

VIRTUALIZED EXCHANGE 2010 DISASTER RECOVERY ENABLED BY EMC SYMMETRIX VMAX AND VMWARE vCENTER SITE RECOVERY MANAGER

A Detailed Review

EMC GLOBAL SOLUTIONS

Abstract

This white paper details an all-VMware solution that leverages the VMware High Availability feature for Microsoft Exchange 2010 for HA and VMware® Site Recovery Manager with EMC® SRDF® for DR. The paper provides best practices and also details performance testing data that will help with the overall design and implementation of this type of solution.

May 2011



vmware®
PARTNER

TECHNOLOGY
ALLIANCE

Copyright © 2011 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

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

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

Part Number H8151.1

Table of Contents

Executive summary	5
Overview	5
Business case	5
Introduction	6
Purpose.....	6
Audience.....	6
Terminology	6
Technology overview	8
Hardware and software resources.....	9
Microsoft Exchange Server 2010 storage design on EMC VMAX	11
Overview	11
Design guidelines.....	11
Building block approach for sizing storage for Microsoft Exchange Server 2010	12
Phase 1 – Collect user requirements	13
Phase 2 - Design the storage architecture based on user requirements	14
Disk requirements calculations based on IOPS.....	15
IOPS calculation example.....	15
Disk space calculation	15
Virtual Provisioning design on the Symmetrix VMAX.....	16
Phase 3—Validate the design	17
vSphere VM Building block.....	17
Virtualized Microsoft Exchange Server 2010 local protection design and guidance on EMC VMAX	20
Overview	20
VMHA rapid recovery from outages.....	20
Deploying Microsoft Exchange for High Availability using VMHA.....	21
Backup and restore of Microsoft Exchange 2010 data with Replication Manager	22
.Replication Manager design	22
Replication Manager design layout.....	23
Rapid restore of Microsoft Exchange 2010 using Replication Manager	26
Virtualized Microsoft Exchange 2010 Server disaster recovery with EMC SRDF and VMware SRM	27
Solution design for disaster recovery.....	27
SRDF / Asynchronous	27
SRDF Adapter for VMware Site Recovery Manager	28
SRDF design	28
Implementation of the SRDF Adapter	31

VMware Site Recovery Manager	32
Protection group and recovery plan layout for Microsoft Exchange 2010	32
Microsoft Exchange 2010 failover steps	34
DR recovery time objective details	35
Microsoft Exchange 2010 failback steps	36
Additional testing for Microsoft Exchange 2010.....	36
Testing SRM recovery plans using TimeFinder/Snap technology	37
Impact on SRDF when local restore is performed using Replication Manager	37
Performance testing and validation results	38
Methodology and tools.....	38
Data collection points	38
Jetstress	38
LoadGen.....	39
Validation results with Jetstress	39
Environment validation with LoadGen	39
24-hour end-to-end validation testing	40
ESX and VM performance results	41
VM High Availability performance testing	42
Microsoft Exchange client results	43
Storage performance	44
Performance results summary for Replication Manager	46
Conclusion	47

Executive summary

Overview

This document outlines the design guidance and architecture for a solution for Microsoft Exchange 2010 enabled by EMC® Symmetrix® VMAX™, VMware vSphere™ 4, EMC SRDF® and VMware® Site Recovery Manager. This document outlines the implementation and test procedures for the Microsoft Exchange 2010 HA and DR solution, based on features in VMware vSphere 4.1, leveraging VMHA features for local HA and SRDF/A for array replication to DR. VMware Site Recovery Manager was used to automate restart of virtual guests in the DR site.

EMC's commitment to consistently maintain and improve quality is led by the Total Customer Experience (TCE) program, which is driven by Six Sigma methodologies. As a result, EMC has built Customer Integration Labs in its Global Solutions Centers to reflect real-world deployments in which TCE use cases are developed and executed. These use cases provide EMC with an insight into the challenges currently facing its customers.

Business case

When designing Microsoft Exchange 2010 architectures, today's customers need to weigh the different high availability (HA) and disaster recovery (DR) options (DAG vs. other methods), and identify the one that best suits their requirements. Today, customers are interested in alternatives to Microsoft Exchange DAGs for local HA and remote DR. To meet these demands, a popular and cost-effective approach would be to leverage VMHA features (vMotion®/DRS) to provide local HA, and utilize EMC replication technology, combined with VMware Site Recovery Manager (SRM) for disaster recovery as an alternative to DAG. This solution rolls out formal design guidance and performance testing data to customers and the field and also provides implementation guidance specific to an all-EMC/VMware HA/DR solution for Microsoft Exchange 2010.

Introduction

Purpose

The purpose of the solution is to:

- Validate the process of creating virtualized Microsoft Exchange 2010 building blocks for the VMAX storage array with VMware vSphere.
- Show how virtual provisioning can successfully and easily be used on a VMAX with Microsoft Exchange 2010 mailbox database volumes to provide unlimited mailbox growth.
- Provide design guidance and recommendations to perform site recovery operations for the Microsoft Exchange 2010 environment in a disaster recovery scenario using VMware Site Recovery Manager and EMC SRDF.
- Provide design guidance and recommendations that will leverage VMHA features for local high availability of the Microsoft Exchange 2010 messaging environment.
- Provide design guidance and recommendations as well as performance details about using Replication Manager and VMAX snapshots to back up and restore Microsoft Exchange 2010 on VMAX.

Audience

The intended audience for the white paper is:

- Internal EMC personnel, who are responsible for designing and selling EMC application related solutions.
- EMC Partners, who are responsible for designing and selling EMC application-related solutions.
- Customers, who are responsible for purchasing and designing application related solutions.

Terminology

Table 1 defines terms used in this document.

Table 1. Terminology

Term	Definition
EMC Replication Manager	The EMC Replication Manager (RM) product provides simplified management of storage replication and integration with critical business applications to take disk-based copies that serve as foundation for recovery operations.
EMC SRDF	The SRDF Family of software is the most powerful suite of remote storage replication solutions available for disaster recovery and business continuity. Fully leveraging the industry-leading high-end Symmetrix hardware architecture, it offers unmatched deployment, flexibility, and massive scalability to deliver a wide range of distance replication capabilities to meet mixed service-level requirements with minimal operational impact.

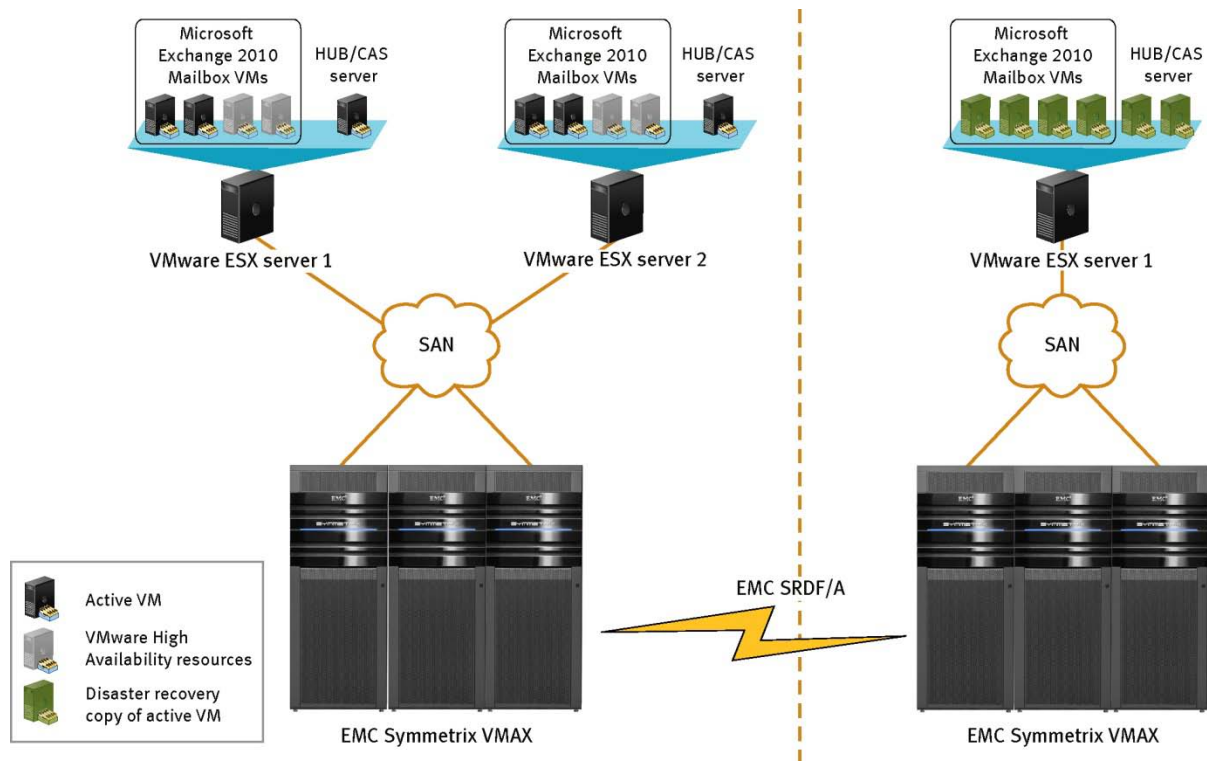
EMC TimeFinder®	The TimeFinder Family of software is the most powerful suite of local storage replication solutions available. Fully leveraging the industry leading high-end Symmetrix hardware architecture, it offers unmatched deployment flexibility and massive scalability to deliver a wide range of in-the-box data copying capabilities to meet mixed service-level requirements with minimal operational impact.
EMC VMAX	The array offering by EMC to provide high-end storage for the virtual data center. Its innovative EMC Symmetrix® Virtual Matrix Architecture™ seamlessly scales performance, capacity, and connectivity on demand to meet all application requirements.
VMware vSphere	VMware vSphere is VMware's most complete and robust virtualization platform, transforming datacenters into a dramatically simplified cloud infrastructure and enabling the next generation of flexible, reliable IT services.
VMware VMHA	VMware High Availability (HA) provides easy-to-use, cost-effective high availability for applications running in virtual machines. In the event of physical server failures, affected virtual machines are automatically restarted on other production servers with spare capacity.
VMware Site Recovery Manager	Accelerate recovery and ensure successful recovery by automating the recovery process and eliminating the complexity of managing and testing recovery plans. VMware vCenter Site Recovery Manager eliminates complex manual recovery steps and removes the risk and worry from disaster recovery.

Technology overview

The enterprise solution for Microsoft Exchange 2010 enabled by VMAX, SRDF, and VMware SRM provides:

- Easy and improved manageability of Microsoft Exchange servers with VMware vSphere and EMC Symmetrix VMAX.
- Reduced total cost of ownership (TCO) by reducing initial allocation of storage capacity and simplifying management. Thin provisioning enabled by the EMC Symmetrix Virtual Provisioning™ feature allows the application to use space only as needed, without reallocation or reconfiguration.
- Local HA capability by leveraging VMware VMHA features that allow for only one copy of the database locally, thus providing server and storage cost reductions.
- An easy and automated solution to recover the Microsoft Exchange environment on an alternate site in the event of disaster recovery using VMware SRM and SRDF.
- Exceptional backup and recovery performance so that data remains highly available. Once a snapshot of the Microsoft Exchange data is taken and presented to the mount host, you can use any backup software to take a backup from that copy, which is isolated from the production environment.

The following illustration depicts the overall solution:



SYM-002531

Hardware and software resources

Table 2 describes the equipment used to validate the solution.

Table 2. Hardware

Hardware Equipment	Quantity	Configuration
Symmetrix VMAX	2	uCode: 5875 4 engine 240 GB cache 240 600 GB 10k FC drives 240 1 TB 7.2k SATA II drives
SAN switch	1	Cisco® MDS 9509
IP switch	1	Cisco® Catalyst 3560
ESX® server	3 (2 for prod and 1 for DR)	Dell PowerEdge R900 6 core, 128 GB RAM 2 dual-port 4 Gb QLogic HBA On prod, each ESX server will host 2 Microsoft Exchange VMs and 1 HUB/CAS On DR, the ESX server will host all 4 Microsoft Exchange VMs and 2 HUB/CAS
HBA	6	2 dual-port, 4 GB HBAs (QLA2562) per host
Total number of disks tested in this solution	88	1 TB 7.2k SATA II drives
Replication Manager Server and Mount Host	1	Dell PowerEdge R900 6 core, 128 GB RAM 2 dual-port 4 GB QLogic HBAs
Storage connectivity		Fibre Channel

Table 3 describes the software used to validate the solution.

Table 3. Software

Software	Configuration
VMAX microcode	5875
HBA driver	QLogic 9.1.7.16
Multi-pathing software	EMC PowerPath®/VE 5.4 (64 bit)
Host OS	Microsoft Windows Server 2008 – R2
Microsoft Exchange Server 2010	14.01.0218.015
VMware vSphere	4.1
Replication Manager	5.3

Microsoft Exchange Server 2010 storage design on EMC VMAX

Overview

Storage design is a critical element to successfully deploy a Microsoft Exchange 2010 building block on the Symmetrix VMAX. The process is essentially the same design for a physical environment as a virtual building block, from a disk perspective, except that in the virtual environment the VM's OS volume also needs to be accounted for.

The virtualized Microsoft Exchange Server 2010 storage design includes a number of different pieces including:

- Storage Design Requirements - Disk requirements and layout.
- Virtual Provisioning – Pool design.
- vSphere VM Building Block – ESX and VM design and resource requirements. Each of these will be discussed in detail in this section.

Design guidelines

The following list provides design guidance for Microsoft Exchange Server 2010 on a VMAX:

- I/O requirements include user I/O (send/receive), any other overhead (growth, BlackBerry, and so on) plus an additional 20 percent. The 20 percent accounts for some overhead as well as log and BDM I/O.
- Database and log I/O should be evenly distributed among the SAN and storage back end.
- For Microsoft Exchange 2010 data it is acceptable to lay out the database and log volumes across the same spindles
- Use larger hyper volumes when creating LUNs to achieve better performance.
- A minimum of two HBAs are required per server to provide for redundancy.
- When using a hypervisor to virtualize the Microsoft Exchange servers a minimum of three IP connections is recommended for each ESX server.
- Always calculate I/O spindle requirements first, then capacity requirements.
- When unsure of the Read/Write ratio, use the default ratio of 3:2 in a Microsoft Exchange DAG environment and 1:1 when in standalone.
- Isolate the Microsoft Exchange server database workload from other I/O-intensive applications or workloads on separate sets of spindles. This ensures the highest level of performance for Microsoft Exchange and simplifies troubleshooting in the event of a disk-related Microsoft Exchange performance issue.
 - Install EMC PowerPath for optimal path management and maximum I/O performance. For more information on installing and configuring the PowerPath application, visit <http://www.emc.com/products/detail/software/powerpath-multipathing.htm>.

- Follow these recommendations to ensure the best possible mailbox server performance:
 - Format new NTFS volumes on Windows Server 2008, to be used for Exchange database and logs, with the allocation unit size (ALU) to 64 KB. This can be done from the drop-down list in Disk Manager or through the command prompt using diskpart.

Note Partition alignment is no longer required when running Microsoft Windows Server 2008 as partitions are automatically aligned to a 1 MB offset.

Visit the following Microsoft links for guidance on determining server memory and CPU requirements for the Microsoft Exchange 2010 Mailbox Server role:

<http://technet.microsoft.com/en-us/library/ee832793.aspx> (Memory)

<http://technet.microsoft.com/en-us/library/ee712771.aspx> (CPU)

Building block approach for sizing storage for Microsoft Exchange Server 2010

Sizing and configuring storage for use with Microsoft Exchange Server 2010 can be a complicated process, driven by many variables and factors, which vary from organization to organization. Properly configured Microsoft Exchange storage, combined with properly sized server and network infrastructure, can guarantee smooth Microsoft Exchange operation and best user experience.

One of the methods that can be used to simplify the sizing and configuration of large Microsoft Exchange Server 2010 environments is to define a unit of measure—a building block. A building block can be defined as the amount of disk and server resources required to support a specific number of Microsoft Exchange Server 2010 users. The amount of required resources is based on:

- Specific user profile type
- Mailbox size
- Disk requirements

Using the building block approach simplifies the implementation of Microsoft Exchange Server 2010 Mailbox Server. Once the initial building block is designed, it can be easily reproduced to support the required number of total users in your organization.

This approach serves as a baseline for Microsoft Exchange administrators to create their own building blocks that are based on their company's specific Microsoft Exchange environment requirements. This approach is very helpful when future growth is expected, as it makes Microsoft Exchange environment expansion much easier, and straightforward. EMC's best practices involving the building block approach for Microsoft Exchange Server design has proved to be very successful throughout many customer implementations.

For a deeper understanding of the process flow used to develop and validate the test environment's storage design, please refer to the [EMC Virtual Infrastructure for Microsoft Exchange 2010 Enabled by EMC Symmetrix VMAX, VMware vSphere 4 and Replication Manager Proven Solution Guide](#).

Note Access to this document requires a Powerlink account.

The following sections describe the storage design process applied for this solution.

Phase 1 – Collect user requirements

The user requirements used to validate both the building block storage design methodology and VMAX performance are detailed in Table 4:

Table 4. User requirements

Item	User Equipment
Total Number of users	20,000
User I/O profile	100 messages sent\received per day = .10 IOPS per user per day in a DAG setup
User Mailbox size	Start with 1 GB, grow to 2 GB
Deleted item retention	14 days
Concurrency	100 percent
RPO	Remote < 5 minutes, local = 6 hours
RTO	60 Minutes
Mailbox Resiliency Solution (DAG)	NO
Backup/Restore required	Yes (Hardware VSS)

The requirements include starting with a user mailbox size of 1 GB with the ability to seamlessly grow to 2 GB. This document shows how this can be easily accomplished using the VMAX Virtual Provisioning feature ([Virtual Provisioning design on the Symmetrix VMAX](#)).

Based on the user requirements a virtual Microsoft Exchange Building block of 5,000 users per server was created. The decision on the number of users per server was based on a number of factors, including:

- Total Number of users - Use this number to find a figure that can be evenly divided by a per-server number.
- User Profile - A larger user I/O profile or large mailbox usually dictates fewer users per building block.
- Recovery Time Objective - With a smaller RTO and depending on the backup and restore technology used, fewer users may be supported within a single building block.
- Array features - The ability of an intelligent array, such as the VMAX, to back up and restore larger amounts of Microsoft Exchange data in seconds makes it much easier to achieve good consolidation.
- Simplicity and ease of design - Fewer larger databases will help and using a SAN and intelligent storage will make this easier to achieve.
- ESX server configuration, memory and number of virtual CPUs available—The building block must fit the hardware resources so that the ESX resources are factored in.

Table 5. Required information for Building Block design

Building Block characteristic	Value
Maximum number of Microsoft Exchange Server 2010 users per server	5000
Number of Databases per server	10
Number of users per database	500
RAID Type	RAID 10
Disk Type	1 TB 7200 rpm SATA II drives
Number of ESX servers for Mailbox VMs	2
Total number of Microsoft Exchange Mailbox VMs	4
Database Read / Write Ratio	1:1
Logs Truncation failure buffer	6

Phase 2 - Design the storage architecture based on user requirements

It is recommended to first calculate the spindle requirements per building block from an I/O perspective, and then for space requirements. The following formula can be used to calculate storage for the Microsoft Exchange Server 2010 database and logs from an I/O perspective on EMC storage:

$$(IOPS * \text{percent R}) + WP (IOPS * \text{percent W}) / \text{Physical Disk Speed} = \text{Required Physical Disks}$$

Table 6. Formula details

Where	Is
IOPS	The number of input/output operations per second
Percent R	The percentage of I/Os that are reads
Percent W	The percentage of I/Os that are writes
WP	The RAID write penalty multiplier (RAID 1=2, RAID 5=4)
Physical Disk Speed	60 IOPS for 7.2k rpm drives or 130 IOPS for 10k rpm drives on VMAX

The calculations for the Microsoft Exchange Server 2010 test environment are summarized below. These requirements were calculated using the formula detailed in the *Microsoft Exchange 2010 Efficiency, Flexibility, Performance, and Availability at Scale Enabled by EMC Symmetrix VMAX, Virtual Provisioning, and VMware vSphere White Paper*, located at: http://powerlink.emc.com/km/live1/en_US/Offering_Technical/White_Paper/h7018-exchange-vmx-vp-vsphere-wp.pdf

Note Access to this document requires a Powerlink account.

Note: When using thick or thin devices with Microsoft Exchange 2010 the initial spindle requirements for the pools must meet the I/O requirements. Once that has been accomplished the sizing calculations should follow.

Disk requirements calculations based on IOPS

- Front-end host IOPS = Number of users * I/O Profile + any additional overhead = User IOPS + 20 percent (Logs, BDM, etc.) = Total user IOPS.
- Back-end Array IOPS including RAID penalty = (Total IOPS *.60 Read) + write penalty 2 R1/0 4 R5 (Total IOPS *.40 Write) = Array IOPS / IOPS per Spindle = Disks required.
- Shown above is the EMC recommended formula for manually calculating disk requirements per server, based on IOPS.
- Additional I/O should be added to the user I/O profile for such things as BlackBerry, ActiveSync and so on, to the I/O Profile.

IOPS calculation example

- 5,000 users * .10 = 500 + 20 percent = 600 IOPS
- RAID 1 FC: R (600*.50) + W 2(600*.50) = 900 / 60 = 15 disks per server

As a result, based on IOPS requirements, 15 spindles are needed.

Disk space calculation

Until the EMC Microsoft Exchange 2010 calculator is available, EMC recommends the Microsoft calculator be used to compute the required number of disks. Current MS Calculations require more than 60 percent capacity over user mailbox target size. EMC's Virtual Provisioning is a key to reducing up-front storage purchase and handling any unforeseen growth.

Capacity calculations based on user requirements: Start with a 1 GB mailbox (initial requirements). The database sizing information below is an example for this solution. It is recommended that the Microsoft Exchange Server 2010 Mailbox Role Calculator be used to size the production databases from a space perspective. Additional information for the Microsoft Exchange 2010 Mailbox Server Role Requirements Calculator may be found at:

<http://msexchangeteam.com/archive/2009/11/09/453117.aspx>

Capacity calculations need to be performed as a two-fold process; once to determine the capacity required for the initial requirements, and then to calculate the LUN size for the final capacity.

Capacity requirements = Total Database capacity per server + Total Log Capacity per server

Database size = <Number of Mailboxes> x <Mailbox Size on Disk> x <Database Overhead Growth Factor>

500 users x 2.25 GB + 20% = 1350 GB

Based on this the database LUN capacity can be calculated as follows:

- Database LUN size = <Database Size> + <Content Index Catalog / (1 - Free Space Percentage Requirement)>

- $(1350 \text{ GB} + 135 \text{ GB}) / 0.8 = 1856 \text{ GB}$
- Log LUN Size = Logs per day per user x number of users x truncation failure tolerance x 20%
- $20 \text{ logs @ } 1 \text{ MB} \times 500 \text{ Users} \times 6 \text{ days} \times .20 = 75 \text{ GB}$

Total Database capacity per server = $\langle \text{Database LUN size} \rangle * \langle \text{Number of databases} \rangle$

$= 1856 * 10 = 18560 \text{ GB}$

Total Log capacity per server = $\langle \text{Log LUN size} \rangle * \langle \text{Number of databases} \rangle$

$75 * 10 = 750 \text{ GB}$

Total capacity = $\langle \text{Total Database capacity per server} \rangle + \langle \text{Total Log capacity per server} \rangle$

$18560 + 750 = 19310 \text{ GB}$

Usable capacity available per 1 TB SATA drive = 917 GB

Spindle requirement per server = $\langle \text{Total capacity} \rangle / 917$

$19290 / 917 \sim 22 \text{ disks}$

The capacity spindle calculation was made for a 2 GB mailbox. When starting with a 1 GB mailbox, the number of disks required will be much less, thus providing significant cost savings.

Virtual Provisioning design on the Symmetrix VMAX

This section details how VMAX Virtual Provisioning was used to provide a well-performing, easy-to-use, and economical design for Microsoft Exchange Server 2010. One of the main advantages of using VMAX Virtual Provisioning for Microsoft Exchange is the ability to easily increase the database volumes' storage as the users' mailboxes grow, without any server interruptions. This can deliver tremendous savings in disk cost, power, cooling, and reduced footprint. It also provides for outstanding storage flexibility for your Microsoft Exchange environment. VMAX Virtual Provisioning works with VMware, Hyper-V, and storage technologies such as snaps, clones and SRDF. For more information on VMAX Virtual Provisioning, see the following on Powerlink (access required):

http://powerlink.emc.com/km/live1/en_US/Offering_Technical/Technical_Documentation/F_S_Symm_Virtual_Provisioning.pdf

With Virtual Provisioning for Microsoft Exchange on a VMAX there are three main approaches that can be taken with regards to the thin pool design; they include:

- A thin pool for each Microsoft Exchange Mailbox Server, which provides for more granular (building block) design, easier troubleshooting and analysis.
- One large thin pool for all Microsoft Exchange database and logs, which is the simplest and provides for best utilization of space but may make troubleshooting and performance analysis more challenging.

- One thin pool supporting multiple applications, which is not recommended and should only be done when the I/O of all applications is well understood and will not change.

Other recommendations for the design of the disk layout for Microsoft Exchange Server 2010 with thin pools on VMAX are outlined below:

- Use large data devices (120 - 240 GB) for the thin pools.
- Use concatenated thin device metas, since the data devices are already striped, from a disk group perspective.

The design used for our testing environment included:

- A single thin pool per Microsoft Exchange Mailbox Server.
- Each thin pool was supported by a disk group containing 22 1 TB 7200 rpm SATA II drives.
- 240 GB data devices were used with mirrored protection.
- 20 devices per pool, 10 thin devices at 1.9 TB for the database, 10 thin devices at 75 GB for the log.

This calculation is performed for spoofing storage requirements to hosts.

Database thin device details are as follows:

- A 1.9 TB DB thin device was used to accommodate the 2 GB mailbox requirement.
- 10 DBs per server - 500 users per database.
- Database LUN size = $\langle \text{Database Size} \rangle + \langle \text{Content Index Catalog} \rangle / (1 - \text{Free Space Percentage Requirement}) = 1.9 \text{ TB}$.

In order to prevent a situation where a thin pool runs out of space, it is recommended that the thin pool utilization threshold tool be used to monitor each pool's disk space utilization. These settings can be found within the Symmetrix Management Console. The tool can be set to send warnings and alerts when the pools space utilization reaches a certain level.

Phase 3—Validate the design

Microsoft Exchange Jetstress and Microsoft Exchange LoadGen tools were used to validate the storage design. For a complete summary of Jetstress and LoadGen findings, see Performance Testing and Validation Results.

vSphere VM Building block

Once the user per building block and disk calculations are complete the virtual machine and ESX requirements can be calculated. The memory and CPU requirements are based on Microsoft best practice for Microsoft Exchange Server 2010 Mailbox Servers. Additional information may be found at:

<http://technet.microsoft.com/en-us/library/dd346700.aspx>

CPU and memory requirement calculations start with the mailbox server role. Based on the requirements, the building block must be able to support 5,000 users per server. Provisioning sufficient megacycles so that mailbox server CPU utilization does not exceed 80 percent is

also required. Table 7 lists the mailbox server megacycle and memory requirement calculations.

Table 7. Mailbox CPU requirements

Parameter	Value
Active megacycles	5000 mailboxes x 3 megacycles per mailbox = 15,000
Passive megacycles	0 (sizing for all active)
Replication megacycles	0 (sizing for all active)
Maximum mailbox access concurrency	100%
Total required megacycles during mailbox server failure	15,000

Using megacycle capacity to determine the number of mailbox users that a Microsoft Exchange Mailbox Server can support is not an exact science. There are a number of factors that can result in unexpected megacycle results in test and production environments. Therefore, megacycles should only be used to approximate the number of mailbox users that a Microsoft Exchange Mailbox Server can support. Also, per VMware recommendations, a 10 percent overhead needs to be factored-in for hypervisor overhead. Please refer to the VMware paper for Microsoft Exchange 2010 below for more details.

http://www.vmware.com/files/pdf/Exchange_2010_on_VMware_-_Best_Practices_Guide.pdf

Remember that it is always better to be a bit conservative rather than overly aggressive during the capacity planning portion of the design process.

Based on Microsoft guidelines and server vendor specifications we have determined CPU and memory requirements for each VM role. Table 8 provides the summary of the VM CPU and Memory configurations, which has taken into consideration an ESX server failure or a DR scenario when all the mailbox and HUB/CAS servers need to run on a single ESX server.

Table 8. VM CPU and memory configurations summary

Virtual Machine Role	vCPUs per VM	Memory per VM
Mailbox (to support 5000 users during failover)	4	16 GB
HUB/CAS	4	16 GB

A 120 GB VMFS volume for the mailbox server OS was provisioned. Raw Device Mappings (RDM) disks were used for the database and log volumes primarily to accommodate the requirement for a hardware VSS snapshot. Table 9 describes those values.

Table 9. Mailbox VM resources

Item	Description
------	-------------

Number of users supported	5000
User profile supported	0.10 (100 messages / user / day)
Database LUN	10 1.9 TB thin LUNs (RDM)
Log LUN	10 75 GB thin LUNs (RDM)
OS LUN	120 GB (VMFS)
VM CPU	4 vCPU Xeon X7460 2.66 GHz
VM memory	16 GB

Table 10. ESX server configuration

Item	Description
Maximum number of users supported	20,000 Heavy
Server type	Dell PowerEdge R900
Total memory	96 GB
Total vCPUs	24
Number of HBAs	2 dual-port 4 Gb QLogic

Virtualized Microsoft Exchange Server 2010 local protection design and guidance on EMC VMAX

Overview

Today, Microsoft Exchange customers have the option of deploying Native Database Availability Groups inside of Exchange Server to provide application-aware high availability. However, other options exist to provide cost-effective and enhanced availability for Exchange Servers such as Hypervisor-based high availability. VMware HA, enabled by vSphere, makes it possible for these organizations to meet their requirements and design for high availability without incurring the additional costs of multiple databases and extra infrastructure. VMware HA leverages multiple ESX/ESXi™ hosts configured as a cluster to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines. Leveraging VMHA for high availability in a standalone Microsoft Exchange 2010 environment is a fully supported configuration by VMware.

VMHA rapid recovery from outages

VMware HA provides for application protection as follows:

- Protection against server failure by restarting the virtual machine(s) on other hosts within the cluster.
- Protection against application failure by continuously monitoring a virtual machine and resetting it in the event that a failure is detected.

One main advantage of VMHA is that it provides a single method to protect all workloads. There is no need to install special software within the application or virtual machine. All workloads are protected by VMware HA. VMHA configuration is a one-time setup and no further action is required to protect the Microsoft Exchange virtual machines. All the VMs are automatically protected. To protect against failures and provide load balancing across the hosts within a cluster can be achieved by combining VMware HA with VMware Distributed Resource Scheduler (DRS).

VMware HA provides several advantages over other failover solutions, they are:

- Basic setup – Once the VMware HA cluster is set up, there is no additional configuration required and all virtual machines in the cluster get automatic failover support.
- Reduced hardware cost and easy setup - Only a single copy of the data is required, which will be a significant savings from a hardware software cooling and energy perspective. The virtual machine acts as a portable container for the Microsoft Exchange application and it can be moved among hosts. Hence, Microsoft Exchange servers are no longer bound by hardware, which can enhance availability in several ways. When using VMware HA, there must be sufficient resources available to fail over the number of hosts that need protection with VMware HA. However, the VMware vCenter™ Server system automatically manages resources.
- Increased application availability – The Microsoft Exchange virtual machines have access to increased availability. Because the virtual machine is able to recover from hardware failure and Microsoft Exchange services starts up at boot, the application will

have increased availability without increased computing needs, even if it is not clustered. By monitoring and responding to VMware Tools heartbeats and resetting nonresponsive virtual machines, it protects against guest operating system crashes.

- DRS and vMotion integration – In the case of an ESX host failure, the virtual machines are restarted on other hosts. DRS can then provide migration recommendations or migrate virtual machines for balanced resource allocation. If one or both of the source and destination hosts of a migration fail, VMware HA can help recover from that failure.
- VMHA is not an ideal solution during a patching operation since there is only a single copy of the mailboxes and the mailboxes cannot be moved to a different mailbox server. This can result in additional downtime. But many customers feel this is offset by the savings and simplicity VMHA can offer.

Deploying Microsoft Exchange for High Availability using VMHA

VMware HA operates in the context of a cluster of ESX/ESXi hosts. To configure high availability for Microsoft Exchange 2010 using the VMHA feature offered by vSphere, follow the standard best practices that are applicable to any other application.

- There must be at least two hosts in a VM cluster and all hosts must be licensed for VMware HA.
- All hosts must be given a unique host name and configured with static IP addresses. When using DHCP, ensure that the IP addresses of the host(s) persist upon reboot.
- All hosts must have access to the same management networks. There must be at least one management network in common among all hosts, and best practice is to have at least two.
- All Microsoft Exchange VMs must have access to the same virtual networks and datastore so that they can run on any host in the cluster.
- All hosts in a VMware HA cluster need to have DNS configured so that the short host name can be resolved to the appropriate IP address from any other host in the cluster.

The following needs to be performed to configure VMHA for the virtualized Microsoft Exchange environment:

- Select the Hosts & Clusters view.
- Right-click the cluster and click Edit Settings.
- On the Cluster Features page, select Turn On VMware HA.
- Configure the VMware HA settings as appropriate for your cluster.
 - Host Monitoring Status
 - Admission Control
 - Virtual Machine Options
 - VM Monitoring
 - Click OK to close the cluster's Settings dialog box

It is required to install VMware tools on the machines for VM monitoring to work. VM monitoring restarts the individual Microsoft Exchange virtual machines if the VM's VMware

Tools heartbeats are not received within a set time. You can also determine the order in which you want the VMs to come up on another host when encountered with ESX failure. The VM restart priority setting within the vSphere configuration will determine the relative order in which virtual machines are restarted after a host failure. In the virtualized Microsoft Exchange environment, it is recommended to always give the higher restart priority for the HUB/CAS servers.

Performance testing with VMHA was conducted as part of this solution and failover metrics were gathered when Microsoft Exchange guests needed to be moved between ESX hosts. **An ESX server failure was simulated and it took approximately 3 minutes and 28 seconds for the Microsoft Exchange VMs to come online on the other ESX server.** For a complete summary of the findings, see Performance Testing and Validation Results.

Backup and restore of Microsoft Exchange 2010 data with Replication Manager

EMC Replication Manager was used to take disk-based VSS snapshots of Microsoft Exchange 2010 for rapid backup and recovery. Replication Manager relies on disk-based replicas (TimeFinder snaps) as the foundation of the recovery operation. Replicas can be used for granular backup and restore of databases and file-systems or for simplifying and automating data refreshes in development testing, quality assurance, and reporting cycles.

Specifically Replication Manager provides the following:

- Eliminates scripting and makes replication technologies easy to use through an intuitive point-and-click user interface and configuration wizards.
- Enables fast reaction or recovery of data in the event of corrupt or lost information.
- Saves operations time by automating the mounting, dismounting, scheduling and expiration of replicas.

For details on how Replication Manager works with Microsoft Exchange 2010, refer to: [EMC Virtual Infrastructure for Microsoft Exchange 2010 Enabled by EMC Symmetrix VMAX, VMware vSphere 4 and Replication Manager Proven Solution Guide.](#)

Note Access to this document requires a Powerlink account.

.Replication Manager design

Replication Manager is a robust enterprise-level application that can be used to perform a number of functions and provide great benefits in conjunction with TimeFinder local replication technologies. Replication Manager can be used with Microsoft Exchange 2010 and TimeFinder snaps to:

- Create Microsoft Exchange database application sets.
- Create online full and copy replicas of Microsoft Exchange databases and logs, using Microsoft Volume Shadow Copy Services (VSS).
- Create replicas of Microsoft Exchange databases protected as part of a Database Availability Group (DAG), whether it is an active or passive copy of the database.
- Check the consistency of replicated data.
- Start on-demand mount and dismount.

- Perform point-in-time or roll-forward restores to the production databases in the event of a corruption.

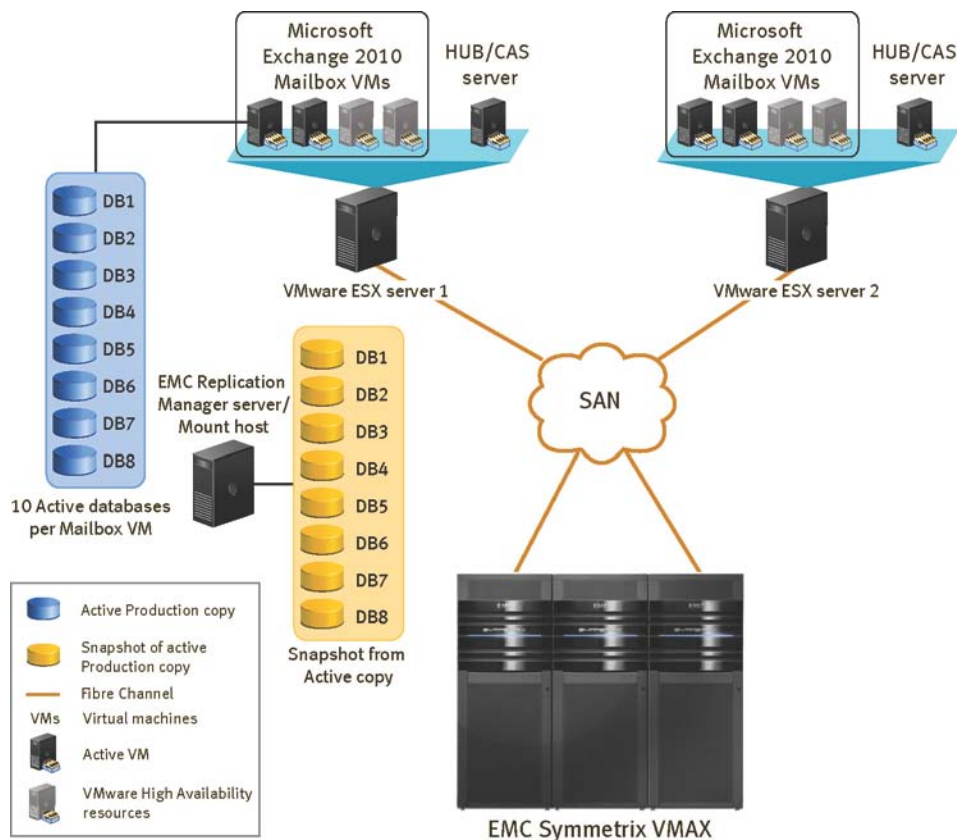
Outlined below are some best practice considerations when designing a Replication Manager environment for Microsoft Exchange 2010:

- Microsoft Exchange 2010 is a capacity bound application more so than a performance application. Hence, as a cost savings approach, it is acceptable to take snaps of the active copy. Performance details are available in the [Performance Testing and Validation Results](#) section.
- It is good practice to separate the database and logs onto different LUNs to take advantage of Replication Manager's roll-forward capability.
- It is recommended to not exceed 5 databases (10 LUNs) per Replication Manager application set. Beyond that, the possibility of VSS timeouts and checksum timeouts increase, which can result in job failures.
- If possible, use a physical mount host as opposed to a virtual machine with physical RDMS as this will reduce the time it takes to mount the volumes.
- In Microsoft Exchange 2010, it is no longer required to run consistency checks on the mounted snap copy. It is however suggested to do this once per week. Since the snap will be taken off the only copy which is active, it is a good idea to run the check during off-hours or during the weekend.
- Use separate dedicated spindles for the save pool for better performance.
- Zone the Replication Manager mount hosts' HBAs to different array ports than the Production Microsoft Exchange server's HBAs for performance reasons.
- Use EMC PowerPath on the mount host for optimal multipathing and load balancing, particularly when running consistency checks.
- It is desired to run consistency checks in parallel.

Replication Manager design layout

The following components outline the environment that will be part of the Replication Manager design:

- VMAX running 5875 microcode.
- TimeFinder snaps for local replication.
- 2 ESX servers; each hosting 2 VM mailbox servers.
- Microsoft Exchange 2010 in a VMHA configuration for high availability.
- Single stand-alone server to run Replication Manager server software. This server also functions as the mount server.



SYM-002548

Figure 2. Replication Manager design diagram

The following are the high-level steps performed by Replication Manager when taking a snap of the Microsoft Exchange 2010 databases:

1. The Replication Manager jobs are configured to first check the event logs for corruption, verify if circular logging is disabled, and map the Microsoft Exchange databases to the back-end storage volumes.
2. The Replication Manager job will then take a TimeFinder snap copy of the databases defined in the application set. This copy is an application-consistent VSS copy of the source databases. The snaps will be taken from the passive DAG copy.
3. Once the snap is taken, Replication Manager will mount the replica (snapshot volumes) into a specified directory on the Replication Manager/mount server.
4. Once successfully mounted, Replication Manager will initiate a checksum process against the mounted volumes to ensure consistency. If this process is successful, Replication Manager will truncate the Microsoft Exchange logs. If it is not successful the Replication Manager job will fail and current transaction logs will not be truncated until the next scheduled backup.
5. After the process has completed, the snap volumes will remain mounted on the mount server until the next replica rotation. This will allow adequate time for the data to be archived to tape.

6. The jobs were scheduled to run four times a day, every day to provide a worst-case RPO of 6 hours. In most cases, the corruption is only on the databases and a roll-forward recovery can be performed, resulting in no data loss.

Table 11. Replication Manager disk backup layout

Database	Replication Tech	Schedule	AppSet	Job	Storage Pool	Mount host
DB1-DB5	TimeFinder Snap	6:00	DB1-DB5	DB1-DB5	DB1-DB5	R900DMX_05
DB6-DB10	TimeFinder Snap	6:30	DB6-DB10	DB6-DB10	DB6-DB10	R900DMX_05
DB11-DB15	TimeFinder Snap	7:00	DB16-DB20	DB11-DB15	DB11-DB15	R900DMX_05
DB16-DB20	TimeFinder Snap	7:30	DB16-DB20	DB16-DB20	DB16-DB20	R900DMX_05
DB21-DB25	TimeFinder Snap	8:00	DB21-DB25	DB21-DB25	DB21-DB25	R900DMX_05
DB26-DB30	TimeFinder Snap	8:30	DB26-DB30	DB26-DB30	DB26-DB30	R900DMX_05
DB31-DB35	TimeFinder Snap	9:00	DB31-DB35	DB31-DB35	DB31-DB35	R900DMX_05
DB36-DB40	TimeFinder Snap	9:30	DB36-DB40	DB36-DB40	DB36-DB40	R900DMX_05

Table 11 highlights the Replication Manager design for this solution. There were several considerations taken into account to design the Replication Manager backup environment:

- Five databases were included in every application set. Since the databases all reside on separate RDM LUNs, it allows for single database as well as roll-forward restores if needed. The five databases per application set helps reduce the number of Replication Manager jobs and essentially results in simplified management.
- Each Replication Manager job is directly associated with an application set and hence configured to take a snap of five databases. The job is configured to take snaps from the active Microsoft Exchange copy. The performance results, when running LoadGen under full load and with snaps active, showed that there was not any significant performance impact on the active database copy. For additional information, refer to the [Performance Testing and Validation Results](#) section.
- A single mount host was required to mount all snap copies.
- The local replication technology used for this solution is TimeFinder snaps. Microsoft recommends hosting multiple full copies of Microsoft Exchange databases on separate servers to provide local recoverability capability. But in most environments with large mailboxes, Microsoft Exchange 2010 is capacity driven more so than performance-driven. For this reason, using snaps will provide huge space savings. TimeFinder snaps are pointer-based copies of the source LUN that store the changed tracks, only; hence they use minimum space on the VMAX array. Finally, Replication Manager integrates well with the TimeFinder/Snap technology and provides for instant restore capabilities, both point in time and roll forward.
- The jobs were scheduled using Replication Manager to minimize administrative tasks and automate the solution. This worked well in both roll-forward and point-in-time restores (since the databases and logs were located on separate LUNs, it also provided for a roll forward restore capability and no data loss solution when the logs were intact).

Rapid restore of Microsoft Exchange 2010 using Replication Manager

The various scenarios and procedures for recovering Microsoft Exchange 2010 with Replication Manager are detailed in [EMC Virtual Infrastructure for Microsoft Exchange 2010 - Enabled by EMC Symmetrix VMAX, VMware vSphere 4 and Replication Manager](#). Please refer to this Proven Solution Guide for step-by-step procedures.

Note Access to this document requires a Powerlink account.

Virtualized Microsoft Exchange 2010 Server disaster recovery with EMC SRDF and VMware SRM

Solution design for disaster recovery

Incorporating disaster recovery can be challenging for Microsoft Exchange 2010 environments. In order to make the failover as fast and fool-proof as possible it is critical to leverage automated built-in solutions to protect the messaging environment from disaster. Most environments require application and hardware-agnostic methods for protecting their mission-critical applications such as email. This solution is designed to meet those demands by leveraging VMware Site Recovery Manager and EMC SRDF remote replication technology to provide for enterprise-level disaster recovery for the messaging environment. VMware vCenter Site Recovery Manager (SRM) leverages VMware vSphere and EMC SRDF remote replication technology to deliver centralized site failover of the Microsoft Exchange 2010 environment and enable dramatically improved testing. A vSphere plug-in SRDF Adapter provides tight integration between SRDF and SRM and helps extend the disaster recovery capabilities of SRM for Microsoft Exchange 2010 to the VMAX storage environment. The solution design is outlined in detail in Figure 3.

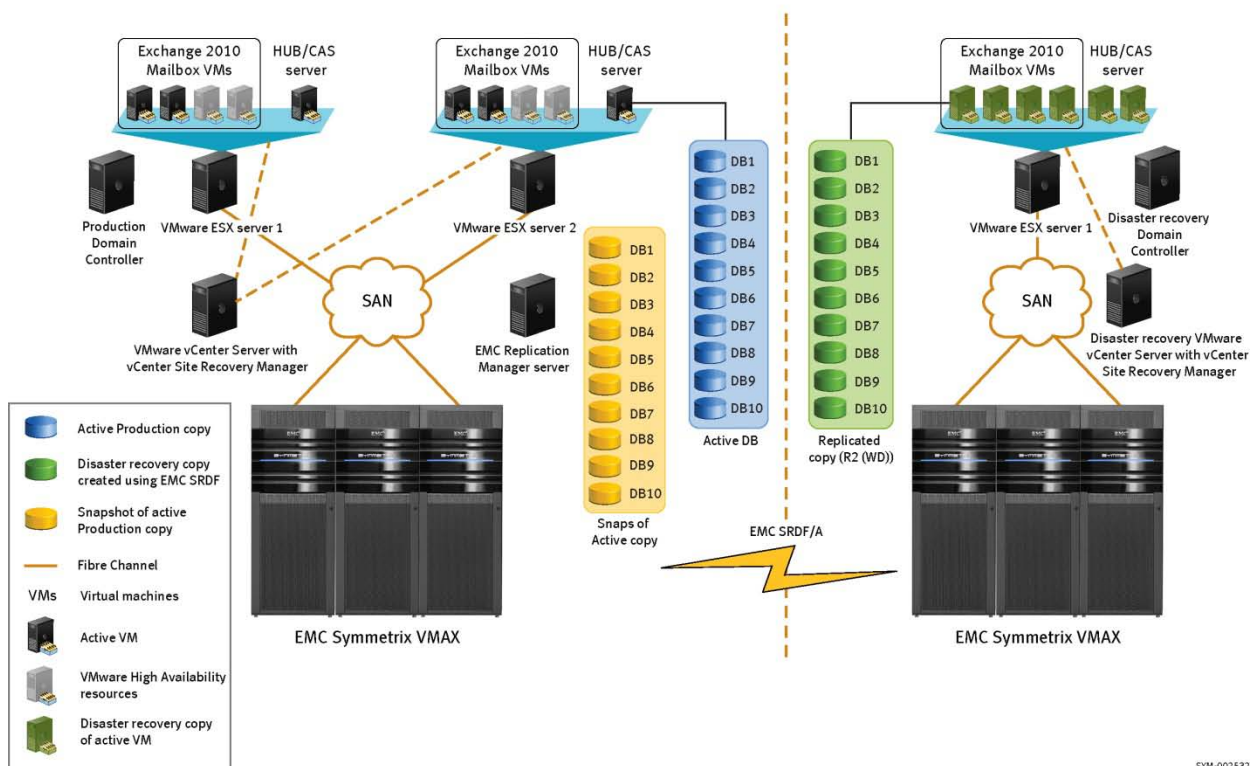


Figure 3. Disaster recovery for Microsoft Exchange 2010 using VMware SRM and EMC SRDF

SRDF / Asynchronous

EMC SRDF/A – SRDF Asynchronous – replicates data from a Symmetrix array in the local site to another in the remote site. With SRDF/A the local Symmetrix acknowledges the host I/O, accumulates these in delta sets, and delivers them to the remote Symmetrix in defined

intervals. By separating the host I/O with replication I/O, it ensures that the latency on the replication link (distance between the sites) will not affect the host response time. This will ensure that data can be replicated over very large distances to a remote site without any production impact and is thus an ideal remote replication solution for enterprise applications like Microsoft Exchange.

SRDF Adapter for VMware Site Recovery Manager

The SRDF SRA adapter is a vSphere plug-in that integrates SRDF with VMware SRM to facilitate simplified management of disaster recovery for mission-critical applications. The VSI plug-in that the adapter leverages to manage the array is now a larger bundle, comprising of multiple plug-ins, including SRA, Storage Viewer, SPM, and path management. The EMC SRDF Adapter for VMware vCenter Site Recovery Manager utilizes Solutions Enabler software to perform the discovery and management of the Symmetrix DMX and VMAX storage arrays on behalf of VMware vCenter Site Recovery Manager.

SRDF design

The solution depicts the implementation of SRDF for remote replication to protect Microsoft Exchange 2010 in a disaster recovery scenario. The EMC SRDF Adapter supports protection of one or more datastores that reside on different devices in a single SRDF/A group enabled for consistency. This functionality was used in this solution by configuring one single device group for all Microsoft Exchange mailbox VMs and the HUB/CAS VMs. The device group consists of all the VMFS volumes for the OS for both the mailbox servers and the HUB/CAS servers, and the database and log RDM volumes.

Table 12 shows the shows the SRDF-related information for the Microsoft Exchange environment. As can be seen below, one device file comprises all Microsoft Exchange database and log RDM volumes and the VM OS volumes.

Table 12. SRDF grouping

Directors	RDFG No.	RDFG Label	RF	Device File	Vol Type	Pairings	Size
7h, 8h	11	EM1	7h, 8h	C:\SRDF\Exchange.txt	DB	397 59F	1920
					Log	539 6DF	75
					DB	39F 5A7	1920
					Log	53A 6E0	75
					DB	3A7 5AF	1920
					Log	53B 6E1	75
					DB	3AF 5B7	1920
					Log	53C 6E2	75

Directors	RDFG No.	RDFG Label	RF Directors	Device File	Vol Type	Pairings	Size (GB)
7h, 8h	11	EM1	7h, 8h	C:\SRDF\Exchange.txt	DB	3B7 5BF	1920
					Log	53D 6E3	75
					DB	3BF 5C7	1920
					Log	53E 6E4	75
					DB	3C7 5CF	1920
					Log	53F 6E5	75
					DB	3CF 5D7	1920
					Log	540 6E6	75
					DB	3D7 5DF	1920
					Log	541 6E7	75
					DB	3DF 5E7	1920
					Log	542 6E8	75
					OS	1CC 70D	120
					OS	562 70E	120
7h, 8h	12	EM2	7h, 8h	C:\SRDF\Exchange.txt	DB	449 5EF	1920
					Log	543 6E9	75
					DB	451 5F7	1920
					Log	544 6EA	75
					DB	459 5FF	1920
					Log	545 6EB	75
					DB	461 607	1920
					Log	546 6EC	75
					DB	469 60F	1920
					Log	547 6ED	75
					DB	471 617	1920
					Log	548 6EE	75
					DB	479 61F	1920
					Log	549 6EF	75
					DB	481 627	1920
					Log	54A 6F0	75
					DB	489 62F	1920
					Log	54B 6F1	75
					DB	491 637	1920
					Log	54C 6F2	75
OS	563 713	120					
OS	1D0 714	120					

Directors	RDFG No.	RDFG Label	RF Directors	Device File	Vol Type	Pairings	Size (GB)
7h, 8h	13	EM3	7h, 8h	C:\SRDF\Exchange.txt	DB	499 63F	1920
					Log	54D 6F3	75
					DB	4A1 647	1920
					Log	54E 6F4	75
					DB	4A9 64F	1920
					Log	54F 6F5	75
					DB	4B1 657	1920
					Log	550 6F6	75
					DB	4B9 65F	1920
					Log	551 6F7	75
					DB	4C1 667	1920
					Log	552 6F8	75
					DB	4C9 66F	1920
					Log	553 6F9	75
					DB	4D1 677	1920
					Log	554 6FA	75
					DB	4D9 67F	1920
					Log	555 6FB	75
					DB	4E1 687	1920
					Log	556 6FC	75
OS	1CF 715	120					
OS	564 716	120					
7h, 8h	14	EM4	7h, 8h	C:\SRDF\Exchange.txt	DB	4E9 68F	1920
					Log	557 6FD	75
					DB	4F1 697	1920
					Log	558 6FE	75
					DB	4F9 69F	1920
					Log	559 6FF	75
					DB	501 6A7	1920
					Log	55A 700	75
					DB	509 6AF	1920
					Log	55B 701	75
					DB	511 6B7	1920
					Log	55C 702	75
					DB	519 6BF	1920
					Log	55D 703	75
					DB	521 6C7	1920
					Log	55E 704	75
DB	529 6CF	1920					
Log	55F 705	75					

Directors	RDFG No.	RDFG Label	RF Directors	Device File	Vol Type	Pairings	Size (GB)
7h, 8h	14	EM4	7h, 8h	C:\SRDF\Exchange.txt	DB	531 6D7	1920
					Log	560 706	75
					OS	565 717	120
					OS	395 718	120
7h, 8h	15	EHUB1OS	7h, 8h	C:\SRDF\Exchange.txt	OS	1CE 70F	120
					OS	561 710	120
7h, 8h	16	EHUB2OS	7h, 8h	C:\SRDF\Exchange.txt	OS	1CB 711	120
					OS	396 712	120

Consistency has been enabled on the device file for all applicable devices. By enabling consistency the SRDF adapter will tag all devices that are part of the device file for grouping when responding to VMware Site Recovery Manager's request for replicated devices. This will show the datastores and the RDM volumes in VMware SRM as being related. Thus when performing a full site failover using SRM, all objects are in sync with each other. This enables for a smooth failover of the Microsoft Exchange mailbox and HUB/CAS servers. For this solution to facilitate this easy approach for a full site failover, all devices were clustered in the same RDFG group. If the requirement is to demonstrate and test single server failover for DR test or other scenarios, it is possible to create separate RDFG groups for each server. In this case, for a full site failover, consistency needs to be enabled across all the devices for the different datastores to show up within SRM as being related.

SRDF Asynchronous mode was chosen for this design, since it provides multiple advantages. SRDF/A provides a long distance solution with minimal performance impact and at the same time preserves data consistency on the DR side, allowing for a restartable R2 copy. This is an important requirement for a tier one application like Microsoft Exchange to counter site failures. The mode also provides write order processing that allows for using lower link bandwidth as compared to other write ordering techniques. SRDF/A implements asynchronous host writes from the source to the target, using predetermined timed cycles called delta sets. Each delta set contains groups of I/Os for processing. SRDF/A then transfers sets of data, one cycle at a time between the R1 and the R2. If the same track is written to more than one time within an active set, SRDF/A will then send the update over the link only once. This mode also allows for achieving dependent write consistency, which ensures that all writes in the R2 are processed sequentially, thus ensuring that the data is always consistent. This solution was tested without introducing any latency on a 1 GB pipe. Depending upon the change rate and distance, it is recommended to size the link appropriately to meet the business SLA.

Implementation of the SRDF Adapter

The EMC SRDF Adapter for VMware vCenter Site Recovery Manager needs to be installed on the server running VMware vCenter Site Recovery Manager. VMware vCenter Site Recovery Manager has to be installed before the EMC SRDF Adapter can be installed. The EMC SRDF Adapter for VMware vCenter Site Recovery Manager leverages the Solutions Enabler software to perform the management tasks on the Symmetrix VMAX storage arrays on behalf of VMware vCenter Site Recovery Manager. It serves this information to VMware SRM where the protection groups and recovery plans are configured for disaster recovery. The host running

Solutions Enabler should have gatekeepers assigned to it to run the in-band commands to manage the Symmetrix storage arrays. This host can be a separate machine or one of the ESX servers can be configured to run SE.

Solutions Enabler version 7.1 and later include a virtual appliance that can be deployed on the vSphere landscape. The version needs to be verified according to the support matrix. Optimally this appliance can be deployed in this setup since it includes the Solutions Enabler version and is preconfigured to provide API services. The entire process of configuring and managing the Solutions Enabler services can be simplified through a web interface.

For detailed instructions on the installation and configuration of the SRA adapter and the SE virtual appliance, please refer to the [Using EMC SRDF Adapter for VMware Site Recovery Manager](#) TechBook. This document is located in Powerlink.

Note Access to this document requires a Powerlink account.

VMware Site Recovery Manager

VMware vCenter Site Recovery Manager (SRM) is a tool that simplifies the disaster recovery process, makes it reliable, manageable, and affordable. SRM integrates well with partners' storage replication software such as EMC SRDF to provide a single stop for management of the recovery process of mission-critical applications like Microsoft Exchange. It transforms the complex hardcopy run books associated with traditional disaster recovery into an integrated element of virtual infrastructure management. VMware vCenter SRM provides functionality to configure site recovery plans for the Microsoft Exchange environment and enables organizations to take the risk and worry out of disaster recovery.

Protection group and recovery plan layout for Microsoft Exchange 2010

This section outlines the VMware SRM design layout. As mentioned in the previous section, the SRDF Adapter presents the replicated SRDF volumes as datastore and RDM volume information to SRM. Since all of the devices are part of a single RDF group, the datastores and RDM volumes show up as related in the SRM GUI. The array managers of SRM will also recognize all of the devices from different groups as a single entity if consistency is enabled on those devices. So to facilitate a full site failover for the Microsoft Exchange environment, one protection group was created encompassing all the required VM OS volumes and Microsoft Exchange database and log LUNs. Under the VMs' section in the protection group, all the applicable VMs will show up. Figure 4 outlines the protection group details for the Microsoft Exchange environment.

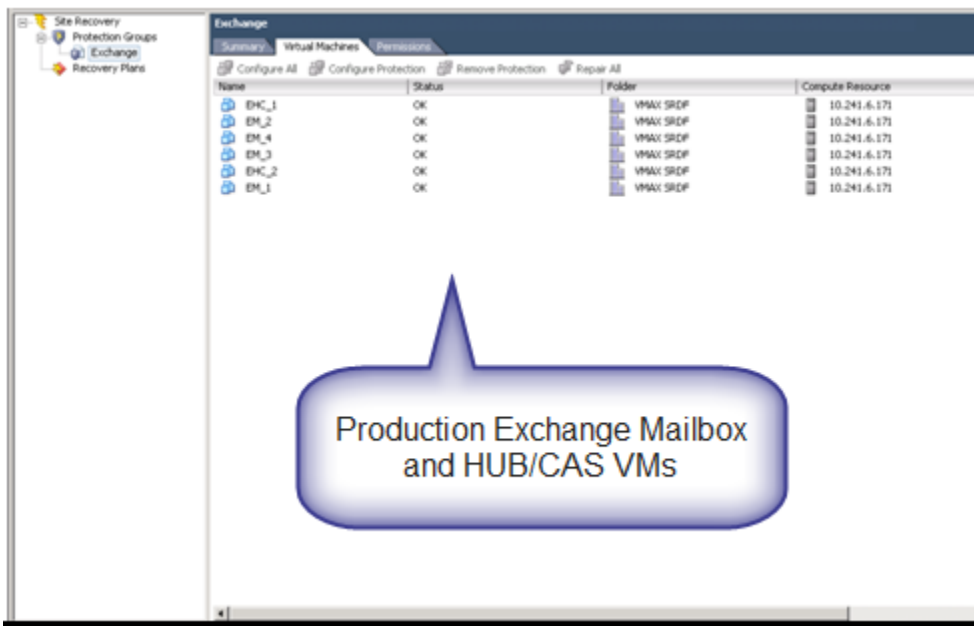


Figure 4. Protection Group layout

VMware Site Recovery Manager uses recovery plans to test and execute the recovery of the protection groups. Once the protection group has been created, a corresponding recovery plan that outlines the recovery procedures was created. The recovery plan was created on the DR vSphere server. The recovery plan includes information about the virtual machines and assigns all of the VMs to the recovery network on the recovery site. It also lists the recovery steps (Figure 6) and once completed gives a detailed history of these steps. Figure 5 shows the recovery plan setup for the Microsoft Exchange environment.

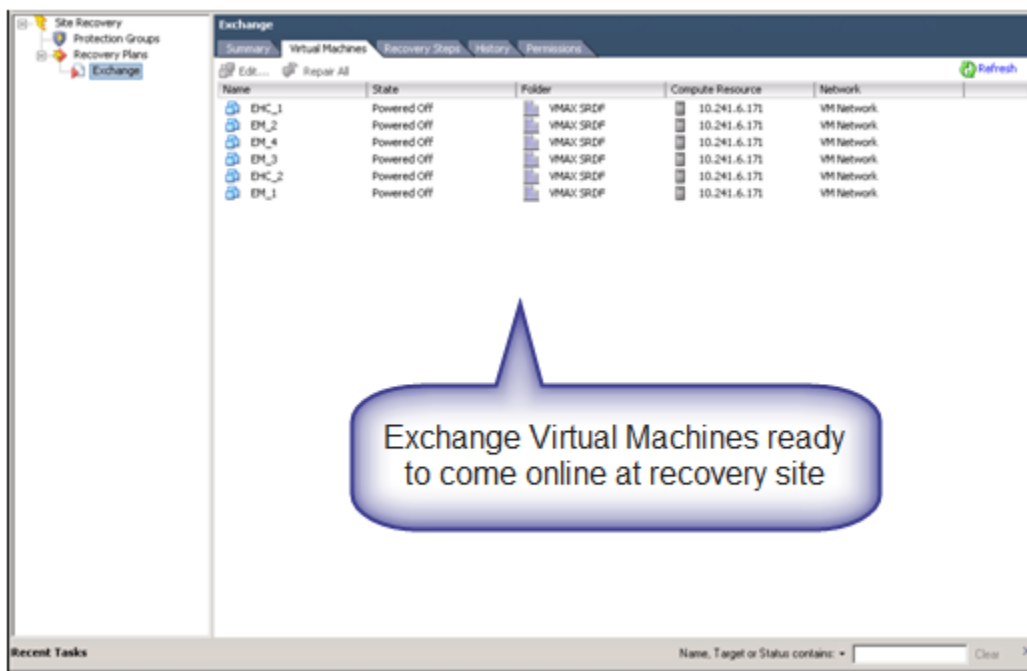


Figure 5. Recovery plan (VM)

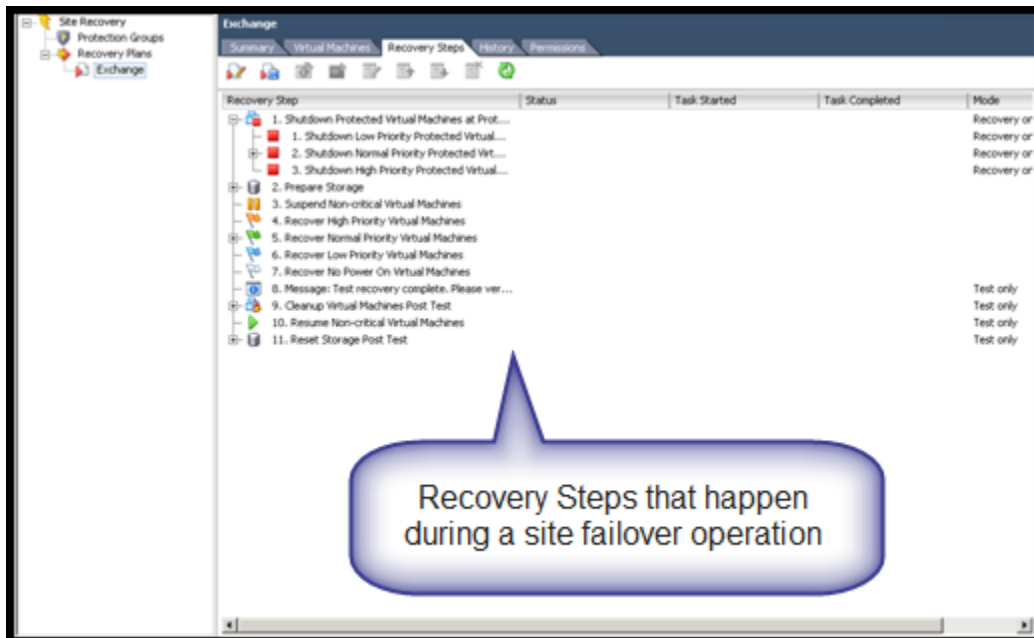


Figure 6. Recovery plan (Recovery Steps)

This section does not detail the creation and configuration of the protection groups and recovery plans. Please consult the corresponding documentation available on the VMware site at:

http://www.vmware.com/pdf/srm_admin_4_1.pdf

Microsoft Exchange 2010 failover steps

The following are the high-level steps that are performed to execute a full site failover of the Microsoft Exchange 2010 environment. The failover will be performed automatically by VMware SRM.

The Microsoft Exchange setup in the solution was configured for long distance disaster recovery and hence the subnet was not extended. The Microsoft Exchange configuration includes one Microsoft Exchange forest and two different AD/DNS sites for the two production and recovery sites. The Microsoft Exchange forest configuration is replicated between the sites. When a disaster occurs, the production site is completely down. At this time, the VMware Site Recovery Manager recovery plan (design detailed in the previous section) is executed. The following recovery steps are performed by VMware SRM when the plan is executed:

- All of the virtual machines (both mailbox and HUB/CAS) are shut down, if they are not already. These VMs can be tagged with priority to be shut down or brought online in a particular order. It is recommended to always tag the HUB/CAS servers with high priority so that these machines are shut down in the end and brought up first, which will facilitate a smooth Microsoft Exchange site recovery.
- At this time, the EMC SRDF Adapter for VMware Site Recovery Manager provides flexibility to the process. The adapter first tries to perform a dynamic personality swap

of the SRDF replicated devices and then converts the recovery sites to read/write enabled. So essentially the R2 devices will now become R1 and vice versa. The replication is still left in a suspended state and can be resumed manually. This is done by setting the **SwapAfterFailover** option to “No”, in the *EMCSrdfSraGlobalOptions.xml* file, in the recovery site.

Note: This file is located in the vCenter server under “C:\Program Files (x86)\VMware\VMware vCenter Site Recovery Manager\scripts\SAN\EMC Symmetrix”.

- The tool then executes tasks to prepare storage on the recovery site. As part of this process, the disks for the protection group (design detailed in the previous section) are attached to the VMs on the recovery site. This essentially does a rescan on the ESX server on the recovery site and presents the volumes to the VMs.
- Once the disks have successfully been attached, the recovery process of the VMs is started. All the virtual machines tagged as high priority are first brought online and then the lower priority VMs are powered up.
- As part of the recovery process, the IP addresses of the Microsoft Exchange virtual machines need to be changed to accommodate the new subnet on the recovery site. The IP addresses were changed manually for this solution, but to completely automate the process, SRM includes a tool called *dr-ip-customizer.exe* that allows specifying the IP addresses of the applicable virtual machines in a recovery plan. The tool edits a comma generated CSV file that is used as input to a command that creates custom specifications for the placeholder virtual machines. The tool is installed in the bin sub-directory of the SRM installation directory. The specified customizations are applied to all of the VMs named in the CSV file during a recovery. The following is an example of the CSV file:

VM ID	VM Name	Adapter ID	MAC Address	DNS Domain	Net BIOS	Primary WINS	Secondary WINS	IP Address	Subnet Mask	Gateway(s)	DNS Server(s)	DNS Suffix(es)
shadow-vm-85974	EM_1	0										
shadow-vm-85974		1	00-50-56-97-00-06	vmax.com				192.168.6.231	255.255.255.0	192.168.6.1	192.168.6.120	
shadow-vm-85970	EM_2	0										
shadow-vm-85970		1	00-50-56-ac-00-01	vmax.com				192.168.6.232	255.255.255.0	192.168.6.1	192.168.6.120	
shadow-vm-85971	EM_4	0										
shadow-vm-85971		1	00-50-56-ac-00-16	vmax.com				192.168.6.234	255.255.255.0	192.168.6.1	192.168.6.120	
shadow-vm-85972	EM_3	0										
shadow-vm-85972		1	00-50-56-ac-00-03	vmax.com				192.168.6.233	255.255.255.0	192.168.6.1	192.168.6.120	
shadow-vm-85969	EHC_1	0										
shadow-vm-85969		1	00-50-56-ac-00-09	vmax.com				192.168.6.35	255.255.255.0	192.168.6.1	192.168.6.120	
shadow-vm-85973	EHC_2	0										
shadow-vm-85973		1	00-50-56-ac-00-07	vmax.com				192.168.6.36	255.255.255.0	192.168.6.1	192.168.6.120	

- Once all the virtual machines are up and running, the CAS array ‘A name’ entry in the DNS server needs to be modified. The IP address needs to be changed to reflect the new subnet. Once this is done, the cluster IP address and the host IP address on the Network Load Balancer (Windows tools) on the HUB/CAS servers need to be changed to reflect the new CAS array cluster IP and host IP.

Now Microsoft Exchange should be up and users should be able to send/receive email.

DR recovery time objective details

The site failover process executed by SRM took approximately 36 minutes to complete. It took an additional 10 minutes to manually change the IP addresses on all of the applicable Microsoft Exchange virtual machines. Overall it took less than an hour to bring all of the

Microsoft Exchange servers online and be able to send/receive email, well within the RTO of 1 hour defined in the customer requirements section. The failover was done multiple times to validate the failover timings. But the timings will vary depending on several factors. Some factors are outlined below:

- VM, ESX host, network, DNS propagation times, storage-related tasks such as presenting disks to the host and SRDF operations timings are not constant always.
- IP change and Active Directory operations can be performed manually or be scripted as part of the SRM failover process, timings of which can vary.
- Application functionality validation after site failover (if all the databases were mounted and email verification, etc.) can vary from environment to environment.

Microsoft Exchange 2010 failback steps

VMware Site Recovery Manager does not provide automation for failback from the disaster recovery site to the production site. The failback process from a full site failover is a manual process. The high-level steps of the process are outlined as follows:

- Ensure that the production site is back up and running.
- If not done already, swap the personalities of the R1 (prod) and R2 (DR) volumes and start synchronization from DR to prod. So essentially the recovery site now assumes the R1 personality and the production site is R2.
- If the production site is a completely new setup, create a new SRDF relationship between the disaster recovery site and production site volumes.
- Wait for the synchronization to complete.
- Once it is complete, shut down the DR servers in site 2 (mailbox and HUB/CAS).
- Allow the leftover tracks to become replicated to the target side. Then after ensuring that the replication is in a consistent state, split the link between the two sites. This will make the volumes on the production site Read/Write enabled.
- Prepare storage to the ESX servers and ensure that they can be accessed by the servers. Then attach the disks to the corresponding Microsoft Exchange VMs.
- Bring up the servers on the production site. Log in and change the IP addresses on all the production machines and the CAS array to reflect the new subnet.
- Confirm the DNS entries for all servers have been updated within Active Directory.
- Ensure that all Microsoft Exchange services are started and that the databases are mounted.
- Verify user email access.
- Swap the personalities of the replicated volumes and start synchronization from production to DR.

Additional testing for Microsoft Exchange 2010

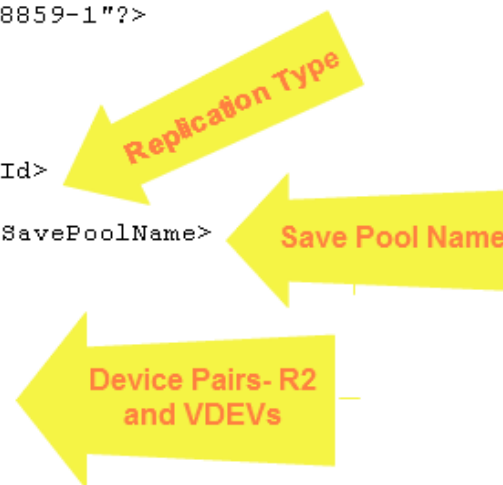
This section describes additional testing and restore scenarios and capabilities, which can be incorporated with this design. These were not performed as part of this solution but are explained in detail.

Testing SRM recovery plans using TimeFinder/Snap technology

The EMC SRDF Adapter for VMware SRM allows for isolated testing of the Microsoft Exchange environment on the DR site, by leveraging TimeFinder snaps of the R2 volumes. This feature is available only with the 5875 microcode. The following are the high-level steps:

- It is essential that the write pacing feature of SRDF/A is turned on for this testing to work. This is because a snap of a R2 volume can cause the SRDF/A session to drop (due to R2 lagging behind too far from R1), and write pacing prevents this from occurring.
- Ensure that the VDEVs are presented to the recovery side ESX cluster to perform the test failover.
- VSI currently cannot be leveraged to create device pairs between the R2 volumes and the VDEVs. So to perform test failures with the snap technology, the `EmcSrdfSraOptions.xml` file on the recovery side needs to be edited to define the device pairs. The replication type also needs to be defined as SNAP technology in the file. The following is a sample file:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<EmcSrdfSraOptions>
  <Version>2.0</Version>
  <TestFailoverInfo>
    <CopyInfo>
      <ArrayId>000190260000</ArrayId>
      <CopyType>SNAP</CopyType>
      <SavePoolName>DEFAULT_POOL</SavePoolName>
      <DeviceList>
        <DevicePair>
          <Source>0767</Source>
          <Target>0867</Target>
          <Source>0768</Source>
          <Target>0868</Target>
        </DevicePair>
      </DeviceList>
    </CopyInfo>
  </TestFailoverInfo>
</EmcSrdfSraOptions>
```



- Highlight the recovery plan and click Test. The EMC SRDF Adapter will then use the information in the options file and perform the test failover on the snap devices. The adapter successfully creates a local copy of the R2 volumes, and SRM will then use this copy and powers on the virtual machines hosted on those datastores.
- It is essential to isolate the IP network of this setup from the production environment to prevent any business interruption.

Impact on SRDF when local restore is performed using Replication Manager

When a local restore operation is performed using Replication Manager from a snapshot, it overwrites the production volume. SRDF is unaware of this operation and hence for the applicable volumes a full resynchronization is required. It is recommended to move the

applicable devices to a different device group. Use the SRDF adaptive copy mode to perform the initial synchronization. This will prevent any host impact while the full synchronization is taking place. Once the tracks have caught up, convert the mode to asynchronous and move the devices back to the original device group. The solution design allows for only the applicable volumes to be moved between device groups and fully synchronized without affecting the other databases.

Performance testing and validation results

The following section describes the approach and methodology used to validate this solution. To ensure that the design and guidance presented in this document deliver good performance, end-to-end validation testing was done on the entire infrastructure. The performance testing was conducted in four sections:

- 24-hour end-to-end performance validation testing of entire solution using LoadGen
- 24-hour end-to-end performance validation testing to test high availability using LoadGen
- Replication Manager Backup and Restore performance results
- TimeFinder/Snap performance and impact with Microsoft Exchange Server 2010

Methodology and tools

The solution involved a combination of a number of components. Each component needs to perform well for the solution to be deemed successful. There were three main focuses for the testing validation; overall 24-hour life cycle testing of Microsoft Exchange under normal conditions, overall 24-hour life cycle testing of Microsoft Exchange to counter the loss of an ESX server, and the backup and restore of Microsoft Exchange Server 2010. Each test was run multiple times to ensure that the results were consistent. Replication Manager testing included LoadGen testing when TimeFinder snaps are active.

Data collection points

To validate the complete health of the solution and identify any bottlenecks, results were collected and analyzed from a number of places. The performance results are grouped into the following areas:

- VMware, ESX and Mailbox VM performance
- Microsoft Exchange Server & Microsoft Exchange Client performance
- Storage and SAN performance
- SRM and VMHA performance
- Replication Manager Backup and Restore performance

Jetstress

The Microsoft Jetstress tool is used to validate the Microsoft Exchange storage design. The tool simulates Microsoft Exchange I/O at the database level by interacting directly with the Extensible Storage Engine (ESE) database technology (also known as Jet), on which Microsoft Exchange is built. Jetstress can be configured to test the maximum I/O throughput available to the disk subsystem within the required performance constraints of Microsoft Exchange.

Jetstress can accept a simulated profile of specific user count and I/Os per second (IOPS) per user to validate that the disk subsystem is capable of maintaining an acceptable performance level by the metrics defined in that profile. It is strongly recommended that Jetstress testing is used to validate storage reliability and performance prior to the deployment of the Microsoft Exchange production environment.

LoadGen

Microsoft Exchange Load Generator (LoadGen) is used for a full end-to-end assessment of the Microsoft Exchange Server 2010 environment. LoadGen can be used to perform pre-deployment validation and stress testing tasks that introduce various workload types into a test (non-production) Microsoft Exchange messaging system. This test simulates the delivery of multiple MAPI, Microsoft Outlook Web access, IMAP, POP, and SMTP client messaging requests to a Microsoft Exchange server.

Important!

LoadGen should only be used in a test lab configuration and in non-production Microsoft Exchange environments. For more information on LoadGen go to the following website:

<http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=cf464be7-7e52-48cd-b852-ccfc915b29ef>

Validation results with Jetstress

In this solution, Jetstress version 14.01.0180.003 was used to simulate an I/O profile of 0.1 IOPS per user. Jetstress was run for four building blocks (20,000 users). The building blocks were validated using 2-hour and 24-hour performance tests. Table 13 shows the I/Os and the average latency across all servers, which reflects the aggregate performance of this solution.

Table 13. Average performance per mailbox VM 5,000 users

Database I/O	
DB Disk Transfers/Sec	885
DB Disk Read/Sec	640
DB Disk Write/Sec	245
Average DB Disk Read Latency (ms)	14.29 ms
Average DB Disk Write Latency (ms)	7.44 ms
Transaction Log I/O	
Log Disk Writes/sec	211
Average Log Disk Write Latency (ms)	1.34
BDM IOPS	299

Environment validation with LoadGen

The LoadGen tool was used to simulate MAPI workload against the entire Microsoft Exchange infrastructure. LoadGen testing is necessary to determine how each Microsoft Exchange component performs under a real, close-to-production user load. LoadGen requires full

deployment of the Microsoft Exchange environment for validation testing. All LoadGen validation testing should be performed in an isolated lab environment, where there is no connectivity to production data. LoadGen generates users and the workloads against the entire Microsoft Exchange environment including Microsoft Exchange Server VM, network and storage components.

LoadGen simulates the entire mail flow, helping to determine any bottlenecks in the solution. It is the only tool that helps in determining the CPU and memory resources that are necessary to sustain the load for which the Microsoft Exchange environment was designed.

In our tests, Microsoft Exchange Server Load Generator 2010 (LoadGen) is used to simulate Microsoft Outlook 2007 online mode mailboxes with the following characteristics:

- The Microsoft Outlook 2007 profile with 100 messages produced 131 Microsoft Outlook tasks per user, per day, and was used for all 20,000 users.
- Each mailbox is 1024 MB in size.
- Eight LoadGen Client machines were used in the testing. Each generating load for 2,500 heavy users.

UserGroups						
Name	Succeeded	Client Type	Action Profile	User Count	Tasks per User Day	TasksCompleted
⊕ UserGroup1	Succeeded	Outlook 2007 Online	Outlook_100	500	131	82455
⊕ UserGroup1_1_	Succeeded	Outlook 2007 Online	Outlook_100	500	131	82215
⊕ UserGroup1_2_	Succeeded	Outlook 2007 Online	Outlook_100	500	131	82319

Figure 7. LoadGen view of the load generated

LoadGen was run a number of times in 10-hour and 24-hour increments to simulate a normal workday's operation and change; the simulated workday duration was set to 9 hours.

During testing the following user load was produced in a 24-hour period:

- The Microsoft Outlook 2007 100 message profile generated a total of 8.25 million tasks per day across all servers.
- The load produced an average of 9.1 GB of logs per DB during 24 hours.
- Approximately 20 logs per day were produced per user, which is a good measure to compare to customer environments.

24-hour end-to-end validation testing

The objective of this test was to validate the entire solution under normal operating conditions for a regular work day. All aspects of the solution were evaluated, including the ESX server and VM's performance, Microsoft Exchange server and client experience, as well as the VMAX array.

In this test, all Microsoft Exchange databases and VMs were placed under normal operating conditions. Replication Manager was set to take snapshot backups of the databases every day. The LoadGen tool was configured to simulate 24-hour heavy load with a nine-hour workday.

ESX and VM performance results

During LoadGen, two of the primary performance counters, ESX and VM CPU utilization, were analyzed. The figures represent the CPU utilization on the two ESX servers and four Mailbox VMs during the daytime hours of the 24-hour test. During this test, LoadGen ran at a heavy load for the first nine hours. As can be seen from the ESX Performance tab:

- The mailbox VMs for the building block on ESX server 1 averaged around 60 percent CPU utilization with spikes up to 71 percent.
- The second ESX server under a similar load averaged approximately 53 percent CPU utilization with spikes to 74 percent.
- As seen below both ESX servers were very similar from a CPU utilization standpoint.
- Overall, the VMware virtual environment performed well under the load and had some additional headroom to handle spikes.




Figure 8. ESX 1 and VM CPU Utilization





Figure 9. ESX 2 and VM CPU Utilization

VM High Availability performance testing

VMware High Availability (HA) testing was performed and an ESX server failure was simulated. The time taken for the Microsoft Exchange mailbox and HUB/CAS VMs on the failed ESX server to fail over and start up on the other ESX server was measured. The server failure simulation was done under load. Once all of the VMs were running on a single ESX server, normal load was simulated using LoadGen and the VM and ESX CPU utilization was measured. The following results were observed:

- The environment handled the failover very efficiently and the Microsoft Exchange VMs on the failed ESX server were able to come online on the other ESX server in approximately 3 minutes and 28 seconds. This was calculated based on the event logs of the servers showing the time when the server went down and the time when the server came up online.
- All mailbox VMs for the building block on the ESX server averaged 54 percent to 70 percent CPU utilization with spikes around 89 percent.
- This indicated that on average, the servers were performing well, though there were some odd spikes.
- One fact that can be extracted from the graph is how indexing affects CPU performance. The following graph shows that there is a decrease of almost 18 percent in CPU utilization on an average with indexing disabled (EM2) as compared to when it is enabled (EM1).



Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	10.241.7.22	Usage	Average	Percent	73.99	84.76	69.02	76.73
	EM_1	Usage	Average	Percent	67.46	91.47	56.32	71.311
	EM_2	Usage	Average	Percent	54.02	71.79	41.93	53.952
	EM_3	Usage	Average	Percent	56.62	89.09	56.62	70.205
	EM_4	Usage	Average	Percent	61.11	82.16	52.74	66.65

Figure 10. ESX & VM performance results

Microsoft Exchange client results

Mailbox server response times for client requests are tracked to determine the amount of time it takes the mailbox server to respond to a client request and gauge the overall client experience. The response time average per request should not exceed 10 milliseconds. Use the following performance monitor counter on the mailbox server to monitor response time:

- MExchangeIS\RPC Averaged Latency

Use the following performance monitor counters on the mailbox server to monitor the message sent and delivered rates:

- MExchangeIS Mailbox (_Total)\Messages Sent/sec.
- MExchangeIS Mailbox (_Total)\Messages Delivered /sec.

The validity of each test run, from a Microsoft Exchange client perspective, was determined by comparing the results of select performance counters to a Microsoft specified criteria. Performance counter data was collected at 10-second intervals for the duration of each test run. The results of the first and last hours were discarded. Results were averaged over the remaining duration of test.

Table 14 lists the Primary counters and validation criteria

Table 14. Primary counters and validation criteria

Counter	Target	Test Results
Avg Processor(_Total)\% Processor	<80%	34% - 39%

Counter	Target	Test Results
Time		
MSExchange Database\I/O Database Reads Average Latency	<20ms	12
MS Exchange Database\I/O Database Writes Average Latency	<20ms	7
MSExchange Database\I/O Log Reads Average Latency	<20ms	0
MSExchangeIS\RPC Requests	<70	5
Average MSExchangeIS\RPC Averaged Latency	<10ms	3

From a Microsoft Exchange client perspective all performance counters were well below the target and indicate a healthy client experience.

For additional information about monitoring Microsoft Exchange Server 2010 performance and other key performance counters, visit "Performance and Scalability Counters and Thresholds" on Microsoft's TechNet site at:

<http://technet.microsoft.com/en-us/library/dd335215.aspx>.

Storage performance

As part of performance testing, all components of the storage were analyzed and the VMAX performance was measured. From a storage perspective the array and thin pools performed well. During the nine-hour day the disk pools reached 75 percent maximum utilization and averaged approximately at 60 percent. It should be noted that when testing Jetstress against the same configuration with the same user profile the disk utilization averaged around 80 percent.

Figure 11 shows a well-performing VMAX. Note that during the nine-hour day the utilization increases and during the off-hours the utilization decreases.




Figure 11. VMAX utilization during 24-hour test run

Another useful view of the array is the disk heat map for the 24-hour test. The view shows the load on all disks. As seen below, the utilization is fairly even across all disks. While a few disks were reaching 70-80 percent utilization overall the Microsoft Exchange layout with snaps performed well.

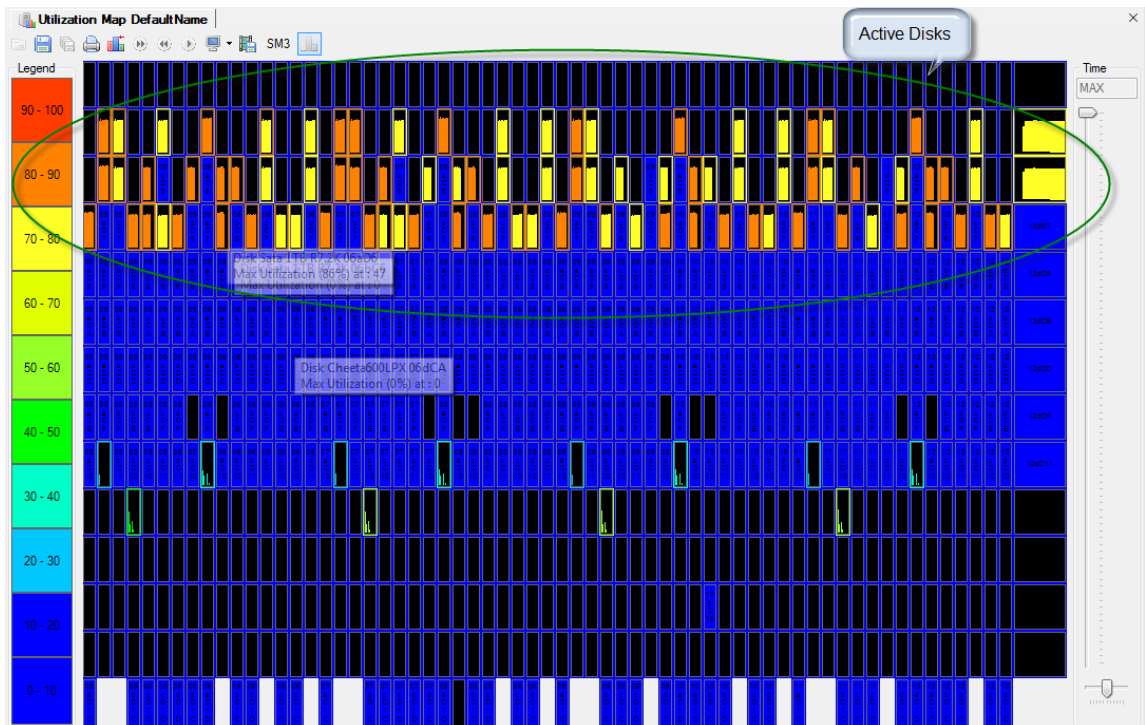


Figure 12. VMAX disk heat map

As part of the performance tests, the front-end utilization of the array was measured. As seen in Figure 13, the front-end directors are performing excellently and are averaging approximately 15 percent in utilization.

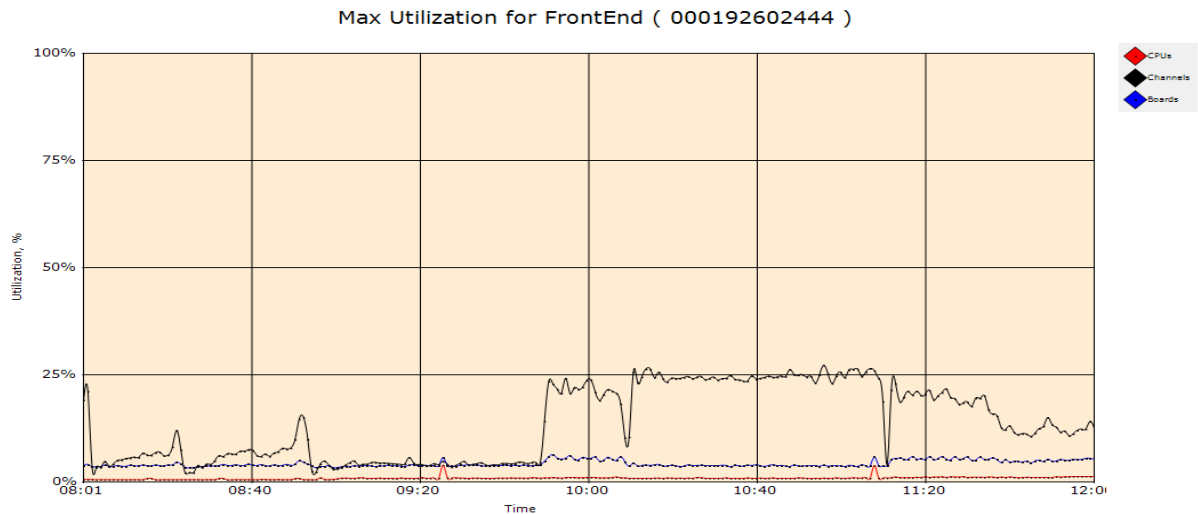


Figure 13. VMAX front-end utilization

SRDF was used to replicate the Microsoft Exchange data to the remote side for disaster recovery. The SRDF workload throughput was measured while full load was generated using LoadGen. As observed in Figure 14, the SRDF throughput was constant at approximately 60 MB/sec.




Figure 14. SRDF workload throughput

Performance results summary for Replication Manager

Replication Manager and snap performance testing have been extensively performed in the previous solutions.

For detailed Replication Manager and snap performance results please refer to the [EMC Virtual Infrastructure for Microsoft Exchange 2010 Enabled by EMC Symmetrix VMAX, VMware vSphere 4 and Replication Manager Proven Solution Guide](#).

Note Access to this document requires a Powerlink account.

Conclusion

The virtualized Microsoft Exchange 2010 building block approach adopted in this solution allows scaling the messaging environment in a cost-effective and flexible manner by combining EMC's Virtual Provisioning technology for storage, EMC Replication Manager for local backup and recovery, VMware VMHA for high availability, and VMware Site Recovery Manager and EMC SRDF for disaster recovery.

Leveraging these proven technologies provides for the following benefits:

- Users can achieve a 4:1 server consolidation ratio by incorporating VMware vSphere as the server virtualization platform.
- The Symmetrix VMAX Virtual Provisioning technology is true virtual provisioning for Microsoft Exchange Server 2010. Unlike DAS virtual provisioning, which involves building new servers, provisioning new storage, and performing mailbox moves to the new servers, EMC Virtual Provisioning eliminates this operational overhead and allows the database volumes to grow seamlessly.
- With EMC Virtual Provisioning, customers purchase only the storage required for the initial mailbox size. As user mailboxes grow, more storage can be seamlessly added with no effect on the users or Microsoft Exchange server performance. The only additional cost is the purchase of additional disk space.
- VMware HA provides easy-to-use, cost-effective, high availability for the messaging environment running in virtual machines. It saves on the cost of purchasing additional hardware for high availability and in the case of an ESX server failure, automatically restarts the affected Microsoft Exchange VMs on the spare ESX server. It also detects OS related failures and restarts the affected Microsoft Exchange VM on the same ESX server.
- Leveraging VMware SRM and EMC SRDF for disaster recovery provides for a single push button solution to counter site failures. This integrated approach delivers centralized management of the recovery of the messaging environment, automates the recovery process, and transforms the complicated run books associated with traditional disaster recovery into an integrated element of virtual infrastructure management.
- With EMC Replication Manager, users can gain a very small backup window regardless of the database size with little to no impact on the production Microsoft Exchange Mailbox Servers. A small percentage of the production space is required for the snap space compared to clones. During testing, with a heavy Microsoft Exchange Server 2010 load, as little as 2 percent of production storage was required to protect the databases for a 24-hour period.
- The automated and rapid recovery capabilities that Replication Manager and VMAX snapshots provide are unmatched. By leveraging the protected restore functionality in the VMAX, a typical database restore and recovery takes only six minutes regardless of the size of the mailbox.