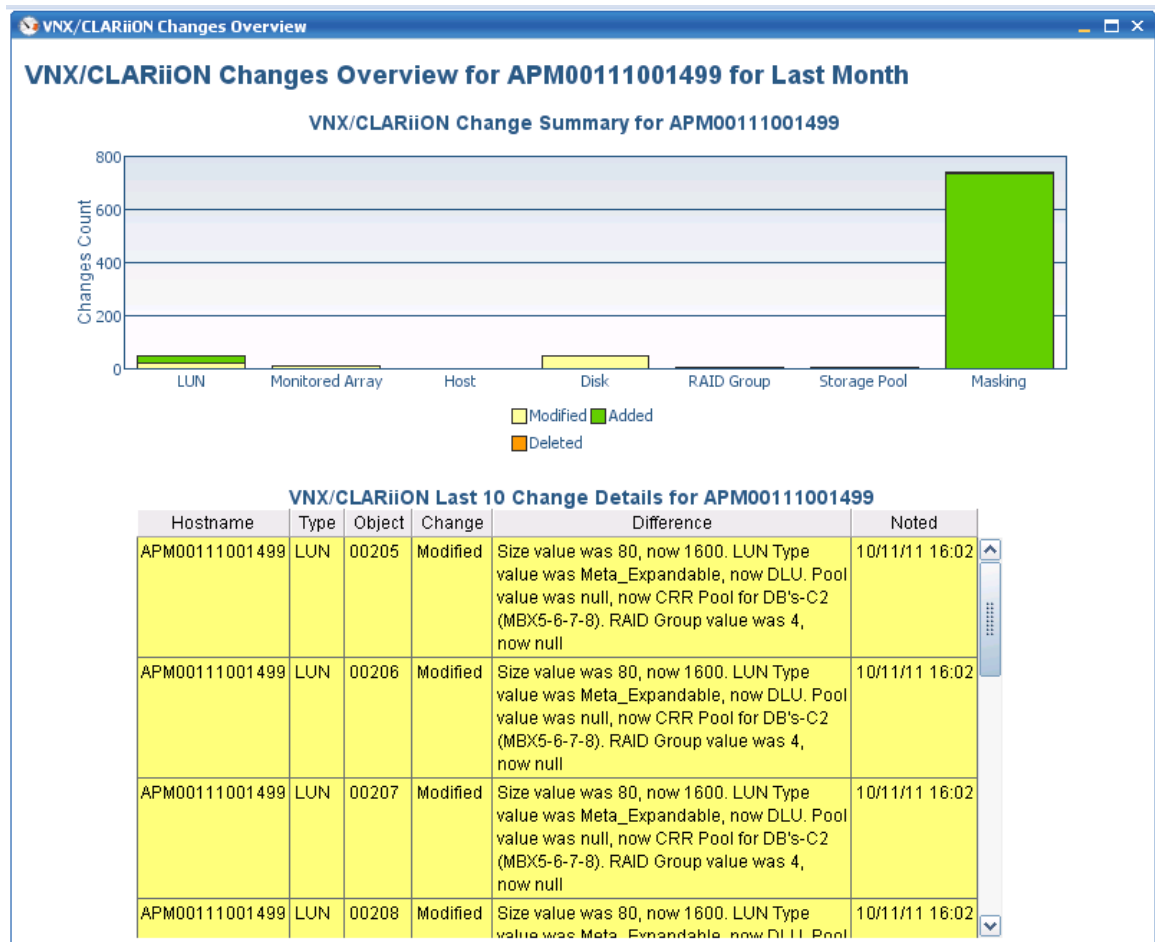


Figure 24. DPA/VNX sample report: VNX change summary report

Figure 25 shows an overview of overall storage capacity, raw capacity, and user capacity.




Figure 25. DPA/VNX sample report: VNX capacity overview

DPA integration with Microsoft Exchange Server

Microsoft Exchange Server can be monitored for recoverability from a collector installed on the same machine as Exchange or from a collector installed remotely.

Note: Exchange can be monitored for recoverability analysis and for system information only from the Exchange Server host.

The account used to connect DPA to the Exchange server must be a domain user with Exchange read-only administrator rights and local administrator rights. DPA supports recoverability analysis for only one Exchange information store in a cluster.

Solution validation and test results

The following sections provide details and results for tests performed as part of this solution validation. The test results include:

- RecoverPoint journal usage under load
- SRM recovery plan testing
- Exchange failover with SRM
- Exchange failback with SRM
- Exchange database recovery with Replication Manager

RecoverPoint journal usage under Exchange load

One of the most important questions that need to be answered during a RecoverPoint deployment for Exchange 2010 is how much journal space is required to protect the Exchange data. If you allocate too little journal space, you cannot meet your RPO requirements and might lose the bookmarks (point-in-time snapshots), because new bookmarks are likely to overwrite the current ones. If you allocate too much journal space, you might waste valuable storage resources. As part of testing, a 24-hour Microsoft Load Generator (LoadGen) test was run to measure the journal usage under load.

About Microsoft LoadGen

LoadGen does not represent real-world workloads. LoadGen is a simulation tool to measure the effects of MAPI, OWA, ActiveSync, IMAP, POP, and SMTP clients on Exchange 2010 and 2007 servers. The tool allows you to test how a server running Exchange responds to email loads. To simulate the delivery of messaging requests, LoadGen tests run on client computers. LoadGen is a useful tool for administrators when sizing servers and validating a deployment plan. Specifically, LoadGen helps you determine if each server can handle the load it is intended to carry. Another use for LoadGen is to help validate an overall solution. For additional information about using LoadGen, visit Microsoft TechNet at <http://technet.microsoft.com/en-us/library/dd335108.aspx>.

RecoverPoint journal space usage during testing

Table 6 shows journal space usage in four-hour increments observed during a 24-hour LoadGen test. 125 GB (25 percent) of the 500 GB journal was reserved space, thus only about 375 GB of capacity was available for use by the test. During the test, especially during the first eight to 10 hours of busy “daytime” activity generated by LoadGen, almost 46 percent of the journal was utilized. By the end of the test, the journal had nearly reached its maximum usable capacity. The test results showed that the journals could provide only a 24-hour maximum protection window because the change rate generated by LoadGen was almost three times more than the 5 Mb/s change rate originally used to size the journals. Thus, the journal size established for testing was not sufficient to provide the three days of rollback protection originally planned. To provide three days of protection, 1.5 TB journals need to be deployed for an Exchange environment where the change rate is 15 Mb/s.

The numbers presented in [Table 6](#) are intended to serve only as an approximate guideline to size RecoverPoint journals for Exchange 2010. It is recommended to size journals based on actual change rates in your production environment. The results presented in the table are based on the following:

- LoadGen workload for 5,000 users per server with 150 sent/received messages per user
- Hourly application-consistent (VSS) bookmarks for a RecoverPoint consistency group with all of the LUNs for a single Mailbox server with 10 databases
- Automatic crash-consistent (VSS) bookmarks for a RecoverPoint consistency group with all of the LUNs for a single Mailbox server with 10 databases
- Nightly scheduled backup job initiated by Replication Manager

Percentages are based on a 375 GB usable journal space (500 GB minus 25 percent reserved by RecoverPoint) provisioned for each consistency group and 13.6 TB of replicated Exchange data (80 percent of 17 TB provisioned storage).

Table 6. RecoverPoint journal usage with Exchange Server 2010 under load

Test timeline	Journal usage (GB)	Journal usage per user (MB)	Journal usage (percentage)
Start	0.20 GB	----	0.05%
4 hrs. (day time)	113 GB	23 MB	30.1%
8 hrs. (day time)	174 GB	36 MB	46.4%
12 hrs. (night time)	184 GB	38 MB	49.0%
16 hrs. (night time)	228 GB	47 MB	60.8%
20 hrs. (day time)	279 GB	57 MB	74.4%
24 hrs. (day time)	345 GB	71 MB	92.0%
Note: Performance is measured based on 375 GB of usable journal space and a 15 Mb/s change rate generated by the LoadGen tool.			

SRM recovery plan testing

Having a disaster recovery plan in place is the first step toward ensuring that you are prepared for the worst-case scenario. However, the only way to know if your disaster recovery plan works is to test it. The ability of SRM to execute recovery plans in test mode allows you to do this.

- Recovery plans can be tested at any time without interrupting replication or the business's RPO.
- Accurate time measurement of recovery plan testing enables administrators to predict how much time the recovery of the virtual environment is likely to take.
- Administrators can quickly access a replica of the entire protected environment for testing purposes.
- Automatically generated summary reports can be used to demonstrate compliance with regulatory requirements.

Recovery plan testing

All recovery plans were successfully tested as part of the validation of this solution. To demonstrate the failover process with SRM and RecoverPoint, both failover and failback of each site in its entirety was performed.

When testing a recovery plan, SRM creates an isolated test network at the recovery site and enables image access for the RecoverPoint consistency group. SRM also starts the virtual machines replicated by RecoverPoint and performs any reconfiguration necessary to access the disks on the array at the recovery site.

Note: During testing, some virtual machines could not power on automatically in vCenter as part of the SRM recovery plan. As a workaround, the virtual machines were powered on by connecting the vSphere client directly to the vSphere server. EMC has opened a support request with VMware (SR 11105524210) and VMware is currently investigating the issue. EMC recommends checking the VMware support site periodically for hotfixes related to this issue.

For the MBX1 protection group failover test, the status view in the RecoverPoint management application shows that RecoverPoint image access is enabled, as shown in [Figure 26](#) (on the next page). This means that the vSphere host and MBX1 Exchange Mailbox Server Role virtual machine have access to the RecoverPoint point-in-time bookmark initiated by SRM. The entire test took just over 17 minutes including IP address changes. The recovery test for all four Mailbox servers took just 25 minutes.

[Table 7](#) provides additional details about results from the multiple recovery plan tests.

Table 7. SRM Recovery Plan testing results

SRM Recovery Plan	Time elapsed (minutes)	RTO Met?
Single Mailbox server	17	Yes
Four Mailbox servers in a DAG	25	Yes

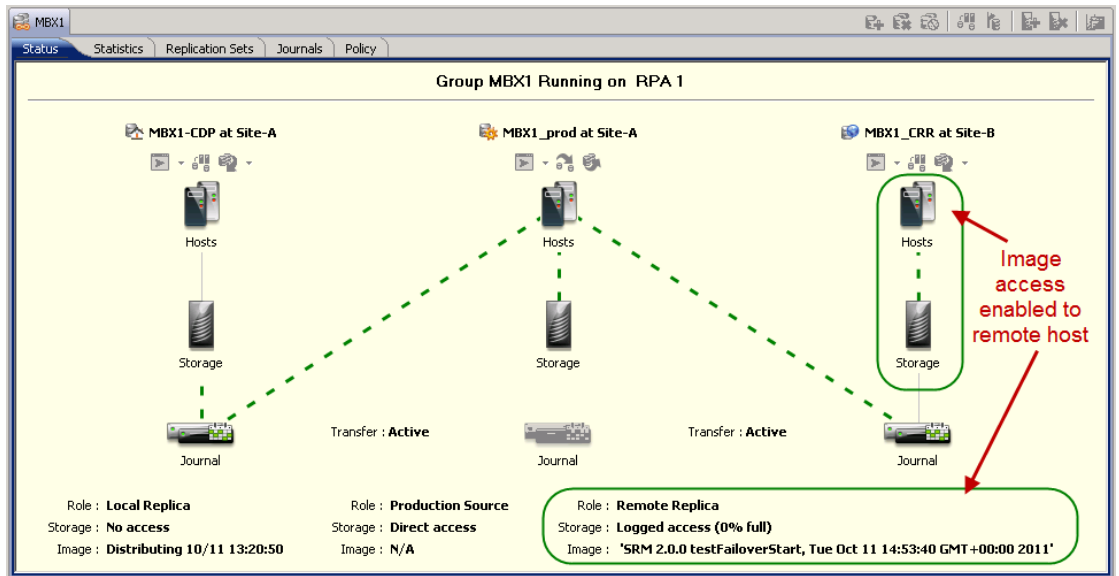


Figure 26. RecoverPoint management application showing image access enabled on the host at the recovery site

Following the recovery test, SRM generated a Recovery Plan History report. This report can be used to demonstrate accurate timing of recovery plans and recovery reliability for auditing, compliance, system administration, and other business purposes. Figure 27 shows an extract from the Recovery Plan History report.

Recovery Plan History Report
 VMware Site Recovery Manager 5.0

Plan Summary				
Name:	MBX1 Recovery			
Description:				
Protected Site:	Site-A			
Recovery Site:	Site-B			
Run Summary				
Operation:	Test			
Storage Options:	Synchronize storage when plan runs			
Started By:	Administrator			
Start Time:	2011-11-08 06:51:09 (UTC 0)			
End Time:	2011-11-08 07:08:40 (UTC 0)			
Elapsed Time:	00:17:31			
Result:	Warnings			
Errors:	0			
Warnings:	23			
Recovery Step	Result	Step Started	Step Completed	Execution Time
1. Synchronize Storage	Success	2011-11-08 06:51:09 (UTC 0)	2011-11-08 06:51:45 (UTC 0)	
1.1. Protection Group MBX1	Success	2011-11-08 06:51:09 (UTC 0)	2011-11-08 06:51:45 (UTC 0)	
Consistency Group 'MBX1':				

Figure 27. SRM Recovery Plan History Report sample

Exchange failover with SRM

This section describes a failover process for Exchange Mailbox servers deployed in a DAG at Site A to the active recovery Site B, where another production Exchange DAG is already running. (The section [vCenter Site Recovery Manager configuration](#) in this white paper describes how this recovery plan was created and tested.) The successful execution of an SRM recovery plan instills confidence that recovery from a real-world disaster is likely to be successful. EMC strongly recommends periodic testing of recovery plans to ensure that they operate as expected.

For this solution, remote site recovery occurs as the result of the integration of RecoverPoint Continuous Remote Replication (CRR) with VMware vCenter SRM. SRM automates the recovery process so that executing failover becomes as simple as pressing a single button. There is no need to interact with the RecoverPoint console. An SRM plug-in, installed on the vCenter instance at the recovery site, is all that is required to trigger the execution of a recovery plan.

In this solution, Microsoft Exchange is designed and configured for long distance disaster recovery; therefore, the subnet at Site A does not extend to Site B. The Exchange configuration includes a single Active Directory forest and two different AD/DNS sites. The Active Directory configuration automatically replicates between the sites during normal operation. When a disaster occurs at one of the sites, the replication is interrupted. Site failover then occurs as a result of the execution of a VMware SRM recovery plan for the entire Exchange DAG.

Failing over all protection groups in a protected Exchange DAG (at Site A) to its recovery site (Site B) involves activating the “Recovery” option on the SRM server at the recovery site, as shown in [Figure 28](#).

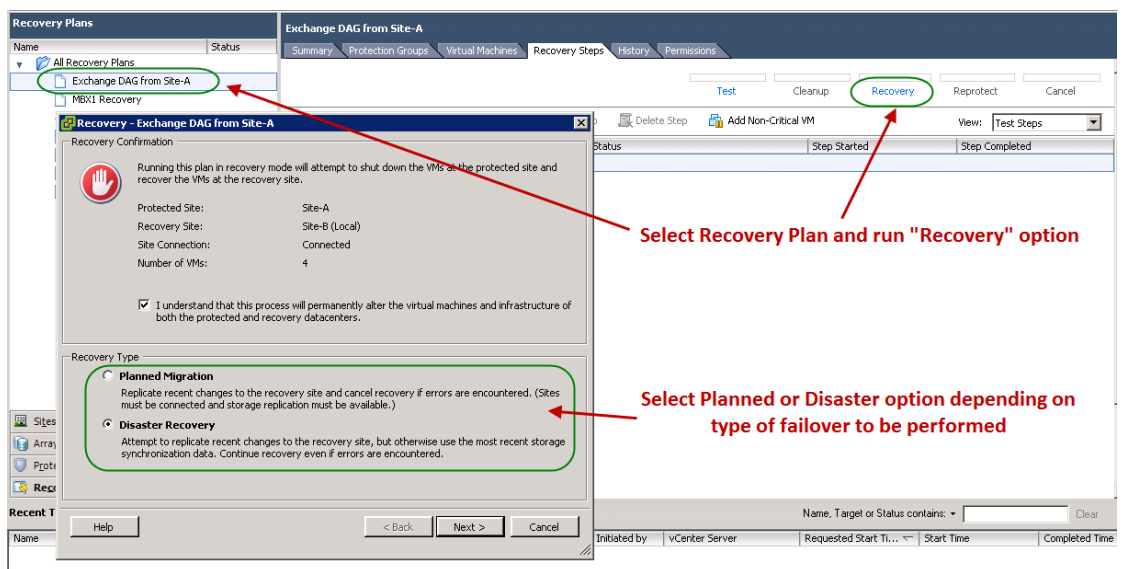


Figure 28. Running site failover to recover Exchange DAG from Site-A

Changing the DAG IP address

It is essential to ensure that the Exchange DAG IP address is updated in the DNS before the first Mailbox server member of the DAG is powered on at the recovery site. If the DNS is not updated, the underlying Windows Failover Cluster DAG resource name is not brought online. The DAG node with the “Cluster Group” acts as the Primary Active Manager. The Primary Active Manager is the DAG node responsible for choosing which databases are activated in a failover. For more information on the Primary Active Manager, refer to the Microsoft TechNet article at <http://technet.microsoft.com/en-us/library/dd776123.aspx>.

To change the DAG IP address, run the following Exchange PowerShell command on the SRM server.

```
Set-DatabaseAvailabilityGroup -Identity <DAG Name> -  
DatabaseAvailabilityGroupIPAddresses <new DAG ip-address at  
recovery site>
```

<DAG Name> is the name of the Exchange DAG from Site A (in this solution it is <DAG-Site-A>), and <ip-address> is the address provided for the recovery site subnet that is different from the protected site subnet.

Note: Exchange Management Tools must be installed on the SRM server in order to run Exchange Management Shell cmdlets.

To simplify the DAG IP address change process during SRM failover, an additional step that calls an Exchange PowerShell script can be added to a recovery plan. Figure 29 shows a step that issues a PowerShell command to call a batch file to change the DAG IP address.

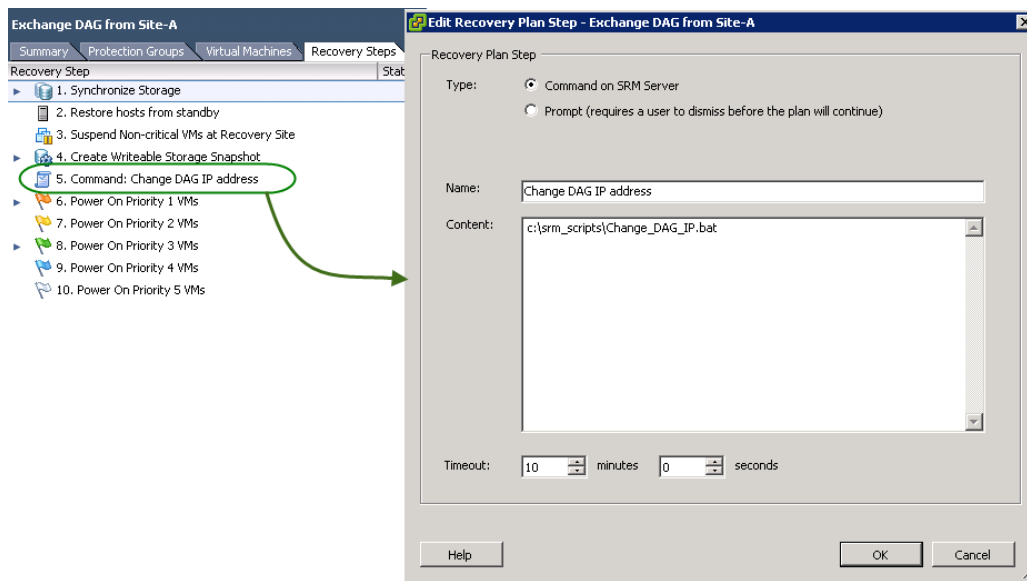


Figure 29. Script added to a recovery plan to change the DAG IP address

Validating client access

When all virtual machines are running, modify the CAS array entry on the DNS server and change the IP address to reflect the new subnet. Next, use Windows tools to change the cluster IP address and the host IP address on the Network Load Balancer on the HUB/CAS servers to reflect the new CAS array cluster IP and host IP addresses. Verify that Exchange is running and that users can now send and receive email.

Note: In this solution Exchange CAS/HUB servers were not replicated. Instead, CAS/HUB servers deployed at each site provided enough resources to support users from the other site during the failover event.

Table 8 provides details about test results obtained from using SRM to perform failover of an entire Exchange environment from Site A to Site B.

Table 8. SRM recovery plan test results

SRM recovery plan	Time elapsed (min)	RTO met?
Four Mailbox servers in a DAG	38	Yes

Exchange failback with SRM

VMware Site Recovery Manager version 5 now provides automatic failback functionality. This functionality enables you to perform an automated failback operation after a planned or unplanned failover. A “reprotect” process automatically reverses replication and synchronizes any changes from the recovery site back to the protected site.

Figure 30 shows a screen shot from the SRM Management Console, from which you can initiate a failback operation. vCenter SRM, integrated with the RecoverPoint SRA plug-in, automatically initiates storage replication.

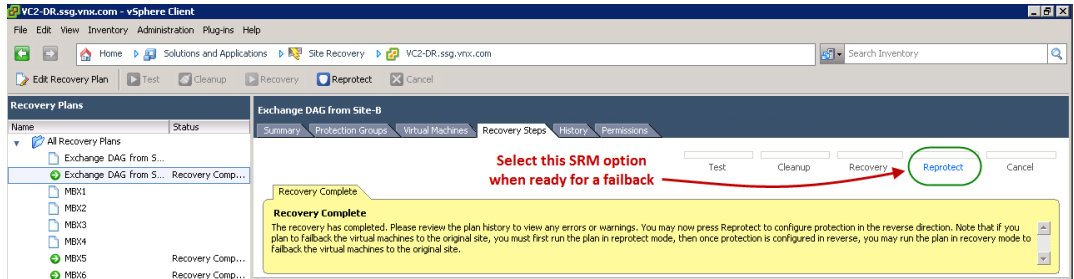


Figure 30. SRM Failback operation available from vCenter SRM

Resynchronization options

In the case of an unplanned failover following an actual disaster, if all storage was lost at the stricken site (the site required a total rebuild), it is likely that a very large amount of data will need to be replicated from the recovery site back to the newly rebuilt, original protected site. This process is time consuming and must be very carefully planned. Even though you can use the SRM automatic reprotect process to initiate resynchronization of the data and configuration settings automatically, you can instead exercise greater control over the resynchronization by initiating replication using the RecoverPoint or Unisphere management consoles.

Note: For the purposes of validating this solution, the protected data was not destroyed following the failover. Thus, the failback and reprotect operation required RecoverPoint to replicate only minimal changes to the data. The SRM failback process for the entire Exchange DAG with two copies took just over 30 minutes.

Failback steps

To perform failback, follow these high-level steps:

1. Verify all network communications between the recovery site and the original protected site.
2. Verify connectivity between vCenter servers.
3. Ensure all resource mappings are in place.
4. Ensure no changes need to be made to SRM protection groups or protected virtual machines.
5. Modify the recovery plan to update any custom scripts. For example, update Exchange DAG IP addresses that have changed. Also update CAS array IP addresses if the CAS array was previously replicated.

Note: For the purposes of validating this solution, Exchange CAS/HUB servers were not replicated. Instead, the CAS/HUB servers deployed at each site provided enough resources to support users from other site following a failover event.

6. To perform failback incrementally, specifying which data to replicate and in what sequence, use the RecoverPoint or Unisphere management console to replicate data and resynchronize configuration settings from the recovery site back to the originally protected site.

Alternatively, to have SRM perform the failback automatically, initiate “Reprotect” from the SRM management interface.

7. When the failback process is complete, ensure that all Microsoft Exchange services are started and the databases are mounted.
8. Verify email flow and user access.

Exchange database recovery with Replication Manager

In this solution, Replication Manager is used to automate RecoverPoint application-consistent replicas and to perform Exchange Server 2010 log truncation.

Restore options

With Exchange 2010, Replication Manager allows you to restore a single database with or without its log volume. This option applies only if you have isolated the database from its logs by placing them on separate devices or volumes.

Figure 31 shows the RM restore options for one Mailbox server in the context of the **Restore Wizard**. You can select any application-consistent bookmark (snapshot) from a local or remote copy, and then select specific databases for recovery. If point-in-time recovery is required for a bookmark other than those created by RM jobs, use the RecoverPoint Console instead. For additional information on using Replication Manager with Exchange 2010, refer to the *Replication Manager Product Guide* and *EMC Replication Manager Support for Microsoft Exchange 2010 Technical Note* available on EMC Powerlink.

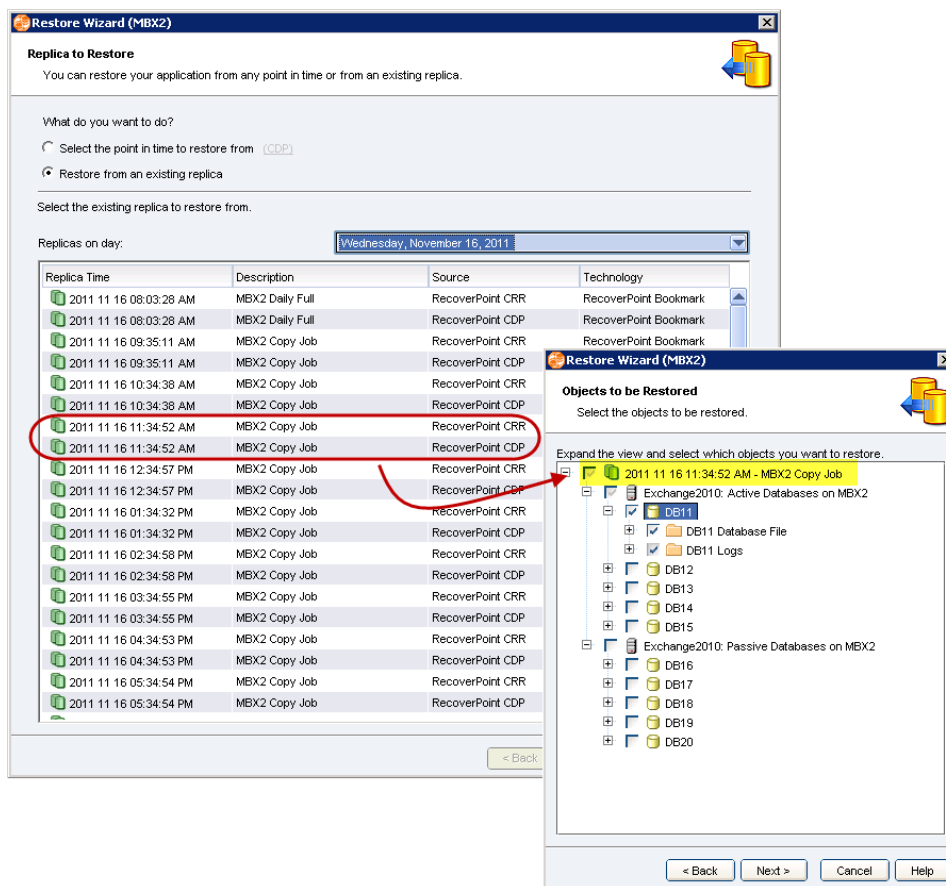


Figure 31. Replication Manager restore options for Exchange 2010

Consistency group restore options

Before Exchange can be restored, you must change the policy option for the RecoverPoint consistency group from “Group is managed by SRM...” to “Group is in maintenance mode...” [Figure 32](#) shows the RecoverPoint consistency group policy options.

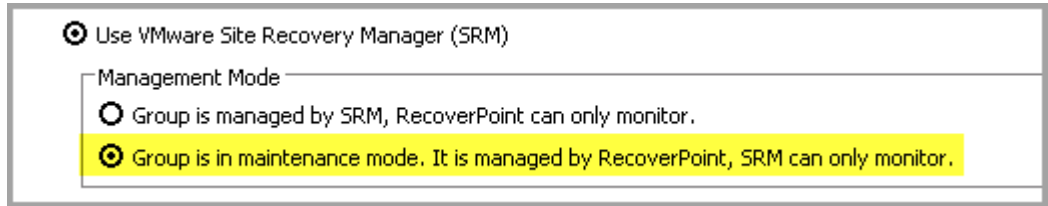


Figure 32. RecoverPoint consistency group policy options

Conclusion

Summary

The EMC Protection for Microsoft Exchange Server 2010 solution is an Exchange business continuity solution designed for enterprises with two or more active data centers at different geographic locations. The solution offers high availability at every location and near-instantaneous recovery from a disaster at any location.

The solution uses EMC VNX5700 storage with the VNX Total Protection Pack (EMC RecoverPoint, EMC Replication Manager, and EMC Data Protection Advisor) and VMware vSphere 5 with vCenter Site Recovery Manager (SRM) 5.

The EMC Protection for Microsoft Exchange Server 2010 solution demonstrates the replication and activation of an entire Exchange Server 2010 Database Availability Group (DAG) and its associated Exchange virtual machines across a WAN between multiple, active/active data centers.

The solution can help achieve the most ambitious recovery point objectives (RPO), recovery time objectives (RTO), and total cost of ownership (TCO) goals.

Specifically, the solution demonstrates the advantages of deploying Exchange on:

- EMC VNX series storage
- VMware vSphere platform

Further, the solution demonstrates various protection options for a multi-site, active/active Exchange deployment:

- High availability at all sites
- Backup and restore at any site
- Recovery from a disaster at any site

Findings

Rigorous testing of this solution achieved the following results:

- Each site provided enough performance and capacity to support up to 20,000 active users (10,000 users during normal operations and another 10,000 users during a disaster recovery event).
- Recovery for all four Mailbox servers completed in just 25 minutes with SRM.
- RTO of less than one hour was achieved during site failover testing with SRM.

References

If you do not have access to <http://powerlink.emc.com>, contact your EMC representative.

White papers

EMC and VMware white papers:

- *EMC Replication Manager and EMC RecoverPoint* (<http://www.emc.com>)
- *Replication Manager Support for Exchange 2010 Technical Notes* (<http://powerlink.emc.com>)
- *EMC RecoverPoint Storage Replication Adapter for VMware Site Recovery Manager Version 2.0 Release Notes* (<http://powerlink.emc.com>)
- *Microsoft Exchange 2010 on VMware Best Practices Guide* (<http://www.vmware.com>)
- *Microsoft Exchange 2010: Storage Best Practices and Design Guidance for EMC Storage* (<http://www.emc.com>)
- *EMC Replication Manager Support for Microsoft Exchange 2010 Technical Notes* (<http://powerlink.emc.com>)

Product documentation

EMC and VMware product documentation:

- *EMC VNX Series Unified Storage Systems Specification Sheet* (<http://www.emc.com>)
- *EMC Replication Manager Product Guide* (<http://powerlink.emc.com>)
- *EMC RecoverPoint Administrator's Guide* (<http://powerlink.emc.com>)
- *VMware vCenter Site Recovery Manager Administration Guide 5.0* (<http://www.vmware.com>)

Other documentation

Microsoft TechNet articles:

- Exchange Server 2010 deployments for high availability and site resiliency, <http://technet.microsoft.com/en-us/library/dd638121.aspx>
- Microsoft guidelines considered for this solution's design, <http://technet.microsoft.com/en-us/library/ee712771.aspx>
- Exchange database corruption types and lagged copy limitations, <http://technet.microsoft.com/en-us/library/dd335158.aspx>
- LoadGen, <http://technet.microsoft.com/en-us/library/dd335108.aspx>
- Configuring Exchange 2010 DAGs and all underlying cluster and Active Directory components, <http://technet.microsoft.com/en-us/library/dd638215.aspx> and <http://technet.microsoft.com/en-us/library/dd297985.aspx>
- Understanding Active Manager, <http://technet.microsoft.com/en-us/library/dd776123.aspx>