



EMC[®] Replication Manager

Version 5.5.1

Product Guide

P/N 302-000-663
REV 01

EMC Corporation

Corporate Headquarters:
Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright © 2008-2014, EMC Corporation. All rights reserved.

Published March, 2014

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.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

Preface

Chapter 1

Introduction

Replication Manager overview	26
Determining your replication goals	27
Repurposing	27
Backup and recovery	28
Disaster restart	28
Copies of replicas	28
Replication Manager product overview	29
Data management with Replication Manager	30
Replication Manager architecture	32
Console software	32
Server software	33
Agent software	33
User roles	36

Chapter 2

Getting Started

Preliminary setup	40
Preparing the Replication Manager environment	40
Replication Manager software components	40
Starting the Replication Manager Console	41
Components of the main window	43
Performing commands on objects	44
Tree panel	44
Content panel	45
Icon descriptions	46
Getting help	54

Online help	54
Product log files	54

Chapter 3 **Configuring Replications**

Available replication technologies	58
Symmetrix TimeFinder/Clone	58
Symmetrix TimeFinder/Clone (remote)	60
Symmetrix TimeFinder/Mirror	60
Symmetrix TimeFinder/Mirror (remote).....	61
Symmetrix TimeFinder/Snaps	63
Symmetrix TimeFinder/ Duplicate Snaps.....	65
Symmetrix TimeFinder/Snaps (remote)	66
Symmetrix Open Replicator (Symmetrix to CLARiiON or VNX).....	67
CLARiiON SnapView clone.....	68
VNX SnapView clone.....	69
CLARiiON SnapView clone (remote).....	70
VNX SnapView clone (remote).....	71
CLARiiON SnapView snap.....	71
VNX SnapView snap.....	72
CLARiiON SnapView snap (remote).....	73
VNX SnapView snapshots (remote).....	73
CLARiiON ATA disk support.....	74
CLARiiON and VNX MirrorView coexistence.....	74
Full SAN Copy	74
Incremental SAN Copy.....	77
Celerra SnapSure	78
VNXe SnapSure	79
Celerra and VNX Replicator	80
VNXe Replicator	81
RecoverPoint	82
VMware support.....	84
Hyper-V support	104
AIX VIO LPAR support	107
Number of replicas supported.....	110
Managing application sets.....	111
Federated application set support	112
Composite application set support	113
RecoverPoint consistency groups	113
MirrorView /A and MirrorView /S consistency groups.....	113
Application sets and VMware	114
Defining a new application set	116

Validating an application set	117
Modifying an application set.....	118
Deleting an application set	119
Removing application objects from an application set.....	119
About user access to application sets.....	120
Granting or revoking access to an application set	120
Special application set considerations	120
Managing jobs.....	124
Defining a new job	126
Job name and settings panel.....	126
Replication Storage panel	131
Mount Options panel	132
Starting the Job panel	133
Notification panel	134
Completing the Job Wizard panel	135
Topology view	136
Understanding link and copy jobs.....	137
Creating replicas of RecoverPoint targets.....	141
Step-by-step	141
Managing existing jobs.....	141
Simulating a job.....	141
Running a job on demand.....	143
Changing job schedules	143
Modifying a job	146
Changing notifications	146
Restrictions on simultaneous replications.....	149
Deleting a job	149
Effect of failed job on clone synchronization	150
Managing replica rotations	150
Managing job schedules	152
Scheduling a job	152
Modifying a schedule	152
Deleting a schedule.....	153
Stopping a schedule.....	153
Restarting a schedule.....	153

Chapter 4 Daily Operations

Introduction	156
Running an existing job on demand.....	157
Monitoring running tasks	157
Canceling a task.....	157
Storage arrays and devices	158

	Unmounting replicas on demand	160
Chapter 5	Mount, Restore, and Recovery	
	Mounting replicas.....	162
	Mount restrictions and limitations.....	163
	Mounting replicas to an alternate host.....	168
	Alternate mount considerations.....	168
	Mounts to an alternate mount host with Windows	170
	Mounts to an alternate mount host with VxVM.....	170
	Configuring mount hosts to work with storage arrays	171
	Mounting Windows replicas to an alternate host.....	172
	Mounting data in the same location as production.....	172
	Mounting using alternate path.....	174
	Mounting using path mapping	177
	Mounting replicas to the production host.....	182
	Production mount considerations.....	182
	Mounting to production using alternate path.....	183
	Mounting multiple replicas simultaneously.....	184
	Using mount and backup scripts	186
	Mount scripts	186
	Backup scripts	188
	CLARiiON or VNX static mount.....	189
	Clones.....	189
	Snaps	189
	Snaps of clones.....	190
	Mounting replicas to an alternate MSCS cluster.....	191
	Unmounting from an alternate cluster after failover	191
	Mounting replicas to a production MSCS cluster	192
	Mount and restore on UNIX Clusters.....	193
	Overview	193
	Prerequisites	195
	Procedures for mounting a replica (UNIX cluster).....	195
	In case of mount failure	195
	Mounting a replica multiple times.....	195
	Post-mount steps	196
	rm_hacmp_ pvidupdt.pl command.....	196
	rm_serviceguard_ vgidupdt.pl command syntax	197
	Considerations for UNIX cluster restore.....	199
	VMware mount and restore	200
	General mount considerations for VMware mounts.....	200
	Deploying virtual machines from a mounted VMFS	205
	Before unmounting a deployed VMFS.....	205

Restoring VMware replicas	205
File level restore with VMFS and NFS Datastores	208
Restoring a single virtual machine	210
RecoverPoint mount and restore	215
General mount considerations for RecoverPoint	215
RecoverPoint replicas in mount and restore wizards	216
RecoverPoint image access	216
RecoverPoint replicas under SRM or CE management.....	217
Restore considerations for RecoverPoint.....	218
Non-RecoverPoint restores	219
Troubleshooting mount to a Windows 2008 or Windows 2012 host	219
Unmounting a replica	220
Unmounting from a UNIX cluster node.....	221
Replica mount performance guidelines	222
Mount locks	222
Serialization of mounts	223
LUN surfacing and mounting.....	223
LUN unmounting and submerging	223
Summary	223
Restoring from a replica	224
Restore restrictions and limitations	224
Database restores.....	226
File system restores.....	226
Application restore issues.....	227
Restrictions on restoring to synchronized LUNs	228
Restoring a CLARiiON or VNX protected restore environment.....	228
Restrictions on VDEVs	228
Symmetrix TimeFinder restrictions	229
Celerra restrictions	230
Shared storage issues.....	230
Restoring replicas that contain multiple applications	233
Restoring to the R1 device in an SRDF environment.....	234
Restore restrictions in a MirrorView environment.....	234
Troubleshooting restore in a Microsoft Cluster	234
Using application callout scripts.....	235
Application callout naming conventions	235
Specifying callout script processing time	237
Script performance optimization.....	238
Resolving script failures on Windows Server 2008 and Windows Server 2012	238
Application callout numbers for replication.....	238

	Application callout numbers for mount and restore.....	239
	Application callout numbers for unmount.....	240
	Callout script permissions.....	240
	Running PowerShell commands from callout scripts.....	240
	Gathering information about alternate file locations.....	241
	Troubleshooting mount failures on the array.....	245
	CLARiiON and Symmetrix mount failures.....	245
	CLARiiON only mount failures.....	245
Chapter 6	Using Consistent Split	
	Defining consistent-split technology.....	250
	Symmetrix consistent split.....	250
	CLARiiON and VNX consistent split.....	250
	Creating consistent-split replicas.....	253
	Creating consistent-split replicas in Oracle.....	254
	Creating consistent-split replicas in SQL Server.....	255
	Mounting consistent-split replicas.....	256
	Mounting Oracle replicas using consistent split.....	256
	Mounting SQL Server replicas using consistent split.....	256
	Mounting file system replicas using consistent split.....	257
	Restoring consistent-split replicas.....	258
	Restoring Oracle from consistent-split replicas.....	258
	Restoring SQL Server from consistent-split replicas.....	258
Chapter 7	Configuring Federated Data	
	Federated compared with composite application sets.....	260
	Environments that support federated application sets.....	261
	Creating federated application sets.....	262
	Restrictions on composite or federated application sets ...	264
	Creating federated jobs.....	265
	Running a federated replication.....	269
	Federated mount options.....	271
	Partial mounts of federated databases.....	275
	Restoring from a federated replica.....	276
Chapter 8	Replica Management	
	Viewing replicas.....	278
	General tab.....	279
	Objects tab.....	280
	Storage tab.....	281
	History Log tab.....	282

Replicating data on remote storage arrays	283
Creating an SRDF replica	284
Creating a SAN Copy replica	284
Copy a replica using Full SAN Copy	287
Mounts and restores of SAN Copy replicas	288
Modifying CLARiiON or VNX SAN Copy options	289
Remote SRDF/S replication	290
Replication of MirrorView /A or /S secondary	290
Replicating data on Celerra or VNX network file system	292
Celerra NFS storage failover configurations	292
Configuring application sets for Celerra NFS	293
Configuring jobs for Celerra network file systems	294
Availability of RecoverPoint replicas	298
Controlling replica expiration and deletion	300
Deleting a replica	300
Enable or disable replica expiration	300
Disabling replica expiration	301
Enabling replica expirations	301
Modifying replica rotation	302
Rotating replicas	303
Setting retention periods	305
Setting retention periods for replicas	305
Establishing and reestablishing mirrors	307
Understanding Replication Manager scripting	308
Scripting options	308
Additional references	309
Acknowledging a failed replica	310

Appendix A Oracle Procedures

Configuring Oracle for Replication Manager	312
Oracle configuration notes	312
General Oracle configuration for Replication Manager	314
Configuring Oracle for Windows	315
Configuring Oracle for UNIX/Linux	316
Configuring Oracle ASM for Replication Manager	317
Configuring Oracle RAC for Replication Manager	319
Configuring Replication Manager for RAC awareness	320
Configuring Replication Manager to work with failover standalone database in a cluster	322
Configuring Oracle Celerra NFS DR environment	324
Configuring SAP for Replication Manager	328
BRbackup and SPLITINT	328

Overview	328
Configuration procedure	329
Configuring ASM RAC mount to ASM RAC	334
Creating an application set for an Oracle 11gR2 RAC database	335
Setting the reserve policy on AIX shared disks	336
Oracle configuration detailed concepts	337
A note about Oracle user privileges	337
Collecting important Oracle information	338
Configuring the Oracle TCP/IP listener	338
Adding an undiscovered Oracle instance (UNIX only)	339
Deploying Replication Manager in an MSCS Cluster	342
Choosing an Oracle archive log location to replicate	343
Using Oracle ASM with Replication Manager	344
Understanding Oracle application sets and jobs	349
Comparing Oracle replication choices	349
Performing offline replications in an MSCS Cluster with Oracle FailSafe	356
Comparison of replication options with consistent split	357
Upgraded Oracle jobs	360
Mounting Oracle replicas	361
Mount setup and prerequisites	362
Oracle mounts to an alternate location	365
Oracle mount database recovery options	366
Additional mount considerations	377
Restoring Oracle replicas	379
Using pre- and post-replication Oracle scripts	389
Online backups with user-supplied scripts	389
Offline backups with user-supplied scripts	390
General guidelines for scripts	390
Using pre- and post-replication scripts	391
Using the root user to perform Oracle operations	395
Storage Foundation for Real Application Clusters	397
Prerequisite: SSH must be configured	397
Environment variables for SFRAC	397
Prerequisite: SCSI-3 PR must be enabled	398
Mount considerations in SFRAC environments	398
Restore considerations in SFRAC environments	398
Restore steps in SFRAC configuration	400
Oracle troubleshooting	402
TNS permission denied at application set creation	402
Cannot find log archive destination	402
Unable to mount and recover Oracle replica on raw Linux	

devices	403
Database name in the control file does not match	404
Recover mount of Oracle replica fails when SGA is large .	407

Appendix B UDB Procedures

Configuring the UDB environment (UNIX)	410
A typical UDB environment	411
Understanding UDB application sets and replicas	414
UDB online and offline replication	414
UDB replication steps	416
Mounting and restoring UDB replicas	418
UDB mount host setup	418
UDB mount types	418
UDB restores	422
Using pre- and post-replication UDB scripts	424
Offline replications with user-supplied scripts	424
Using pre- and post-replication scripts	425
General guidelines for scripts	427
UDB troubleshooting	428
UDB agent unable to locate the UDB package	428
Workaround for UDB 8 restore failure	429
Workaround for UDB 8/UDB 9 coexistence	429
UDB database is not accessible on the mount host	429

Appendix C SQL Server Procedures

Configuring the SQL Server environment	432
SQL Server prerequisites	434
Required permissions and rights	435
Support for upgrades to SQL Server 2005 and 2008	438
Understanding SQL Server application sets and jobs	439
SQL Server application sets	439
SQL Server 2005/2008 consistency methods	441
SQL Server 2000 replication types	443
SQL Server online replication (with or without VDI)	443
Backing up SQL Server including VDI metadata files	446
Replicating SQL Server 2005/2008 with mirror sessions ..	448
SQL Server 2008 filestream datatype	448
Dynamic Discovery of SQL Server databases	449
Overview of dynamic discovery support	449
System and User databases	450
Creating application sets that use dynamic discovery	451
Viewing application set properties	453

Creating jobs that use dynamic discovery	454
Running jobs that use dynamic discovery	457
Restoring replicas that use dynamic discovery	460
Modifying application sets that use dynamic discovery ...	461
Upgrading implications for dynamic discovery	462
Mounting and restoring SQL Server replicas	464
Considerations when restoring SQL Server 2005/2008 replicas	478
Restoring SQL Server 2005/2008 mirrored databases	479
Using SQL Server snapshot functionality	481
Restore steps in detail	481
Warm standby server	482
Using pre- and post-replication SQL Server scripts	483
Using pre- and post-replication scripts	483
General guidelines for SQL Server scripts	483
Sample replication scripts	484
Using callout scripts in a SQL Server environment.....	484
Considerations for working with SQL Server in a cluster	485

Appendix D Microsoft Exchange Procedures

Setting up Exchange hosts	490
Setting up Exchange production hosts	491
Setting up Exchange mount hosts	492
Setting up Exchange permissions for Exchange 2010 / Exchange 2013	494
Setting up Exchange permissions for Exchange 2007	494
Exchange 2010/Exchange 2013 and Replication Manager	495
Preparing the Exchange 2010/Exchange 2013 environment....	495
Creating Exchange 2010/Exchange 2013 replicas.....	496
Mounting Exchange 2010/Exchange 2013 replicas	507
Restoring Exchange 2010/Exchange 2013 replicas	507
Restoring a deleted Exchange 2010/Exchange 2013 database.	515
Importing a replica from a backup.....	516
Exchange 2007 and Replication Manager	520
Preparing the Exchange 2007 environment	520
Creating Exchange 2007 replicas	522
Replicating various Exchange environments	525
Running a consistency check	535
Consistency check advanced features	536
Managing Exchange errors	537
Mounting Exchange 2007 replicas.....	539

Restoring Exchange 2007 replicas.....	539
Restoring a differential backup in Exchange 2007	541
Restoring with VSS	541
Mounting in an Exchange 2007 CCR environment.....	545
Restoring in an Exchange 2007 CCR environment	546
Restoring in an Exchange 2007 SCR environment	547
Backing up Exchange 2007 replicas	547
Importing a backup as a replica	548
Exchange mailbox recovery procedures	551
Item level restore	553
After performing item level restore.....	554
Using pre- and post-replication Exchange scripts.....	555
Considerations for Exchange in a cluster	556
Additional Exchange 2007 cluster considerations.....	556
Restoring Exchange replicas to a MSCS cluster.....	556
Troubleshooting Exchange issues	558
Considerations in DAG environments with REE enabled .	558
Exchange Interface Service for Exchange 2007 and 2010 ...	558
Modifying the Exchange Interface Service user account or password	559
Activating diagnostic logging for the Exchange Replication Writer	560
Deactivating diagnostic logging for the Exchange Replication Writer	560
Activating diagnostic logging for Exchange 2007/2010 VSS Writer	560
Deactivating diagnostic logging for Exchange 2007/2010 VSS Writer	561
Logging for the Replication Manager Exchange Interface Service.....	561
Resolving failures when you restore databases without transaction logs.....	561
Exchange restore fails with VDS errors due to “devices in use” error.....	561
Storage group restore error: “volume cannot be locked for exclusive use”	562
Exchange replication error: “waiting 1200 seconds for BackupComplete to complete”	563

Appendix E SharePoint Procedures

Overview of support for SharePoint	566
Configuring the SharePoint environment	567

SharePoint prerequisites	567
SharePoint 2010 restrictions	569
Replication Manager deployment in a SharePoint farm ...	569
Host registration in a cluster environment	573
Thin SharePoint replicas	573
SharePoint application sets and jobs.....	576
SharePoint application sets	576
Farm changes that require an application set update	577
SharePoint jobs.....	578
Search activity is paused during replication	579
Overriding thin SharePoint replicas	580
Mounting SharePoint replicas.....	581
SharePoint mount capabilities	581
SharePoint mount prerequisites	581
SharePoint mount options.....	582
Mount options for multi-host SharePoint configurations .	582
Partial mounts of multi-host SharePoint replicas	583
Restoring from SharePoint replicas.....	584
Restore capabilities for SharePoint replicas	584
Using SharePoint with RecoverPoint.....	587
Interoperability considerations.....	587
Creating SharePoint application sets	588
RecoverPoint jobs mount options	588
Running RecoverPoint jobs.....	589
Any Point in Time (APIT) mount and restore	589
GUI intelligence for restoring SharePoint RecoverPoint replicas	590
Troubleshooting SharePoint issues	591
Unable to perform operations for SharePoint	591
Restore fails with VSS_UNKNOWN_ERROR	591
Replication or restore failure after SharePoint server reboot ...	591
Replica fails after restore of a content database	592
Full farm restore fails with “replica does not contain detailed information” error	592
Unable to detect SharePoint 2010 farm when configuring an application set	592
Crawl did not resume after replication	593

Appendix F UNIX File System Procedures

Creating UNIX file system replicas.....	596
UNIX file system mount and restore	597
Restoring UNIX logical volumes.....	597

UNIX raw partition concepts.....	599
Restore limitations for data in file systems	599
Limitations with MPIO devices	599
Linux logical volume manager support.....	600
Using pre- and post-replication application scripts	601
General guidelines for scripts	602

Appendix G NTFS Procedures

Creating NTFS replicas.....	604
NTFS replication operation	604
NTFS mount and restore functions	604
Windows host system layer concepts.....	606
Windows host data layer concepts	606
Using pre- and post replication application scripts	608

Glossary

Index

As part of an effort to improve its product line, EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your EMC technical support professional if a product does not function properly or does not function as described in this document.

Note: This document was accurate at the publication time. Go to EMC Online Support <https://support.emc.com> to ensure that you are using the latest version of this document.

Replication Manager product information

This product guide describes how to operate Replication Manager. Replication Manager that operates in the environments described in the *EMC Replication Manager Support Matrix* on <http://elabnavigator.EMC.com> is the authoritative source for information on supported storage services, operating systems, and applications. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

The appendices describe supported platforms and applications.

Purpose of this guide

This guide describes how Replication Manager conducts the following functions in your storage environment:

- ◆ Create local or remote disk-based replicas of production data stored on:

- **Symmetrix** — Uses TimeFinder/Mirrors, clones, or SnapView snapshots and clones.
 - **CLARiiON** — Uses SnapView clones, or snapshots.
 - **VNX** — Uses SnapView clones and snapshots for VNX block storage, and SnapSure local snapshots and Replicator remote snapshots for VNX file.
 - **Celerra** — Uses SnapSure local snapshots or Celerra Replicator remote snapshots.
 - **VNXe** — Uses SnapSure local snapshots or Replicator remote snapshots.
 - **RecoverPoint** — Create application-consistent and crash-consistent replicas of data residing on VPLEX virtual volumes and protected by RecoverPoint.
 - **VPLEX** — Create application-consistent and crash-consistent replicas of data protected on VPLEX arrays.
 - **VMware** — Create replicas in VMware environments including support for replication, mount, and restore of VMFS, virtual disk, and RDM disks in selected environments.
 - **Hyper-V** — Create replicas in Hyper-V environments including support for replication, mount, and restore of child partitions, pass-through and iSCSI Initiator versions.
 - **VIO LPARs** — Create replicas in VIO LPAR environments including support for replication, mount, and restore of logical partitions.
- ◆ Create remote disk-based replicas of production data on CLARiiON storage using full or incremental SAN Copy
 - ◆ Create remote SnapView Clone and SnapView Snapshot replicas on CLARiiON MirrorView Secondary storage
 - ◆ Create remote BCV replicas on Symmetrix storage
 - ◆ Create remote TimeFinder/Clone and TimeFinder/Snap replicas on Symmetrix storage
 - ◆ Mount replicas to an alternate mount host or back to the production host
 - ◆ Restore production data from disk-based replicas
 - ◆ Run pre-replication, post-replication, backup or callout scripts to automate functions performed on replicas

Audience This guide is intended for use by all user roles to learn the basic operation of Replication Manager.

Readers of this guide are expected to be familiar with the following topics:

- ◆ Operation of application software used in conjunction with Replication Manager
- ◆ Operation of all operating systems on hosts attached to Replication Manager
- ◆ Hardware and software components of the storage arrays that are part of your storage environment
- ◆ Any third-party software that you use with Replication Manager, such as volume managers or virtualization software such as VMware

Organization Here is an overview of where information is located in this guide.

Chapter 1, "Introduction," provides an overview of Replication Manager functionality. This chapter describes some possible ways to use the product, the roles of the users, and what functions each role is authorized to perform.

Chapter 2, "Getting Started," explains some preliminary setup information, how to install and start the Replication Manager Console, the components of the main window, and how to use the Help functionality.

Chapter 3, "Configuring Replications," explains how to create and maintain application sets and jobs to replicate one or more databases or file systems automatically. This chapter also describes how to manage replica rotations and scheduled replication activities.

Chapter 4, "Daily Operations," defines how to perform routine tasks such as running existing jobs, monitoring jobs in progress, examining storage information, unmounting replicas, and cancelling tasks that are already running.

Chapter 5, "Mount, Restore, and Recovery," explains how to mount a replica, restore a replica, recover application data, and use callout scripts to perform additional tasks during the replication process.

Chapter 6, "Using Consistent Split," explains how to use consistent-split technology with Replication Manager to create, mount, and restore replicas. Much of the product functionality differs

when you use consistent split to create a replica. This chapter clarifies those differences.

Chapter 7, "Configuring Federated Data," describes how to define a federated application set, and how to replicate and mount such an application set. This chapter also highlights restrictions you should keep in mind when working with federated application sets.

Chapter 8, "Replica Management," describes how to view replicas, create remote replicas, control replica expiration, automatically rotate replicas, and protect replicas from inadvertent expiration after they have been created.

Appendix A, "Oracle Procedures," describes Replication Manager features that are specific to the Oracle application and managing Oracle data.

Appendix B, "UDB Procedures," describes Replication Manager features that are specific to the UDB application and managing UDB data.

Appendix C, "SQL Server Procedures," describes Replication Manager features that are specific to the Microsoft SQL Server application and managing SQL Server data.

Appendix D, "Microsoft Exchange Procedures," describes Replication Manager features that are specific to Microsoft Exchange applications, including managing Exchange data.

Appendix E, "SharePoint Procedures," describes Replication Manager features that are specific to Microsoft Office SharePoint Server.

Appendix F, "UNIX File System Procedures," describes Replication Manager features that are specific to the replication of UNIX/Linux file systems (Solaris, IBM AIX, HP-UX, and Red Hat) and managing UNIX/Linux file system data.

Appendix G, "NTFS Procedures," describes Replication Manager features that are specific to the replication of NTFS and managing NTFS data.

Related documentation

Related documents include:

- ◆ *EMC Replication Manager Product Guide* (this document) — Provides an overview of the Replication Manager product along with a description of how to perform general tasks once the

Replication Manager product has been installed and configured. This document also provides information specific to Replication Manager's integration with various applications.

- ◆ *EMC Replication Manager Administrator's Guide* — Provides information about installing Replication Manager and configuring the product and related storage services to integrate with one another.
- ◆ *EMC ItemPoint™ for Microsoft® Exchange™ Server User Guide* — Provides an overview of the mailbox recovery software that recovers Microsoft® Exchange Server data.
- ◆ *EMC ItemPoint™ Extract Wizard User Guide* — Provides step-by-step instructions for completing the extraction process of both private and public Exchange Information Store data from tape, disk backups, or other data to any alternate location.
- ◆ *EMC Replication Manager Release Notes* — Provides information about fixed and known defects in the release and also provides information about installation of the release.
- ◆ EMC Replication Manager online help — Provides detailed context-sensitive information about each screen of the product to help customers learn and understand how to use Replication Manager.

Conventions used in this guide

EMC uses the following conventions for notes and caution notices.

Note: A note presents information that is important, but not hazard-related.



CAUTION

A caution contains information essential to avoid damage to the system or equipment. The caution may apply to hardware or software.



WARNING

A warning contains information that when heeded can prevent serious damage or physical injury.

Typographical conventions

EMC uses the following type style conventions in this guide:

Normal	<p>In running text:</p> <ul style="list-style-type: none"> • Interface elements (for example, button names, dialog box names) outside of procedures • Items that user selects outside of procedures • Java classes and interface names • Names of resources, attributes, pools, Boolean expressions, buttons, SQL statements, keywords, clauses, environment variables, filenames, functions, menu names, utilities • Pathnames, URLs, filenames, directory names, computer names, links, groups, service keys, file systems, environment variables (for example, command line and text), notifications
Bold	<ul style="list-style-type: none"> • User actions (what the user clicks, presses, or selects) • Interface elements (button names, dialog box names) • Names of keys, commands, programs, scripts, applications, utilities, processes, notifications, system calls, services, applications, and utilities in text
<i>Italic</i>	<ul style="list-style-type: none"> • Book titles • New terms in text • Emphasis in text
Courier	<ul style="list-style-type: none"> • Prompts • System output • Filenames • Pathnames • URLs • Syntax when shown in command line or other examples
Courier bold	<ul style="list-style-type: none"> • User entry • Options in command-line syntax
<i>Courier italic</i>	<ul style="list-style-type: none"> • Arguments in examples of command-line syntax • Variables in examples of screen or file output • Variables in pathnames
<>	Angle brackets for parameter values (variables) supplied by user.
[]	Square brackets for optional values.
	Vertical bar symbol for alternate selections. The bar means 'or.'
...	Ellipsis for nonessential information omitted from the example.

Where to get help EMC support, product, and licensing information can be obtained as follows.

Product information — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to EMC Online Support (registration required) at:

<https://support.emc.com>

Technical support — For technical support, go to Service Center on EMC Online Support. To open a service request through EMC Online Support, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Your comments Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to techpubcomments@emc.com with your opinions of this guide.

This chapter introduces Replication Manager in the following sections:

- ◆ Replication Manager overview 26
- ◆ Replication Manager product overview 29
- ◆ Replication Manager architecture 32
- ◆ User roles 36

Replication Manager overview

The EMC® Replication Manager product can simplify management of storage replication, integrate with critical business applications, create, mount, and restore point-in-time *replicas* of databases or file systems residing on supported storage arrays, perform automatic discovery of changes to the storage or application environment, and allow you to delegate tasks to appropriate resources. Some important factors that make Replication Manager unique include:

- ◆ Automated management of point-in-time replicas on EMC Symmetrix®, EMC CLARiiON®, EMC VNX™, EMC Celerra®, and EMC VNXe™ storage.
- ◆ Application consistent replication of Microsoft, Oracle, and UDB applications.
- ◆ Reduces or eliminates the need for scripting solutions for replication tasks.
- ◆ Provides a single management console and wizards to simplify replication tasks.
- ◆ Improved recovery and restore features, including application recovery.
- ◆ Integration with physical, VMware, Hyper-V, or IBM AIX VIO virtual environments.

Replicas can be stored on Symmetrix TimeFinder®/Mirrors, Clones; CLARiiON clones or snapshots; VNX SnapView snapshots and clones, Celerra SnapSure™ local snapshots, or EMC Celerra Replicator™ remote snapshots. Replication Manager also supports data using the RecoverPoint Appliance storage service. Replication Manager allows you to perform local and remote replications using TimeFinder, Open Replicator, EMC Symmetrix Remote Data Facility (SRDF®), EMC SAN Copy™, EMC Navisphere®, EMC Celerra iSCSI, Celerra NFS, and/or replicas of EMC MirrorView™/A or MirrorView/S secondaries using EMC SnapView™/Snap and SnapView/Clone replication technologies where they are appropriate.

Determining your replication goals

Replication Manager offers many different ways to create a replica. Some create a replica on the same storage array as the source data. Others create a replica on another array connected using SAN Copy, Celerra Replicator, MirrorView / A, MirrorView / S, Open Replicator, or SRDF. In Celerra environments, you can promote replicas to production or manage planned or unplanned (disaster) failovers of replicated data. This section helps you decide the best technologies to use for replicas based on your goals.

When you create a replica by setting up an application set and a job, there are certain configurable components that alter how the application is quiesced, where the replica resides, what type of storage is used, the duration of the replica, how often a new replica is created and retained, and whether additional processing occurs before, during, or after the replica has been created.

The following sections offer some guidelines that can help you determine the best settings to use when you plan to create replicas for different purposes.

Repurposing

If you plan to use the data in your replica to carry out another function (such as reporting, analysis, or data mining) without impacting the production instance of the application, you should consider using the following settings:

- ◆ Mount the replica to an alternate host.
- ◆ Use a mount recovery method that allows you to start up the application (so you can access the data) in either a read/write or read-only mode.
- ◆ Use post-replication scripts to start the processing that you plan to conduct on the replica.

These settings reduce the stress on the production instance and the production server. They also allow you the freedom to access (and possibly manipulate) the data to complete your repurposing task without affecting your production data.

Backup and recovery

If you plan to use the data in your replica to create a backup and potentially restore that backup at a later time, you should consider using the following settings:

- ◆ If you are using the replica itself as a backup, set the expiration of the replica far enough in the future to provide adequate protection.
- ◆ Consider creating a replica rotation that allows you to preserve a certain number of replicas.
- ◆ If you are backing up to tape, you should mount the replica and use Replication Manager to start a backup script automatically.
- ◆ Use a mount recovery method that is appropriate for the type of backup you are performing.
- ◆ Validate the data (for example, by running **ESEUTIL** on an Exchange replica) before you back up that replica.

Disaster restart

If you plan to use the data in your replica to restart your application after a disaster, you should consider using the following settings:

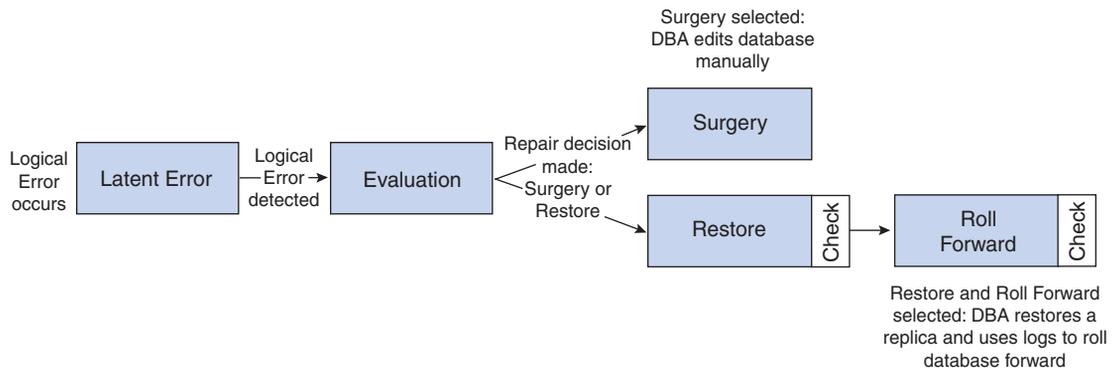
- ◆ Choose a remote replication so that the replicated data is located in a different physical location than the production data.
- ◆ In a Celerra environment, set up your servers for Celerra disaster recovery, as outlined in the administrator's guide. Then you can use the failover feature to startup your environment after a disaster.
- ◆ In VMware SRM environments, set up your environment for disaster recovery as outlined in the administrator's guide. Then you can use failover to change the personality of the environment after a failover has occurred.
- ◆ Mount the replica and verify the validity of the replicated data.
- ◆ Ensure that the replica includes all the data that you will need to restart the application after a disaster occurs.

Copies of replicas

[“Understanding link and copy jobs” on page 137](#) offers specific information on how to create a copy of an existing replica.

Replication Manager product overview

Corrupted databases must be recovered. One option for recovery is to use replicas created using Replication Manager. There are several phases to the recovery process and Replication Manager helps to shorten each of these phases. [Figure 1 on page 29](#) shows the phases of information recovery.



RM-000008

Figure 1 Information recovery process

The following list describes how Replication Manager shortens each of the phases in the data recovery process:

- ◆ **Latent Error phase** — An error occurs in the data, but is not immediately detected. Replication Manager can provide separate replicas to verify the integrity of the data and actively search for errors. Proactive data scrubbing, a process by which Replication Manager creates a point-in-time replica and automated scripts scrub the data to find errors, can reduce or eliminate the latent error phase.
- ◆ **Evaluation phase** — After an error is detected, you must evaluate the data and determine the best way to fix the error. You might choose to perform a *surgical repair*, making manual changes to the database to fix the error. Or, you might decide to *restore* the database from a replica and recover that database by applying the logs. You can shorten the evaluation process by creating a replica of the damaged database and using the replica to perform the evaluation, rather than using the production data.

- ◆ **Surgery phase** — If you decide to perform a surgical repair, you can create a replica of the current data before the surgery. If something goes wrong during the manual database edit, you can restore the replica and attempt the surgery again. Restoring a replica is much faster than restoring from tape after a failed attempt to surgically repair the database. When the system is down, it is important to save as much time as possible.
- ◆ **Restore and Roll Forward phase** — If you decide to restore the data from a replica, you can check each replica and choose the most recent replica that does not have the latent error. After that replica has been restored, use logs to roll the database forward, and then manually restart the database. You cannot perform the same validity check before restoring a tape.

Replication Manager can shorten the overall recovery process. Most other products focus on shortening only the Restore phase, while Replication Manager offers a complete solution that can save time throughout all phases of the recovery process.

Data management with Replication Manager

Replication Manager can help system administrators and Database Administrators (DBAs) set up a data management system within the organization. Because Replication Manager works with different storage arrays, you can produce replicas on several types of storage. The following sections describe your storage options and how to best manage those options to optimize the efficiency of the replicas that you create.

Understanding storage options

Replication Manager can create replicas of data stored on any of the following storage arrays or storage services:

- ◆ Symmetrix series arrays
- ◆ CLARiiON series arrays
- ◆ VNX series arrays
- ◆ Celerra series
- ◆ VNXe series arrays
- ◆ RecoverPoint appliance
- ◆ VPLEX arrays

The *EMC Replication Manager Support Matrix* provides the latest information on the specific models and versions supported by Replication Manager. To access the *EMC Replication Manager Support*

Matrix, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, and scroll down to Replication Manager.

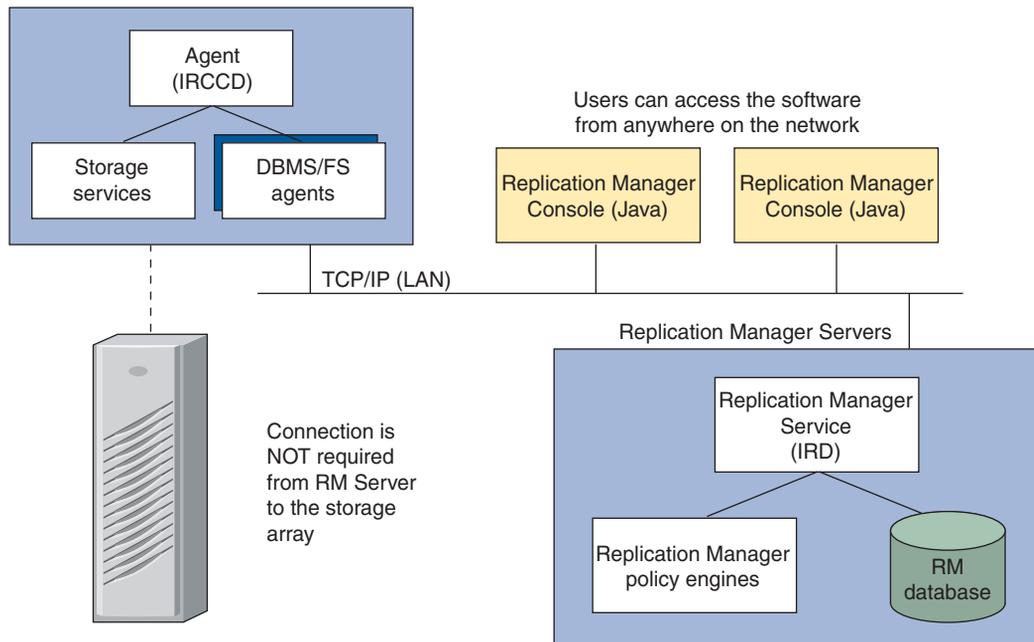
Selecting storage types

Replication Manager gives you the flexibility to store your replicated data on different storage arrays and storage types: Symmetrix TimeFinder/Mirrors (BCVs), TimeFinder/Clones, or TimeFinder/Snaps (VDEVs), Open Replicator; CLARiiON clones or snapshots; VNX SnapView snapshots; Celerra SnapSure local snapshots; and/or RecoverPoint.

In addition, Replication Manager allows you to create replicas of production data or of other existing replicas on a separate storage array using SAN Copy, MirrorView, Celerra Replicator, or VMware SRM functionality.

Replication Manager architecture

Replication Manager uses specialized agent software to communicate with the database or file system that is being replicated. This architecture shown in [Figure 2 on page 32](#) allows Replication Manager to easily support additional databases and file systems as they become available.



RM-000021

Figure 2 Software architecture

Console software

The Replication Manager Console lets you control Replication Manager from Windows systems that have the console installed and have a TCP/IP connection to the server.

The console is a portable Java application that communicates with the Replication Manager Server, over standard TCP/IP sockets. The console is an intuitive product interface that can smoothly integrate the product into the existing storage environment. A command-line

interface installs along with the Replication Manager Console to facilitate scripting.

Server software

The server software controls replication jobs and stores data about each replica. The server software has two distinct components:

- ◆ The *Replication Manager Service (IRD)* controls and coordinates replication and recovery activities for all the storage corresponding to its registered agents and their application sets. The IRD also handles all requests from the console.
- ◆ The *Replication Manager Database* is an embedded data repository that stores data about application sets, jobs, and replicas.

Agent software

The agent software is installed on each host that participates in the replication process, including hosts that manage production data and hosts that are used to mount replicas. The agent software has three distinct components:

- ◆ The Replication Manager Agent (IRCCD) software, which waits for incoming connections from the Replication Manager Server (IRD), and then coordinates all operations on the agent.
- ◆ Storage Services software, which manages the storage relationships between the agent and the storage array where the replica resides.
- ◆ Application Agents for each supported information interface. Each agent is a separate executable, dynamically loaded with Replication Manager at runtime.

Replication Manager agents can install in 64-bit environments. Some Replication Manager agents install as native 64-bit applications, some install 64-bit components with the 32-bit application, and still others install as 32-bit applications, depending upon your environment.

The 32-bit versions of Replication Manager agents require the 32-bit versions of Solutions Enabler¹, and the 64-bit versions of the agents

-
1. Not all configurations require Solutions Enabler. For example, CLARiiON SnapView clone and VNX SnapView replicas, other than replicas of a MirrorView copy, can be created, mounted, restored, and expired without having Solutions enabler or SYMAPI installed on the Windows production, proxy, or mount host.

require the 64-bit versions of Solutions Enabler. When installing Replication Manager agents on Windows Server 2008 or Windows Server 2012 (x64), the native 64-bit agent is installed. Replication Manager can also install 64-bit versions of the Exchange 2007, SQL Server 2005 (x64 or IA64) SQL Server 2008 (x64 or IA64) agents for SRDF/CE and MirrorView/CE environments.

The Replication Manager support information from the E-Lab™ Interoperability Navigator on the EMC Online Support website provides more information.

Database and file system agents

Table 1 on page 34 describes available Replication Manager database and file system agents and the types of information they can replicate.

Note: All replicated data must reside on a supported storage array.

Table 1 Database and file system agents

Agent	Information replicated
Oracle	<ul style="list-style-type: none"> • Oracle information on raw devices or in datafiles • Oracle information in OPS clusters^a (and RAC) • Oracle information using Oracle ASM
UDB	<ul style="list-style-type: none"> • UDB information on raw devices or in datafiles • UDB information in EEE clusters^b
SQL Server	<ul style="list-style-type: none"> • SQL Server 2000/2005/2008 information • SQL Server information in MSCS clusters^a
Exchange	<ul style="list-style-type: none"> • Exchange 2007/2010/2013 information • Exchange information in MSCS clusters^a • Exchange information in Database Availability Groups
SharePoint	<ul style="list-style-type: none"> • SharePoint 2007 and 2010 • SharePoint information in MSCS clusters^a
File system	<ul style="list-style-type: none"> • Data residing on supported file systems. The <i>EMC Replication Manager Support Matrix</i> provides more information on supported file systems. (This agent is always installed when one of the agents listed above is installed.)

^a Data stored in MSCS clusters can be mounted to the passive node of an alternate cluster or to a non-clustered location. OPS is not supported in IBM environments.

^b Replication Manager UDB agent supports EEE in only a single-node instance.

Agent purpose and functions

Each agent provides Replication Manager with a logical view of the data that resides on the storage array; therefore, Replication Manager can:

- ◆ Specify which data to replicate.
- ◆ Ensure that the data can be replicated safely.
- ◆ Quiesce the database.
- ◆ Return the database to normal operation.
- ◆ Recover databases during mounting operations.
- ◆ Shut down or unmount databases during restore operations.

For example, in the case of an Oracle replication, the Oracle agent:

1. Connects to the Oracle database.
2. Obtains a list of the datafiles.
3. Shuts down the database or activates Online Backup mode when you choose the appropriate option.

After replication, the agent:

1. Manages the Oracle control file.
2. Manages the selected initialization files. Including either of these types:
 - text p-file
 - binary sp-file
3. Manages Oracle archived redo log files (or the entire directory).

The appendices at the end of this manual provide detailed information about how each agent works.

User roles

Users must enter a username and password to log in to Replication Manager. The administrator registers each user and assigns one of five different *user roles* to each user, as follows:

- ◆ **Operator** — Has limited privileges to execute existing jobs to which the user has been given access, schedule those jobs, and view properties associated with application sets to which they have been given access.
- ◆ **Power user** — Has all privileges of the Operator role as well as privileges to configure new application sets and jobs.
- ◆ **Database Administrator** — Has all privileges of the Power User role but also has the rights necessary to restore a replica created by a job to which they have access.
- ◆ **Power DBA** — Has all the privileges of the Database Administrator role, but also has rights necessary to configure storage pools.
- ◆ **ERM Administrator** — Can perform every function available within Replication Manager, such as register new hosts, add and modify users, include storage, exclude storage, and access all components without any restrictions.

[Table 2 on page 37](#) defines exactly which functions each role can perform.

Table 2 User roles and allowed actions

Actions / Roles	ERM admin.	Power DBA	DBA	Power user	Operator
Creating New Users	x				
Deleting Users	x				
Registering, Modifying, or Deleting Hosts	x				
Specifying Host Access	x				
Discovering Storage Arrays	x				
Configuring Storage Arrays	x				
Monitoring Snap Cache	x				
Including/Excluding Storage	x				
Discovering Storage Devices	x	x			
Cleaning Up Resources	x	x			
Assigning Storage to Pools	x	x			
Assigning Access to Pools	x	x			
Replica Failover (Celerra)	x	x ^a	x ^a		
Replica Promotion (Celerra)	x	x ^a	x ^a		
Restoring Replicas	x	x ^a	x ^a		
Mounting on Demand ^c	x	x ^a	x ^a	x ^a	
Unmounting on Demand	x	x ^a	x ^a	x ^a	
Deleting Replicas on Demand	x	x ^a	x ^a	x ^a	
Enabling/Disabling Expirations	x	x ^a	x ^a	x ^a	
Configuring Application Set ^c	x	x ^a	x ^a	x ^a	
Configuring Replication Jobs ^c	x	x ^a	x ^a	x ^a	
Running a Job Simulation	x	x ^a	x ^a	x ^a	
Running a Job	x	x ^a	x ^a	x ^a	x ^a
Modifying User Accounts	x	x ^b	x ^b	x ^b	x ^b
Viewing Array Properties	x	x	x	x	x
Viewing Storage	x	x	x ^e	x ^e	x ^e
Viewing Host Properties	x	x ^c	x ^c	x ^c	x ^c
Starting/Cancelling Jobs ^d	x	x ^a	x ^a	x ^a	x ^a
Scheduling Replication Jobs	x	x ^a	x ^a	x ^a	x ^a
Viewing User Properties	x	x	x	x	x

a If given application set access.

b This user can only change their own password.

c Need access to the appropriate hosts to perform the action. (Mount Host for mount on demand, Production Host to configure an Application Set, and Mount and Backup Host to configure a Job.)

d If you can start a job, all the actions that the job can perform are allowed, whether or not you have access to the hosts.

e If given access to the storage through a pool.

This chapter describes how to start using Replication Manager to fulfill all of your replication needs. The chapter includes the following sections:

- ◆ Preliminary setup 40
- ◆ Starting the Replication Manager Console 41
- ◆ Components of the main window 43
- ◆ Getting help 54

Preliminary setup

Before you can start using Replication Manager, you must learn some basics, and your administrator must set up the hardware and software environment to be compatible with Replication Manager.

Preparing the Replication Manager environment

Descriptions of the administrator setup tasks are beyond the scope of this manual; however, the *EMC Replication Manager Administrator's Guide* provides checklists and step-by-step instructions that can help you set up hosts and storage arrays to support Replication Manager.

For specific information about supported applications, operating systems, high-availability environments, volume managers, and other supported software, refer to the *EMC Replication Manager Support Matrix* for the latest support information. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

Replication Manager software components

Replication Manager includes three different software components that all work together to create replicas, mount those replicas, restore replicas, and schedule replication and other tasks. The three components involved are as follows:

- ◆ **Replication Manager Server** — Consists of core software binaries, log files, and an embedded data repository containing configuration data that describes each replica created and the storage associated with that replica. The server component resides on the server host.
- ◆ **Replication Manager Agent** — Consists of an interface to the applications and the storage arrays. The agent is installed on each production host containing data that you want to replicate. Replication Manager includes different kinds of agents designed to create and manipulate replicas for each supported application.
- ◆ **Replication Manager Console** — Consists of software that controls the Replication Manager system from any supported desktop machine or server.

Installing Replication Manager components

For information on installing and/or upgrading the server, agent, and console components, refer to the *EMC Replication Manager Administrator's Guide*.

Starting the Replication Manager Console

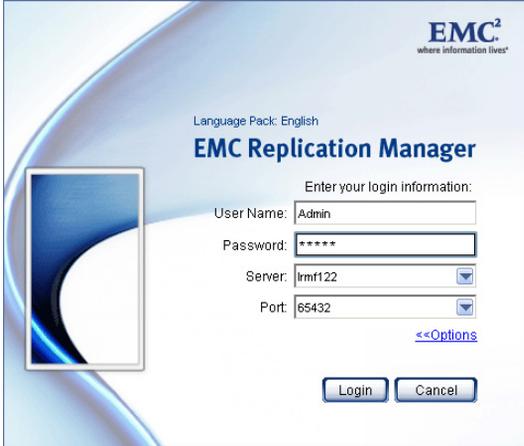
To start the Replication Manager Console:

1. Double-click the Replication Manager Console icon on the desktop, or run **Start > All Programs > Replication Manager > Replication Manager Console**.

Replication Manager first displays the User Login screen shown in [Figure 3 on page 41](#). The screen displays the server and port from your most recent login session by default.

2. Enter the username and password of a Replication Manager user.
3. Choose a Replication Manager Server and port from the list (which displays previously used servers and ports) or type a new server and port number.

Note: The port number here must match the port number the server is using. The default Replication Manager port is 65432. The ERM Administrator chooses the port number during server installation.



The screenshot shows the EMC Replication Manager User Login screen. At the top right is the EMC logo with the tagline "where information lives". Below the logo, it says "Language Pack: English" and "EMC Replication Manager". The main heading is "EMC Replication Manager". Below this, it says "Enter your login information:". There are four input fields: "User Name:" with the value "Admin", "Password:" with masked characters "*****", "Server:" with a dropdown menu showing "lrmf122", and "Port:" with a dropdown menu showing "65432". Below the "Port:" field is a link "<<Options". At the bottom are two buttons: "Login" and "Cancel".

Figure 3 User Login Screen

4. Click **Login**.



CAUTION

Refer to the *EMC Replication Manager Administrator's Guide* for more information about how to assign a password to the Administrator account.

Each time you start Replication Manager, the console ensures that the product resource files are up to date. If the console resources are not current enough to work with the server, Replication Manager displays the following message:

```
This console version x.x is incompatible with the
current version of the Replication Manager Server x.x.
Please update the console.
```

Components of the main window

There are a few distinguishing features of the main Replication Manager window, shown in [Figure 4 on page 43](#) and described on the next few pages.

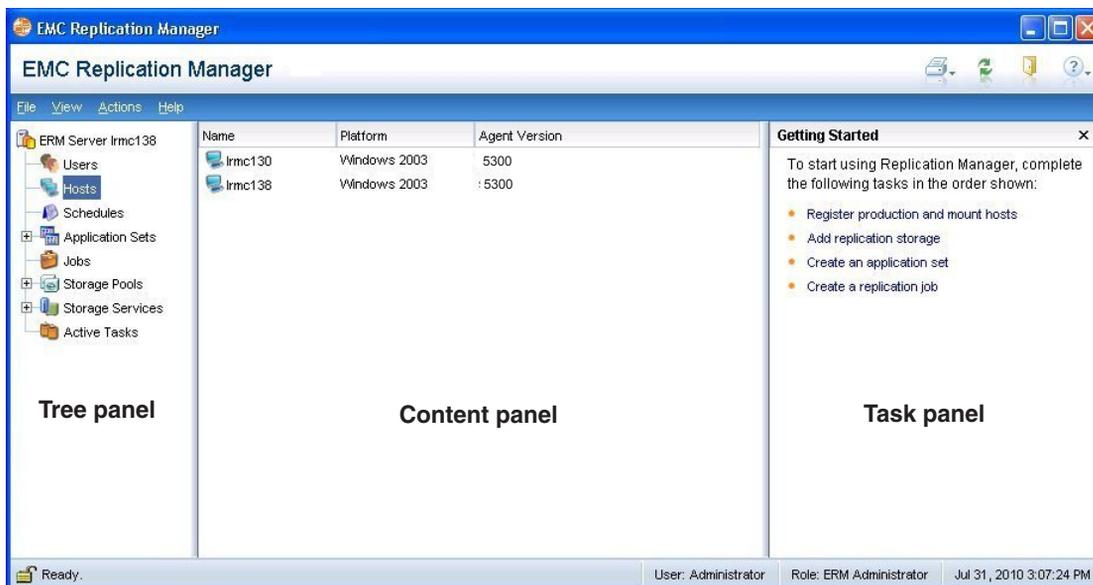


Figure 4 Replication Manager Main Console window

The main window can be separated into distinct parts:

- ◆ **Tree panel** — Left panel containing a folder structure.
- ◆ **Content panel** — Middle panel with data about the selected element.
- ◆ **Task panel** — Right panel with step-by-step tasks to guide the user. This panel can be closed to provide more content panel space.
- ◆ **Menu** — Pull-down menus that can be used to start any task.
- ◆ **Toolbar** — Buttons that can be used to perform certain tasks.
- ◆ **Status bar** — Bar at the bottom of the screen that displays a lock icon to indicate secure communications, status messages, the username of the logged in user, the role of the logged in user, and today's date.

Performing commands on objects

To perform commands on objects in Replication Manager, do one of the following:

- ◆ Right-click the object in the tree panel and select the action from the resulting context menu.

A right-click acts on the object that the mouse pointer is over regardless of whether it is selected. If multiple objects are selected and the mouse pointer is over one of the selected objects, the context menu will contain entries that apply to the topmost selected item and all actions will apply only to the topmost item (with the exception of Delete).

Note: One exception is if you choose Delete from the context menu with multiple items selected. In that case, all items selected will be deleted.

- ◆ Select an object in the tree and use the Actions menu at the top of the screen to select an action.

To get help from the Main window, either click the ? button or select the **Help** menu and then **Contents**.

The status bar displays the following session information:

- ◆ The username of the logged in user, the role of the logged in user, and today's date displayed on the right side of the status bar.
- ◆ The **Padlock** icon on the left is either unlocked or locked. A locked padlock indicates the server is running in Secure Sockets Layer (SSL) mode and the communications between the console and the server are encrypted.
- ◆ The system messages alert you of the current status of the system.

Tree panel

You can right-click items in the tree panel to view context menus containing actions that you can perform on that item, or click an item and select from the Actions menu at the top of the screen.

For example, right-click **Application Sets** to access the Application Set context menu containing commands such as **New Application Set**. Alternatively, you could click **Application Sets**, and then select **New Application Set** from the menu.

Click the (+) symbol next to an object to expand it or the (-) symbol to collapse the object. For example, click the (+) next to **Application Sets** to see a list of application sets.

Click an item in the tree panel to display the children for that item in the content panel.

Each folder in the tree holds a specific category of information:

- ◆ **Users** — Replication Manager accounts for individuals that have access to Replication Manager
- ◆ **Hosts** — Configured production hosts and hosts for mounting replicas
- ◆ **Schedules** — Predetermined time frames when a job runs to create a replica
- ◆ **Application Sets** — Application or file system based sets of information that can be replicated
- ◆ **Jobs** — Sets of actions that are associated with an application set
- ◆ **Storage Pools** — Groups of storage devices that a job can choose from when creating a replica
- ◆ **Storage Services** — Storage arrays, and storage devices visible to Replication Manager
- ◆ **Active Tasks** — Actions currently underway, often due to a user running a job

Content panel

The content panel on the right side of the screen displays the information about the item selected in the tree panel.

Icon descriptions

Replication Manager includes several icons that carry important information. [Table 3 on page 46](#) shows icons used in Replication Manager and describes what they represent.

Table 3 Icon descriptions (page 1 of 8)

Icon	Description	Icon	Description
Toolbar and Status bar icons			
	Print		Refresh console
	Help Toolbar		Exit
	Secure communications (SSL) in use between Replication Manager Server and console		
Tree panel icons			
	Active Tasks folder		Application Sets folder
	Hosts folder		Jobs folder
	Schedules folder		Storage Pools folder
	Storage Services folder		Users folder

Table 3 Icon descriptions (page 2 of 8)

Icon	Description	Icon	Description
	Host that is registered on the Replication Manager Server		Storage pool for disk-based devices (user-defined set of related storage)
	Task that is in progress		
CLARiiON storage icons			
	CLARiiON device		CLARiiON device that is in use
	CLARiiON Snapshot (cache-based) session		CLARiiON storage array that is ready for use with Replication Manager
	CLARiiON Snap Cache		CLARiiON thin LUN
	CLARiiON thin LUNs folder		CLARiiON thin LUN in use
	Pool LUNs		
VNX storage icons			
	VNX device		VNX device that is in use
	VNX Snapshot (cache-based) session		VNX storage array that is ready for use with Replication Manager

Table 3 Icon descriptions (page 3 of 8)

Icon	Description	Icon	Description
	VNX Snap Cache		VNX thin LUN
	VNX thin LUNs folder		VNX thin LUN in use
	Pool LUNs		
Celerra storage icons			
	Celerra device		Celerra device that is in use
	Celerra SnapSure snapshot		Celerra
VNXe storage icons			
	VNXe device		VNXe device that is in use
	VNX SnapSure snapshot		VNXe
Symmetrix storage icons			
	Symmetrix Standard that has been added		Symmetrix Standard that is in use

Table 3 Icon descriptions (page 4 of 8)

Icon	Description	Icon	Description
	Symmetrix BCV that has been added		Symmetrix BCV that is in use
	Symmetrix VDEV Snapshot session		Symmetrix storage array that is ready for use with Replication Manager
	Symmetrix VDEVs Snapshot session (locked)		Thin STD
	Thin STD in use		Thin R1
	Thin R1 in use		

Table 3 Icon descriptions (page 5 of 8)

Icon	Description	Icon	Description
RecoverPoint storage icon			
	RecoverPoint RPA		
Application Set icons			
	Application Set (defines the data to replicate)		Application Set with Unacknowledged Error (red)
	Application Set with Unacknowledged Warning (yellow)		
Replica icons			
	Replica created from an Application Set		Replica that has been mounted
	Replica that has failed		Replica that is in progress
	Replica mount failed		Replica that is invalid
	Simulated Replica		Simulated Replica (Failed)
	Simulated Replica (In Progress)		Replica snapshot (A Replica snapshot may use CLARiiON SnapView, SnapView snapshots, Symmetrix VDEVs, or Celerra SnapSure technology)

Table 3 Icon descriptions (page 6 of 8)

Icon	Description	Icon	Description
	Replica snapshot that is invalid		Replica snapshot that has failed
	Replica snapshot that is in progress		Replica snapshot that has been mounted
	Replica snapshot mount failed		Remote connection to data outside of a given storage array
	Application-consistent RecoverPoint replica		Application-consistent RecoverPoint replica in progress
	Application-consistent RecoverPoint replica failed		Application-consistent RecoverPoint replica mounted
	Application-consistent RecoverPoint replica mount failed		Crash-consistent RecoverPoint replica
	Crash-consistent RecoverPoint replica in progress		Crash-consistent RecoverPoint replica failed
	Crash-consistent RecoverPoint replica mounted		Crash-consistent RecoverPoint replica mount failed
Job Wizard icons			
	Volumes		File Systems

Table 3 Icon descriptions (page 7 of 8)

Icon	Description	Icon	Description
	Datacenter		Datastore
	Oracle Instance		Oracle or SQL Server Database
	SAP environment (SAP BRbackup compliant settings)		
	UDB Instance		UDB Database
	Exchange Instance		SQL Server Instance
	SharePoint Farm Instance		Consistency Method settings
	General settings		Topology View
	Script settings		Fail Job on Error settings
Validation status icons			
	Success		Error

Table 3 Icon descriptions (page 8 of 8)

Icon	Description	Icon	Description
	Warning		Informational
	Validated with warning		
Miscellaneous other icons			
	Folder that contains one or more subcomponents		Valid License
	Expired License		Unlicensed array
	Replica has Unacknowledged Error (red)		Replica has Unacknowledged Warning (yellow)
	Passive node		Print
	Print preview		

Getting help

If you need more help using or troubleshooting Replication Manager, you can check one of these two sources:

- ◆ Online help
- ◆ Product log files

The following sections describe these information sources.

Online help

To get help from the Main window:

- ◆ Click the ? button.
Alternatively, select **Help**, then **Contents** from the menu.
- ◆ For help within a dialog box, click **Help**.
- ◆ For assistance on the Help window itself, choose the topic **Using Help**.

Product log files

You can find important system information in the log files. You can define the log size and the maximum storage allocated to the log directories in the Server Properties and Register New Hosts screens. Refer to the *EMC Replication Manager Administrator's Guide* for more information.

Replication Manager creates log files for the server and agent components of the product as shown in [Figure 5 on page 55](#). The logging options are:

- ◆ **Normal** — Creates logs that describe Replication Manager events (set your logging level to **Normal** in the Server Options or Register New Hosts screen).
- ◆ **Debug** — Creates more verbose logs that describe events and detailed errors or warnings (set your logging level to **Debug** in the Server or Register New Hosts screen).

The default logging level is Debug.

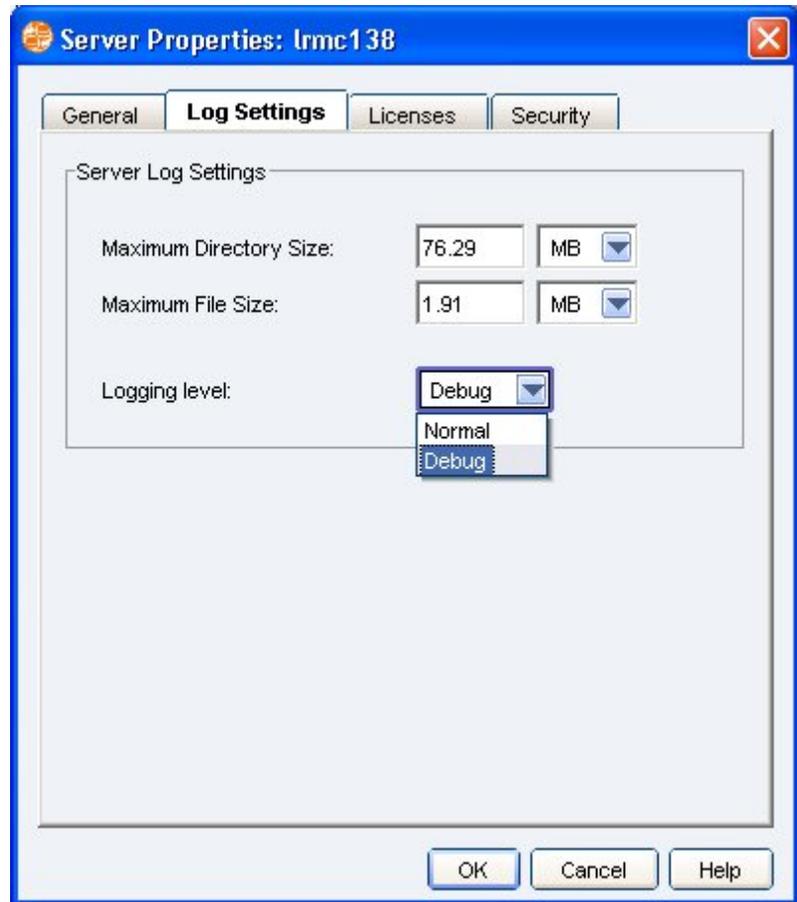


Figure 5 Setting the Replication Manager Server logging level

The filename of each log file is as follows:

`erm_<component><date/time>_<logType>.log`

where:

`<component>` is either server or client (agent).

`<date/time>` is the date and time when the log was started.

`<logType>` is summary or debug.

These log files are stored in the following locations (assuming the product was installed in the default location):

For UNIX installations, logs reside in:

`/opt/emc/rm/logs/client` (agent logs)

For Windows installations, logs reside in:

`C:\Program Files\emc\rm\logs\server` (server logs)

`C:\Program Files\emc\rm\logs\client` (agent logs)

or

`C:\Program Files(x86)\emc\rm\logs\server` (server logs)

`C:\Program Files(x86)\emc\rm\logs\client` (agent logs)

Note: For more information about logging, refer to the *EMC Replication Manager Administrator's Guide*.

This chapter describes how to configure replications and includes the following topics:

- ◆ Available replication technologies 58
- ◆ Number of replicas supported 110
- ◆ Managing application sets 111
- ◆ Managing jobs 124
- ◆ Defining a new job 126
- ◆ Understanding link and copy jobs 137
- ◆ Creating replicas of RecoverPoint targets 141
- ◆ Managing existing jobs..... 141
- ◆ Managing replica rotations 150
- ◆ Managing job schedules 152

Available replication technologies

This section outlines each replication technology that can be used in conjunction with Replication Manager in detail. This section includes the following information for each technology:

- ◆ **Storage Services** — Describes the storage service(s) that must exist in your environment in order to use the given replication technology.
- ◆ **Source** — Describes the valid source storage that can be used with this replication technology.
- ◆ **Target** — Describes the target storage where the replica resides when you select this replication technology.
- ◆ **Storage Requirements** — Describes what storage devices must be available or included in the storage pool you selected when you created the job. This item also describes other storage pool restrictions if they exist.
- ◆ **Mount and Recovery** — Describes whether you can mount and recover data that has been replicated using this replication technology. This item also describes other mount and recovery restrictions.

See each of the following sections for specifics.

Symmetrix TimeFinder/Clone

TimeFinder/Clone technology is controlled via copy sessions, which pair the source and target devices to create copies of a source device on multiple target devices. The target devices can be either standard (STD or STD/R1) or BCV devices as long as they are all the same size and emulation type. Thin STD and Thin STD/R1 are also supported. TimeFinder/Clones have the following characteristics:

- ◆ **Storage Service** — Symmetrix storage arrays.
- ◆ **Source** — Standards (STDs), TimeFinder/Clone, TimeFinder/Mirror BCV (non-emulation mode), R1 from SRDF/A or SRDF/S. Replication Manager supports replication of thin devices with this technology. Thin sources replicate to thin targets while traditional sources replicate to traditional targets.
- ◆ **Target** — Standards (STDs, Thin STDs), BCVs or Thin BCVs using TimeFinder/Clone technology.

- ◆ **Storage Requirements** — Include enough STDs, Thin STDs or BCVs, Thin BCVs to support the source STDs, Thin STDs or BCVs, Thin BCVs in the replica. Clone copies of striped or concatenated meta devices can also be created. Replication Manager supports this by selecting copy devices that are identical in stripe count, stripe size, and capacity.

There are several device type considerations when using BCV, Thin BCV and STD, Thin STD devices to create TimeFinder/Clone replicas. Refer to the *EMC Solutions Enabler Symmetrix TimeFinder API Programmers Manual* for more detailed information.

Creating a TimeFinder/Clone of another TimeFinder/Clone (also known as a cascade) is supported. At Symmetrix Microcode levels of 5874 or greater, the original TimeFinder/Clone session will not be terminated when the cascade is created. Replication Manager does not support a three-level cascade however. If a replica would create a three-level cascade, the second clone session will be terminated. Additionally, federated application sets cannot support the creation of a TimeFinder/Clone of another TimeFinder/Clone.

- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the replica if necessary. For SRDF restore, the RDF link must be suspended.
- ◆ **Uses** — Clone copies are useful in situations where multiple copies of production data are needed for testing, backups, or report generation.

Figure 6 on page 59 provides a graphical representation of this environment.



Figure 6 TimeFinder/Clone environment

Symmetrix TimeFinder/Clone (remote)

Remote TimeFinder/Clones are similar to local TimeFinder/Clones, except that the replication occurs on a remote Symmetrix. Remote TimeFinder/Clones have the following characteristics:

- ◆ **Storage Service** — Local to remote Symmetrix storage arrays.
- ◆ **Source** — SRDF/S R2 (VMAX), SRDF/A R2 (VMAX) or SRDFe/S R2 (VMAXe) device.

Note: Set the environment variable ERM_ALLOW_SRDF_A to 1 to enable SRDF/A R2 (VMAX).

- ◆ **Target** — STD/BCV for R2 traditional STD/BCV or Thin STD/Thin BCV for R2 Thin STD/Thin BCV (using TimeFinder/Clone technology) located on a remote Symmetrix array (connected to the source array using SRDF/S).
- ◆ **Storage Requirements** — Include enough remote STDs or BCVs to support the source STDs in the replica.

Creating a TimeFinder/Clone of another TimeFinder/Clone (also known as a cascade) is supported. At Symmetrix Microcode levels of 5874 or greater, the original TimeFinder/Clone session will not be terminated when the cascade is created. Replication Manager does not support a three-level cascade however. If a replica would create a three-level cascade, the second clone session will be terminated. Additionally, federated application sets cannot support the creation of a TimeFinder/Clone of another TimeFinder/Clone.

- ◆ **Mount** — Can mount the replica, but restore is not supported.

Figure 7 on page 60 provides a graphical representation of this environment.

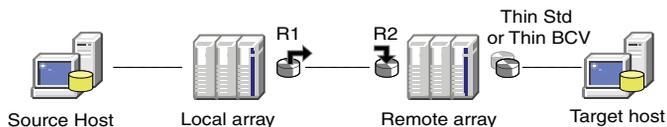


Figure 7

Remote TimeFinder/Clone environment

Symmetrix TimeFinder/Mirror

TimeFinder/Mirror technology creates an exact copy of the data by synchronizing and splitting a BCV. TimeFinder/Mirrors have the following characteristics:

- ◆ **Storage Service** — Symmetrix storage arrays running Symmetrix Microcode prior to version 5874.
- ◆ **Source** — Standards (STDs, Thin STD) RAID5/6 and non-RAID5/6, R1 from SRDF/A or SRDF/S.
- ◆ **Target** — BCVs, Thin BCV (RAID5 and non-RAID5) using TimeFinder/Mirror technology.
- ◆ **Storage Requirements** — Include enough BCVs to support the standards to be included in the replica.

Creating a TimeFinder/Snap or TimeFinder/Clone of a TimeFinder/Mirror is supported when TimeFinder/Mirror emulation mode is not being used. Emulation mode is used for RAID5/RAID6 BCVs, in which case TimeFinder/Snap or TimeFinder/Clone of a TimeFinder/Mirror is not supported.

- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from that replica if necessary. For SRDF restore, RDF link must be suspended.
- ◆ **Uses** — TimeFinder/Mirrors (BCVs) should be used for long-term storage of your most critical production data.

Figure 8 on page 61 provides a graphical representation of this environment.

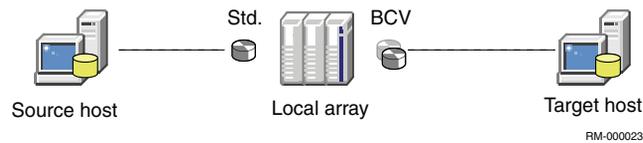


Figure 8

TimeFinder/Mirror environment

Symmetrix TimeFinder/Mirror (remote)

Remote TimeFinder/Mirrors are similar to local TimeFinder/Mirrors, except that the replication occurs on a remote Symmetrix. Remote TimeFinder/Mirrors have the following characteristics:

- ◆ **Storage Service** — Local to remote Symmetrix storage arrays running Symmetrix Microcode prior to version 5874.
- ◆ **Source** — R2 device.
- ◆ **Target** — BCV/Thin BCV of an R2 device located on a single remote Symmetrix array (connected to the source array using SRDF/S).

- ◆ **Storage Requirements** — Include enough remote BCVs, Thin BCVs to support the R1s in the replica. Replication Manager cannot restore these replicas.

Creating a TimeFinder/Snap or TimeFinder/Clone of a TimeFinder/Mirror is supported when TimeFinder/Mirror emulation mode is not being used. Emulation mode is used for RAID5/RAID6 BCVs/Thin BCVs, in which case TimeFinder/Snap or TimeFinder/Clone of a TimeFinder/Mirror is not supported.

- ◆ **Mount** — Can mount the replica, but the replica cannot be restored.

Figure 9 on page 62 provides a graphical representation of this environment.

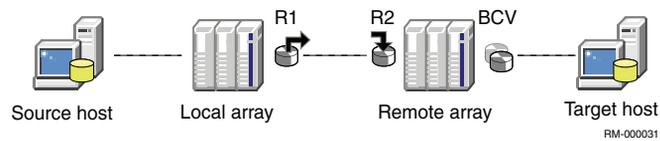


Figure 9 Remote TimeFinder/Mirror environment

Symmetrix TimeFinder/Snaps

TimeFinder/Snaps use space-saving, pointer-based (copy on first write) snapshots. These snapshots are called virtual devices (VDEVs). TimeFinder/Snaps have the following characteristics:

- ◆ **Storage Service** — Symmetrix storage arrays.
- ◆ **Source** — Standards (STDs, Thin STDs) RAID5 and non-RAID5, TimeFinder/Clone, TimeFinder/Mirror, or R1s from SRDF/A or SRDF/S, R2s from SRDF/A (VMAX) or BCVs, Thin BCVs (RAID5 and non-RAID5).

Replication Manager supports replication of thin devices with this technology.

Note: Set the environment variable ERM_ALLOW_SRDF_A to 1 to enable SRDF/A R2 (VMAX).

- ◆ **Target** — Virtual Devices (VDEVs) (using TimeFinder/Snap technology). Target save devices for the snap sessions must reside either in a snap pool that is of the type “default”, or in customized save pools. Also, all the devices in the snap replica chain need to have target save devices from the same pool.
- ◆ **Storage Requirements** — Bind enough save devices to a pool to support the source devices. Creating a TimeFinder/Snap of a TimeFinder/Clone is supported. TimeFinder/Snaps of TimeFinder/Clones no longer terminate the TimeFinder/Clone session if the Symmetrix microcode supports maintaining that session while the snap exists. Federated application sets cannot support the creation of a TimeFinder/Snap of a TimeFinder/Clone and also TimeFinder/Snap of a TimeFinder/Snap.

TimeFinder/Snap of TimeFinder/Clone is not supported if the TimeFinder/Clone is in a Restored/Split state. If that is the case, the replication operation will fail and the clone session needs to be split manually or clone replica must be recreated before creating the snap.

- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the replica if necessary. For SRDF restore, RDF link must be suspended and restore takes place from the source side (the R1), not the remote side (the R2). Restore from a VDEV replica that is mounted read-only is not supported.

- ◆ **Uses** — TimeFinder/Snaps (VDEVs) should be used for short-term storage and working copies of data.

Figure 10 on page 64 provides a graphical representation of this environment.

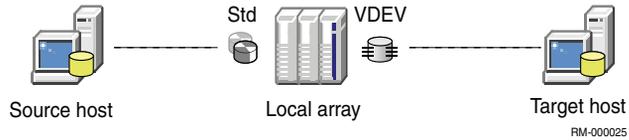


Figure 10 TimeFinder/Snap environment

Customized save pools for TimeFinder/Snaps

Customized save pools can be created on the Symmetrix storage arrays and can be used by the target devices to store the changes made in the source devices. The Symmetrix administrator creates customized save pools, which can be selected by a Replication Manager user for a TimeFinder/Snap job.

If the save pool to be used is specified through the environmental variable, `EMC_ERM_SAVEPOOL_FILE`, then this save pool takes precedence over the save pool specified in Replication Manager console.

To be able to use the customized save pools, make sure the following minimum criteria are met, as shown in [Table 4 on page 64](#):

Table 4 Requirements for customized save pools

Requirement	Version
Solution Enabler	6.0 or higher
Microcode version of Symmetrix array	5x71

Note: After upgrading the Replication Manager client and server, run the storage discovery to be able to view and use the customized save pools.

Device pacing for SRDF/A R2 VP Snap

SRDF/A write pacing extends the availability of SRDF/A by preventing conditions that result in cache overflow on both the R1 and R2 sides. Solutions Enabler provides two types of write pacing: group-level write pacing and device-level write pacing. Both types of write pacing can be activated for an SRDF/A session at the same time.

Here is sample content that you can use to enable the device pacing. In this example, the file is named `srdf_file.cmd`.

```
set rdf group 19 rdfa_dse_pool=DSE_372, emulation=fba;
set rdf group 19 rdfa_dse_threshold=20;
set rdf group 19 rdfa_wpace_autostart = ENABLE;
set rdf group 19 rdfa_devpace_autostart = ENABLE
```

Run the command:

```
symconfigure commit -sid 12 -file srdf_file.cmd.
```

Symmetrix TimeFinder/ Duplicate Snaps

TimeFinder/Duplicate Snaps are a Symmetrix technology that creates a new virtual device (VDEV) from an existing VDEV target. TimeFinder/Duplicate Snaps can only be created in Replication Manager using a copy job to create the duplicate snap from an existing VDEV that is the target of an existing replica.

TimeFinder/Duplicate Snaps have the following characteristics:

- ◆ **Storage Service** — Symmetrix storage arrays.
- ◆ **Source** — Virtual devices (VDEVs) targets created by a previous TimeFinder/Snap replication or a previous TimeFinder/Duplicate Snap copy job. multiple Duplicate snap copy jobs can be chained together as Symmetrix storage requirements and limitations are met.

This type of replica uses TimeFinder/Duplicate Snap functionality available with Symmetrix Microcode levels of 5875 or greater. Replication Manager also supports replication of thin devices (TDEVs) with this technology.

- ◆ **Target** — VDEVs (using TimeFinder/Duplicate Snap technology to create a duplicate of an existing VDEV). Target save devices for the snap sessions must reside in a pool that is of the type “default” and all the devices in the snap replica chain need to have target save devices from the same pool.
- ◆ **Storage requirements** — Bind enough save devices to a pool to support all the virtual devices in the chain.
- ◆ **Mount and recovery** — Mounts and database recovery from the replica are supported. For SRDF restore, RDF link must be suspended and restore takes place from the source side (the R1), not the remote side (the R2). Restore from a VDEV replica that is mounted read-only is not supported.

Restore is supported and restores data to the source of the first VDEV in the chain, if the original source is one of these: Standard (STD, Thin STD), TimeFinder/Mirror, R1s from SRDF/A or SRDF/S, or BCVs, Thin BCVs.

If the original source is a TimeFinder/Clone, restore is not supported.

- ◆ **Uses** — TimeFinder/Duplicate Snap (VDEVs) should be used for short-term storage and working copies of data.

Figure 11 on page 66 provides a graphical representation of this environment.

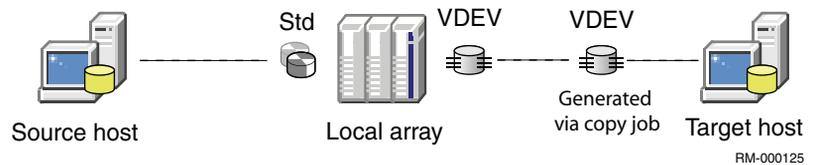


Figure 11 TimeFinder/Duplicate Snap environment

Symmetrix TimeFinder/Snaps (remote)

Remote TimeFinder/Snaps are similar to local TimeFinder/Snaps, except that the replication occurs on a remote Symmetrix. Remote TimeFinder/Snaps have the following characteristics:

- ◆ **Storage Service** — Local to remote Symmetrix storage arrays.
- ◆ **Source** — R2s located on a local Symmetrix array.
- ◆ **Target** — Virtual Devices (VDEVs) (using TimeFinder/Snap technology) of an R2 device located on a single remote Symmetrix array (connected to the source array using SRDF/S). Federated application sets are not supported.
- ◆ **Storage Requirements** — Bind enough save devices to a pool to support the source devices. Creating a TimeFinder/Snap of a TimeFinder/Clone is supported. TimeFinder/Snaps of TimeFinder/Clones no longer terminate the TimeFinder/Clone session if the Symmetrix microcode supports maintaining that session while the snap exists. Federated application sets cannot support the creation of a TimeFinder/Snap of a TimeFinder/Clone.

TimeFinder/Snap of TimeFinder/Clone is not supported if the TimeFinder/Clone is in a Restored/Split state. If that is the case, the replication operation will fail and the clone session needs to be split manually or clone replica must be recreated before creating the snap.

- ◆ **Mount and Recovery** — Can mount the replica but cannot restore.

Figure 12 on page 67 provides a graphical representation of this environment.

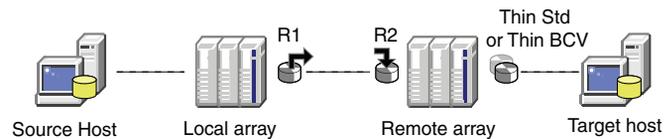


Figure 12 Remote TimeFinder/Snap environment

Symmetrix Open Replicator (Symmetrix to CLARiiON or VNX)

Replication Manager can create EMC Open Replicator Hot Push sessions from Symmetrix control devices (source devices on the Symmetrix) to remote target devices on a CLARiiON or VNX array.

The environment must have the following characteristics:

- ◆ **Storage Service** — Symmetrix to a remote CLARiiON or VNX array.
- ◆ **Source and Target** — Creates a replica of any of the following sources to the corresponding targets:
 - **Source** — Control devices may be any of the following: Standards (STDs), Thin STDs or R1s from SRDF/A or SRDF/S, or mounted TimeFinder clone replicas (refer to “Using a mounted TimeFinder Clone as a source” on page 68).
 - **Target** — May be any of the following:
 - Pool LUNs
 - CLARiiON or VNX LUNs
 - CLARiiON or VNXe iSCSI LUNs
 - CLARiiON or VNX thin LUNs

If thin LUNs are used as a target the benefit of the thin LUN will be negated because the LUN is fully populated on the first transfer from the source.

- ◆ **Storage Requirements** — The following storage requirements are in effect on each source Symmetrix array for Open Replicator replication:
 - Open Replicator must be installed on each source Symmetrix array from which Open Replicator replications are created.
 - Replication Manager can use existing Open Replicator sessions or new sessions.
 - Source and target devices must match in size exactly to use Open Replicator in conjunction with Replication Manager.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the replica if necessary.

Figure 13 on page 68 provides a graphical representation of this environment.

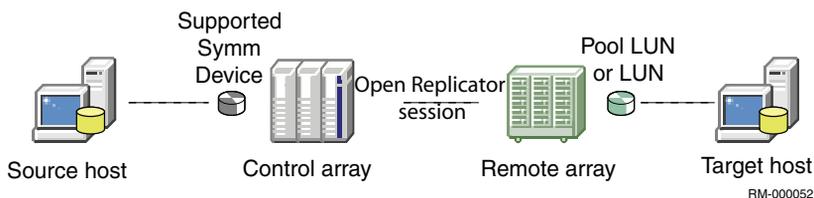


Figure 13 Open Replicator environment (Differential Hot Push only)

Using a mounted TimeFinder Clone as a source

If you are using a mounted TimeFinder clone as a source for an Open Replicator session, note the following restrictions about the TimeFinder clones:

- ◆ Must not have long duration locking enabled.
- ◆ Cannot be mounted as read-only.
- ◆ Cannot be sourced from BCVs.

EMC does not recommend restoring TimeFinder clones which are used as sources for Open Replicator sessions.

CLARiiON SnapView clone

CLARiiON SnapView clones create an exact copy of the data onto a separate LUN or disk. SnapView clones have the following characteristics:

- ◆ **Storage Service** — CLARiiON storage arrays.

- ◆ **Source** — LUNs (including Thin LUNs) or Open Replicator targets using a copy job to create a clone.
- ◆ **Target** — LUNs (including Thin LUNs) using CLARiiON SnapView technology.
- ◆ **Storage Requirements** — Include enough LUNs to support the source LUNs to be included in the replica. If thin LUNs are used, for optimal resource utilization, make sure you have enough traditional and thin LUNs to allow replications to use the same type of LUN as the source.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the replica if necessary.
- ◆ **Uses** — CLARiiON SnapView clones should be used for data that changes a great deal in a short time or storage of more critical data.

Figure 14 on page 69 provides a graphical representation of this environment.

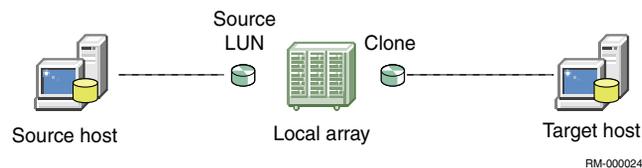


Figure 14 SnapView clone environment

VNX SnapView clone

VNX SnapView clones create an exact copy of the data onto a separate LUN or disk. SnapView clones have the following characteristics:

- ◆ **Storage Service** — VNX storage arrays.
- ◆ **Source** — LUNs (including Thin LUNs) or Open Replicator targets using a copy job to create a clone.
- ◆ **Target** — LUNs (including Thin LUNs) using VNX SnapView technology.
- ◆ **Storage Requirements** — Include enough LUNs to support the source LUNs to be included in the replica. If thin LUNs are used, for optimal resource utilization, make sure you have enough traditional and thin LUNs to allow replications to use the same type of LUN as the source.

- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the replica if necessary.
- ◆ **Uses** — VNX SnapView clones should be used for data that changes a great deal in a short time or storage of more critical data.

CLARiiON SnapView clone (remote)

Remote SnapView clones have the following characteristics:

- ◆ **Storage Service** — Local to remote CLARiiON storage arrays.
- ◆ **Source** — MirrorView /S or MirrorView /A secondary devices (a single application set cannot contain both MirrorView /S and MirrorView /A), or RecoverPoint CDP or CRR target devices.
- ◆ **Target** — Clone LUNs (created using CLARiiON SnapView technology) of a MirrorView /S or MirrorView /A secondary device.

Note: When configuring MirrorView /S and MirrorView /A secondaries, Replication Manager can create snaps or clones of the secondary only when there is a single secondary image per primary volume.

- ◆ **Storage Requirements** — Include enough remote LUNs to support the source LUNs to be included in the replica. The MirrorView link must be established prior to Replication Manager processing. That link can be established manually using Navisphere. MirrorView /A must be set to manual mode to allow Replication Manager to create application consistent replicas in a MirrorView /A environment.
- ◆ **Mount** — Can mount the replica, but the replica cannot be restored to the secondary. It can however be restored to the primary using SAN Copy.

Figure 15 on page 70 provides a graphical representation of this environment.

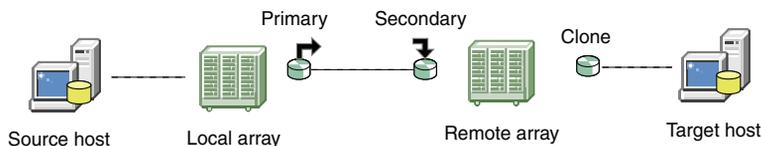


Figure 15 Remote SnapView clone environment

Note: If you are running Replication Manager using native (no SYMAPI) on Windows, you will not see the MirrorView targets offered as an option for the replication source, even if a mirror is correctly configured.

VNX SnapView clone (remote)

Remote SnapView clones have the following characteristics:

- ◆ **Storage Service** — Local to remote VNX storage arrays.
- ◆ **Source** — MirrorView/S or MirrorView/A secondary devices (a single application set cannot contain both MirrorView/S and MirrorView/A), or RecoverPoint CDP or CRR target devices.
- ◆ **Target** — Clone LUNs (created using VNX SnapView technology) of a MirrorView/S or MirrorView/A secondary device.

Note: When configuring MirrorView/S and MirrorView/A secondaries, Replication Manager can create snaps or clones of the secondary only when there is a single secondary image per primary volume.

- ◆ **Storage Requirements** — Include enough remote LUNs to support the source LUNs to be included in the replica. The MirrorView link must be established prior to Replication Manager processing. That link can be established manually using Navisphere. MirrorView/A must be set to manual mode to allow Replication Manager to create application consistent replicas in a MirrorView/A environment.
- ◆ **Mount** — Can mount the replica, but the replica cannot be restored to the secondary. It can however be restored to the primary using SAN Copy.

CLARiiON SnapView snap

CLARiiON SnapView replicas create a copy on first write to disk-based cache memory. Snapshots store only the information from changed tracks, so they use a minimum of cache storage space on the CLARiiON array. SnapView snaps have the following characteristics:

- ◆ **Storage Service** — CLARiiON storage arrays.
- ◆ **Source** — LUNs or SnapView Clone LUNs using a copy job to create a snap of a clone or Open Replicator targets using a copy job to create a snap.
- ◆ **Target** — SnapView snaps (using CLARiiON SnapView technology).

- ◆ **Storage Requirements** — Include enough snap cache space to support the source LUNs to be included in the replica. If a snap is removed due to insufficient snap cache, the replica that includes that snap becomes invalid and cannot be used. Snap cache does not have to be included in the storage pool to be used.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the snaps.
- ◆ **Uses** — SnapView replicas can be used effectively for data storage of data that changes little over time and short-term working copies of the data.

Figure 16 on page 72 provides a graphical representation of this environment.

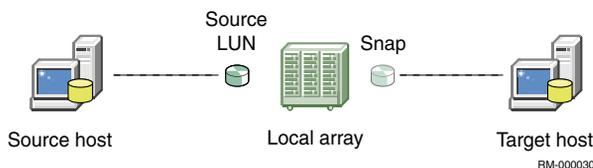


Figure 16 SnapView snap environment

VNX SnapView snap

VNX SnapView replicas create a copy on first write to disk-based cache memory. Snapshots store only the information from changed tracks, so they use a minimum of cache storage space on the VNX array. SnapView snaps have the following characteristics:

- ◆ **Storage Service** — VNX storage arrays.
- ◆ **Source** — LUNs or SnapView Clone LUNs using a copy job to create a snap of a clone or Open Replicator targets using a copy job to create a snap.
- ◆ **Target** — SnapView snaps (using VNX SnapView technology).
- ◆ **Storage Requirements** — Include enough snap cache space to support the source LUNs to be included in the replica. If a snap is removed due to insufficient snap cache, the replica that includes that snap becomes invalid and cannot be used. Snap cache does not have to be included in the storage pool to be used.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform database recovery from the snaps.

- ◆ **Uses** — SnapView replicas can be used effectively for data storage of data that changes little over time and short-term working copies of the data.

CLARiiON SnapView snap (remote)

Remote SnapView snaps have the following characteristics:

- ◆ **Storage Service** — CLARiiON storage arrays.
- ◆ **Source** — MirrorView/S or MirrorView/A primary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices), or RecoverPoint CDP or CRR target devices.
- ◆ **Target** — SnapView snaps (using CLARiiON SnapView technology) of a MirrorView secondary device. (Targets that use multiple secondaries are not supported.)
- ◆ **Storage Requirements** — Include enough snap cache space to support the source LUNs to be included in the replica. If a snap is removed due to insufficient snap cache, the replica that includes that snap becomes invalid and cannot be used. Snap cache does not have to be included in the storage pool to be used. MirrorView/A must be set to manual mode to allow Replication Manager to create application consistent replicas in a MirrorView/A environment.
- ◆ **Mount** — Can mount the replica, but the replica cannot be restored to the secondary. It can however be restored to the primary using SAN Copy.

Figure 17 on page 73 provides a graphical representation of this environment.

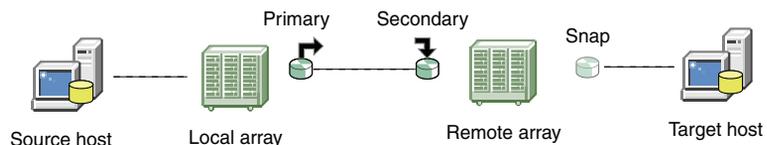


Figure 17

Remote SnapView snap environment

VNX SnapView snapshots (remote)

Remote SnapView snapshots have the following characteristics:

- ◆ **Storage Service** — VNX storage arrays.

- ◆ **Source** — MirrorView/S or MirrorView/A primary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices), or RecoverPoint CDP or CRR target devices.

CLARiiON ATA disk support

Another storage option that Replication Manager supports is the CLARiiON CX series with low-cost Advanced Technology-Attached (ATA) disks. These drives can reduce the cost per megabyte of the storage, and therefore reduce your overall storage investment.

CLARiiON and VNX MirrorView coexistence

Replication Manager supports coexistence with MirrorView/S and MirrorView/A source LUNs. You can take replicas of the LUNs that have MirrorView sessions established, but you cannot restore while the session is in place. You will have to fracture the secondary image of the mirror to be able to do the restore.

Full SAN Copy

Full SAN Copy has the following characteristics:

- ◆ **Storage Services** — Can include any of the following storage services combinations:
 - Symmetrix to remote CLARiiON or VNX clone shown in [Figure 18 on page 75](#)
 - CLARiiON LUN to remote CLARiiON or VNX clone shown in [Figure 19 on page 76](#)
 - CLARiiON Snapshot to remote CLARiiON or VNX clone shown in [Figure 20 on page 76](#)
 - CLARiiON LUN to local CLARiiON or VNX clone shown in [Figure 21 on page 76](#)
- ◆ **Source and Target** — Creates a replica of any of the following sources to the corresponding targets:
 - **Source** — May be any of the following: Symmetrix STDs, Thin STDs, CLARiiON or VNX LUNs, CLARiiON or Snapview snapshots or thin LUNs
 - **Target** — May be any of the following:
 - CLARiiON snapshots and VNX Snapview snapshots
 - CLARiiON or VNX LUNs or thin LUNs local (in-frame).

- Remote LUNs or remote thin LUNs (out-of-frame) (using EMC SAN Copy technology) Target volumes must be equal in size to the source volumes.

Note: thin LUN SAN Copy requires a minimum of EMC FLARE® 29 at both ends of the link.

- ◆ **Storage Requirements** — Include temporary storage areas on the source array and final storage for the replica on the target array. If you are using storage pools the requirements differ as follows depending upon which storage services you are using:

- For Symmetrix to CLARiiON or VNX Full SAN Copy, the storage pool must include enough VDEVs or BCVs on the source array and enough clones on the target array to support the source standards to be included in the replica. The VDEVs or BCVs used as temporary storage must be set up for SAN Copy to the CLARiiON or VNX.

The procedure for setting up these devices is outlined in the *EMC Replication Manager Administrator's Guide*.

- For CLARiiON to CLARiiON Full SAN Copy (or VNX to VNX Full SAN Copy), the storage pool must include enough clones and thin LUNs on the target array to support the source LUNs to be included in the replica.

The local array can be a Symmetrix, CLARiiON or a VNX, but the remote array must be a CLARiiON or VNX.

- ◆ **Mount and recovery** — Can mount the replica on a target host and/or perform direct database recovery from the target to the source.

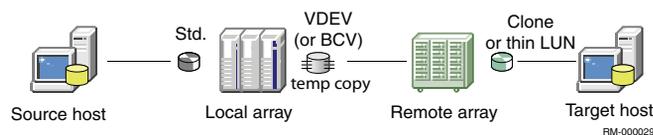


Figure 18 Symmetrix to remote CLARiiON or VNX clone

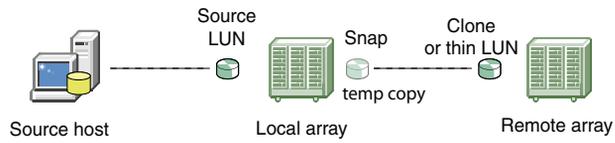


Figure 19 CLARiiON LUN to remote CLARiiON or VNX clone

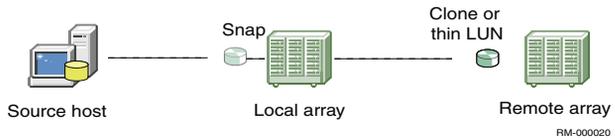


Figure 20 CLARiiON Snapshot to remote CLARiiON or VNX clone

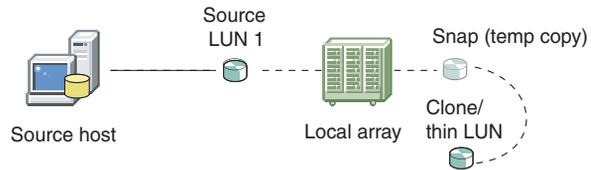


Figure 21 CLARiiON Snapshot to local CLARiiON or VNX clone

Incremental SAN Copy

Incremental SAN Copy has the following characteristics:

- ◆ **Storage Service** — CLARiiON or VNX local storage array in frame or to a CLARiiON or VNX remote array.
- ◆ **Source** — LUNs or thin LUNs in a local CLARiiON or VNX array.
- ◆ **Target** — LUNs or thin LUNs in a remote CLARiiON or VNX array. Target volumes must be equal in size to the source volumes.

Note: thin LUN SAN Copy requires a minimum of FLARE 29 at both ends of the link.

- ◆ **Storage Requirements** — Include enough clones on the remote array to support the source LUNs to be included in the replica.

Note: Replication Manager does not automatically create a snapshot of that clone. To create a separate snapshot of the clone, users must create a copy job to take a snapshot of the resulting clone. Refer to [“Understanding link and copy jobs” on page 137](#) for more information.

- ◆ **Mount and Recovery** — Can mount the replica on a target host and/or perform direct database recovery from the target to the source.

Note: When you perform a mount on Windows, Replication Manager mounts read-only. Choose **Create and mount a snap of the replica** to cause changes you make to the mounted replica to be discarded on unmount.

[Figure 22 on page 77](#) provides a graphical representation of this environment.

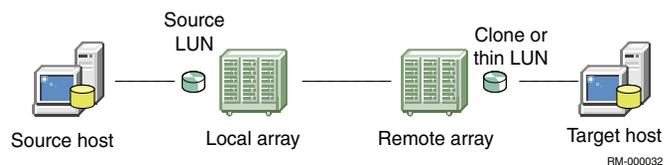


Figure 22 Incremental SAN Copy

Celerra SnapSure

Celerra SnapSure creates a point-in-time copy of all the data on the source LUN. For the initial snapshot, this method creates a full copy of the original LUN, therefore requiring the same amount of space on the file system as the LUN. The file system can be either a local file system or a network file system. Subsequent snapshots space usage depends on how much the data has changed since the last snapshot was taken. Celerra SnapSure has the following characteristics:

- ◆ **Storage Service** — Celerra File Server.
- ◆ **Source** — Celerra LUN or NFS.
- ◆ **Target** — Celerra SnapSure local snapshot.
- ◆ **Storage Requirements** — The following storage requirements apply:
 - The source data must reside on Celerra source LUNs or NFS.
 - Storage must include enough space for the snapshots on the Celerra.
 - If the source data for an application set is a network file system, it must all reside on a single Celerra.
 - Storage pools cannot be defined for Celerra jobs.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and/or perform direct recovery from target to source.

Figure 23 on page 78 provides a graphical representation of this environment.

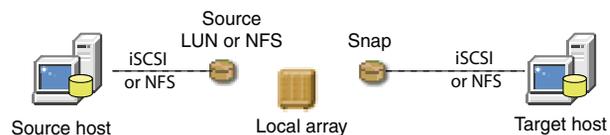


Figure 23 Celerra SnapSure local snapshot

VNXe SnapSure

VNXe SnapSure creates a point-in-time copy of all the data on the source LUN. For the initial snapshot, this method creates a full copy of the original LUN, therefore requiring the same amount of space on the file system as the LUN. Subsequent snapshots space usage depends on how much the data has changed since the last snapshot was taken. SnapSure has the following characteristics:

- ◆ **Storage Service** — VNXe File Server.
- ◆ **Source** — VNXe LUN.
- ◆ **Target** — VNXe SnapSure local snapshot.
- ◆ **Storage Requirements** — The following storage requirements apply:
 - The source data must reside on VNXe source LUNs.
 - Storage must include enough space for the snapshots on the VNXe.
 - Storage pools cannot be defined for VNXe jobs.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and/or perform direct recovery from target to source.

Figure 24 on page 79 provides a graphical representation of this environment.

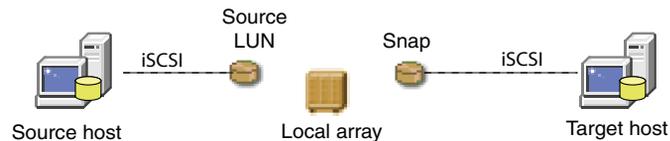


Figure 24 VNXe SnapSure local snapshot

Celerra and VNX Replicator

Celerra and VNX Replicator creates a point-in-time copy of all the data on the source LUN (iSCSI) or network file system (NFS). This replication method creates a full copy of the original LUN or file system on the target as a clone then creates a snapshot of that clone, which becomes the replica.

When replicating file systems, The file system can be either a local file system or a network file system. Subsequent replicas update the target clone incrementally, then create a new snapshot of that clone for the subsequent replicas in the rotation.

Note: For iSCSI application sets, set up a rotation with two or more replicas to take advantage of incremental copies, since a rotation of one results in a full copy every time the replica is run, because the delete of the previous replica for the rotation would also delete the Replicator session.

Celerra Replicator and VNX Replicator have the following characteristics:

- ◆ **Storage Service** — Celerra or VNX File Server.
- ◆ **Source** — Celerra or VNX LUN or network file system.
- ◆ **Target** — Replicator remote snapshot.
- ◆ **Storage Requirements** — The following storage requirements apply:
 - The source data must reside on Celerra source LUNs or network file systems.
 - Storage must include enough space for the snapshots on the Celerra iSCSI or VNXe iSCSI.
 - If the source data for an application set is a network file system it must all reside on a single Celerra.
 - Storage pools cannot be defined for iSCSI Celerra jobs.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform recovery from the replica if necessary.

Figure 25 on page 81 provides a graphical representation of this environment.

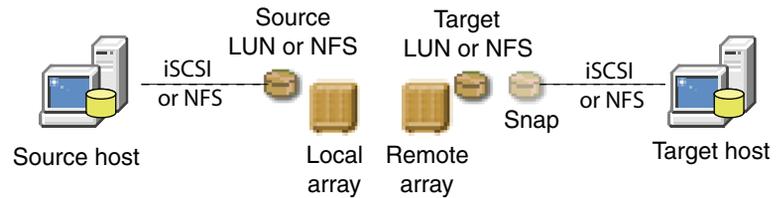


Figure 25 Replicator remote snapshot

VNXe Replicator

VNXe Replicator creates a point-in-time copy of all the data on the source LUN (iSCSI). This replication method creates a full copy of the original LUN on the target as a clone then creates a snapshot of that clone, which becomes the replica.

Note: For iSCSI application sets, set up a rotation with two or more replicas to take advantage of incremental copies, since a rotation of one results in a full copy every time the replica is run, because the delete of the previous replica for the rotation would also delete the VNXe Replicator session.

VNXe Replicator has the following characteristics:

- ◆ **Storage Service** — VNXe File Server.
- ◆ **Source** — VNXe LUN.
- ◆ **Target** — VNXe Replicator remote snapshot.
- ◆ **Storage Requirements** — The following storage requirements apply:
 - The source data must reside on VNXe source LUNs.
 - Storage must include enough space for the snapshots on the VNXe (iSCSI only).
 - Storage pools cannot be defined for iSCSI VNXe jobs.
- ◆ **Mount and Recovery** — Can mount the replica on a target host and perform recovery from the replica if necessary.

Figure 26 on page 82 provides a graphical representation of this environment.

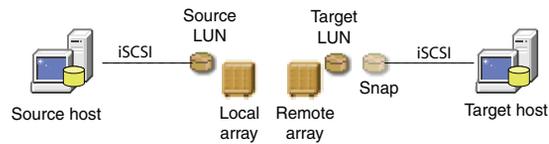


Figure 26 VNXe Replicator remote snapshot

RecoverPoint

Data protection with RecoverPoint has the following characteristics:

- ◆ **Storage Service** — RecoverPoint Appliance (RPA).
- ◆ **Source** — Symmetrix STDs, CLARiiON and VNX LUNs, VPLEX virtual volumes.
- ◆ **Target** — Symmetrix STDs, CLARiiON and VNX LUNs, VPLEX virtual volumes.
- ◆ **Storage Requirements** — RecoverPoint consistency groups contain replication sets which match each volume to be replicated with a volume at the other copy (or copies). RecoverPoint also requires repository volumes and journal volumes as described in the RecoverPoint documentation.
- ◆ **Replica types** — Application-consistent replicas (also known as specific-point-in-time) are created by Replication Manager jobs. Crash-consistent replicas (also known as any-point-in-time) are created using the mount or restore wizards.
- ◆ **Mount and Recovery** — One mount of a RecoverPoint replica at a time per application set, per site (local site or remote site).

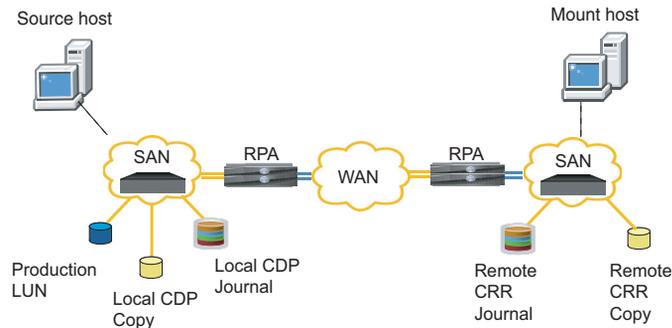


Figure 27 RecoverPoint environment

CLARiiON SnapView replicas of RecoverPoint targets

SnapView replicas of RecoverPoint target LUNs can be used for an extended period, such as for repurposing or backup. By creating a snap or clone of the remote RecoverPoint target, you remove the need to mount the CDP or CRR copy itself. This is important because any writes that occur to a production device when a CDP or CRR copy is mounted also are made to the journal devices, potentially blocking writes to the production LUNs if activity is high enough.

You can create the replica using a single two-phase job, or by creating two jobs, one being a copy job.

- ◆ In a single two-phase job, Replication Manager makes a RecoverPoint bookmark to use as an intermediate, temporary copy that is used as the source for the SnapView replica.
- ◆ In a copy job, you use a replica from a separate RecoverPoint CDP, CRR, or CLR bookmark job, then create a SnapView job that uses the RecoverPoint job as the source. The RecoverPoint source job should not include a read-write mount of the RecoverPoint target. Note that with the copy job method, the replica created by the RecoverPoint job is not removed.

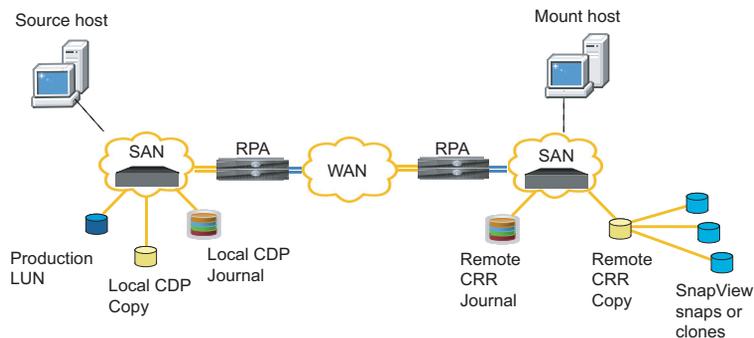


Figure 28 SnapView snap or clones of RecoverPoint target

Requirements for SnapView of RecoverPoint targets

The RecoverPoint target must be on CLARiiON or VNX storage. (The RecoverPoint chapter in the *Replication Manager Administrator's Guide* describes general RecoverPoint configuration details required by Replication Manager.)

SnapView replication of RecoverPoint targets is supported on all platforms, and with all applications, supported by Replication Manager, except RecoverPoint/CE and SRM environments are not supported.

Symmetrix TimeFinder replicas of RecoverPoint targets

Restoring from the replicas of the RecoverPoint targets is not supported.

TimeFinder replicas of RecoverPoint target LUNs can be used for an extended period, such as for repurposing or backup. By creating a snap or clone of the remote RecoverPoint target, you remove the need to mount the CDP or CRR copy itself. This is important because any writes that occur to a production device when a CDP or CRR copy is mounted also are made to the journal devices, potentially blocking writes to the production LUNs if activity is high enough.

You can create the replica using a single two-phase job, or by creating two jobs, one being a copy job.

- ◆ In a single two-phase job, Replication Manager makes a RecoverPoint bookmark to use as an intermediate, temporary copy that is used as the source for the TimeFinder replica.
- ◆ In a copy job, you use a replica from a separate RecoverPoint CDP, CRR, or CLR bookmark job, then create a TimeFinder job that uses the RecoverPoint job as the source. The RecoverPoint source job should not include a read-write mount of the RecoverPoint target. Note that with the copy job method, the replica created by the RecoverPoint job is not removed

Requirements for TimeFinder replica of RecoverPoint targets

The RecoverPoint target must be on Symmetrix storage. (The RecoverPoint chapter in the *Replication Manager Administrator's Guide* describes general RecoverPoint configuration details required by Replication Manager.)

TimeFinder replication of RecoverPoint targets is supported on all platforms, and with all applications, supported by Replication Manager, except RecoverPoint/CE and SRM environments are not supported.

Restoring from the replicas of the RecoverPoint targets is not supported.

VMware support

VMware VMFS or NFS datastores

Replication Manager supports replication, mount, and restore activities in the following VMware configurations.

Replication Manager offers data protection of VMware VMFS or NFS datastores hosting Windows or Linux virtual machines. These virtual environments have the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services:
 - Celerra File Server (iSCSI)
 - Celerra File Server to remote Celerra File Server (iSCSI)
 - Celerra File Server (NFS datastore)
 - Celerra File Server to remote Celerra File Server (NFS datastore)
 - VNXe File Server (iSCSI)
 - VNXe File Server to remote VNXe File Server (iSCSI)
 - VNX File Server (NFS datastore)
 - CLARiiON local clones
 - CLARiiON SnapView snaps
 - CLARiiON-to-CLARiiON SAN Copy
 - CLARiiON-to-VNX SAN Copy
 - VNX local clones
 - VNX SnapView snaps
 - VNX-to-VNX SAN Copy
 - VNX-to-CLARiiON SAN Copy
 - Symmetrix TimeFinder/Clone
 - Symmetrix TimeFinder/Clone Remote
 - Symmetrix TimeFinder/Mirror
 - Symmetrix TimeFinder/Mirror Remote
 - Symmetrix TimeFinder/Snap (VDEVs)
 - Symmetrix TimeFinder/Snap (VDEVs) remote
 - RecoverPoint appliance

- ◆ **Source** — VMware VMFS or NFS datastore within an ESX Server stored on one of the following source devices:
 - Celerra iSCSI LUNs (source for a Celerra SnapSure local replica)
 - Celerra SnapSure local snapshot of iSCSI LUNs (source for a Replicator remote copy replica)
 - Celerra file systems (source for an NFS datastore replication)

Note: For NFS datastores, the entire NFS on Celerra has to be exported and used on the ESX Server; partitioned NFS is not allowed for VMware datastores residing on Celerra.

- CLARiiON LUNs, SnapView Snapshots, MirrorView /S or MirrorView /A primary devices, MirrorView /S or MirrorView /A secondary devices (however a single application set cannot contain both MirrorView /S and MirrorView /A primary devices)
- VNX LUNS, VNX SnapView Snapshots, MirrorView /S or MirrorView /A primary devices, MirrorView /S or MirrorView /A secondary devices (however a single application set cannot contain both MirrorView /S and MirrorView /A primary devices)
- Symmetrix Standards (STDs), TimeFinder /Mirror BCV (non-RAID5), TimeFinder /Clone BCV, R1 from SRDF /A or SRDF /S
- ◆ **Target** — VMware VMFS or NFS mounts listed separately below:

Targets for NFS datastores (replica can be mounted to a Linux server or to an ESX Server)

- SnapSure local snapshot
- Replicator remote snapshot

Targets for VMware VMFS datastores (replica can be mounted to an ESX Server)

- SnapSure local snapshot
- Replicator remote snapshot
- CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)

- VNX LUNS, VNX SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)...
- NFS mount on a Linux server
- Symmetrix STDs, BCVs or BCV of an R2 (depending upon the selected underlying replication technology)
- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply (see sections above). In addition to the following VMware requirements:
 - VMFS requirements:
 - VMware VirtualCenter (also known as vCenter) must be installed. VirtualCenter manages all VMware operations.
 - A Windows-based VMware proxy agent (physical or virtual host) must be identified. The proxy agent has Replication Manager agent software installed and has visibility to the VirtualCenter.
 - NFS datastore requirements:
 - A Windows/Linux-based VMware proxy agent (physical or virtual host) must be identified. The proxy agent has Replication Manager agent software installed and has visibility to the VirtualCenter. This host must be configured for DNS.
 - Host must have the Celerra Control Station discovered under storage discovery in order to use NFS datastores
- ◆ **Mount and Recovery** — After a mount, it is the user's responsibility to bring individual virtual machines online.

For Celerra or VNX environments, these additional requirements apply:

- For VMFS mounts:
 - The ESX Server used for the mount operation must be connected to the target Data Movers IP address using the dynamic discovery tab.
 - The target IQN of the mount must have at least one placeholder LUN masked to the ESX Server.
- For NFS datastore mounts:
 - NFS datastores can be mounted to Linux hosts (physical or virtual) or to an ESX Server.

- The mount host name should be typed if mount host is ESX Server or selected from a list box if the mount host is a Linux host.

Refer to [Figure 29](#) on page 88 through [Figure 38](#) on page 91 for graphical representations of these configurations.

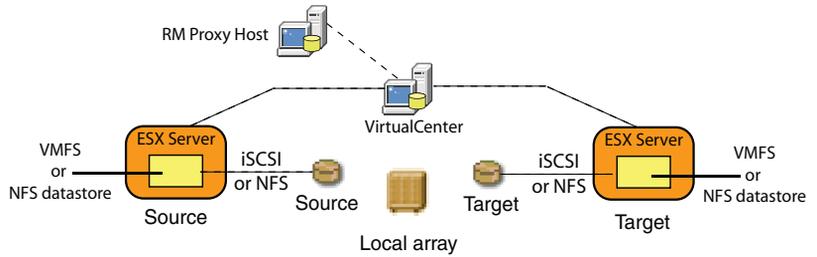


Figure 29 VMFS or NFS datastore on Celerra or VNX File (local)

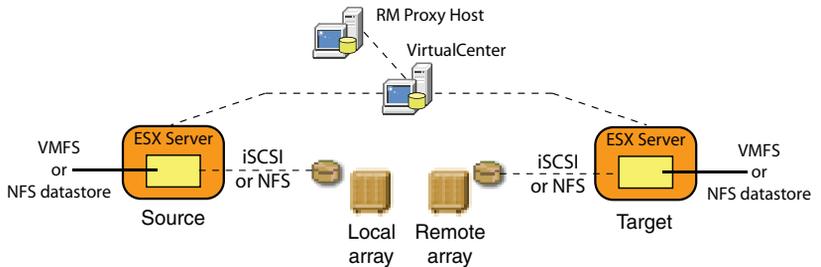


Figure 30 VMFS or NFS datastore on Celerra or VNX File (remote)

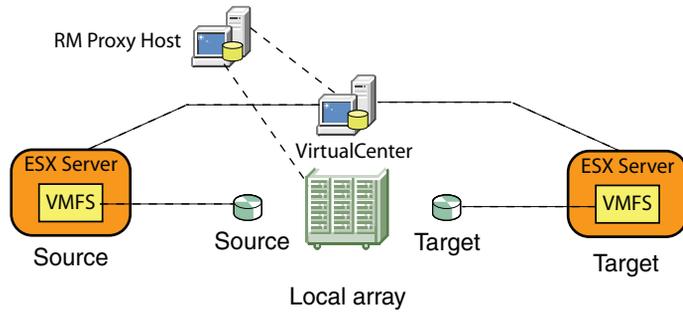


Figure 31 VMFS on CLARiiON or VNX local

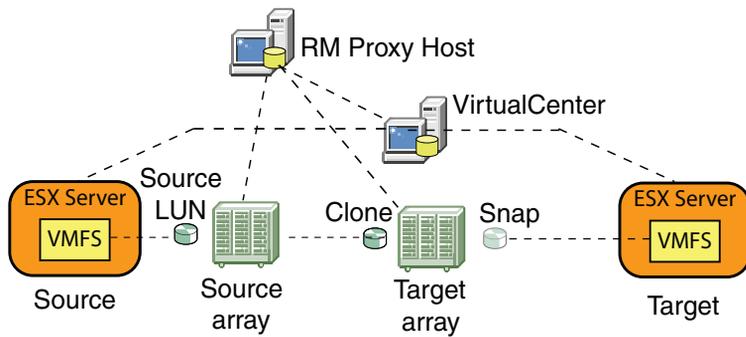


Figure 32 VMFS on CLARiiON or VNX using SAN Copy replication

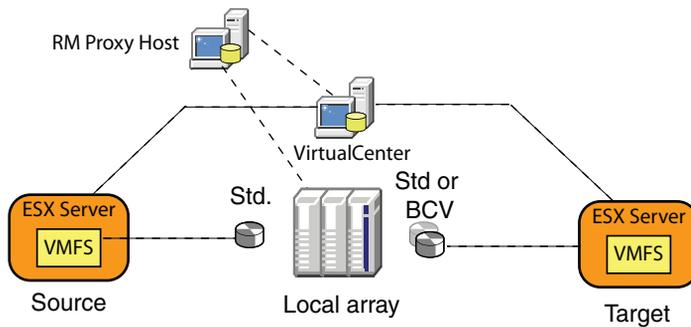


Figure 33 VMFS on Symmetrix TimeFinder/Clones

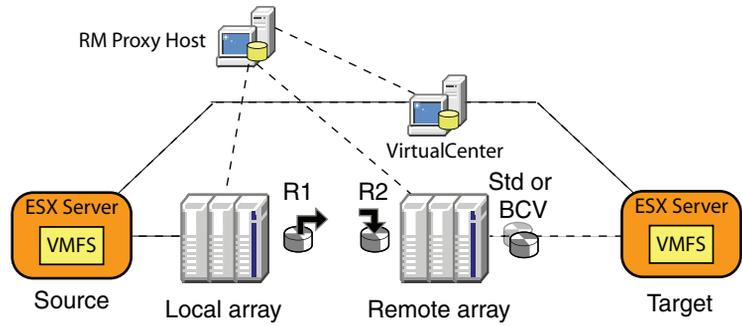


Figure 34 VMFS on Symmetrix TimeFinder/Clones (remote)

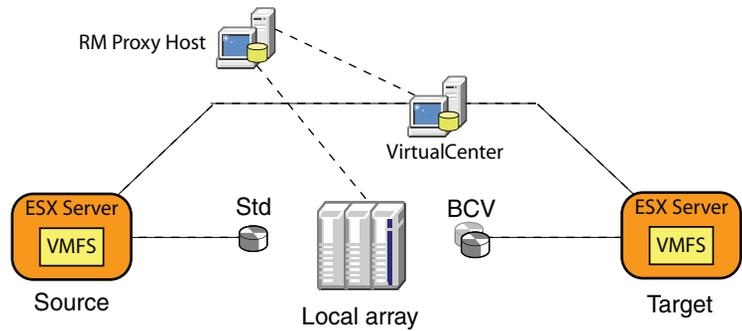


Figure 35 VMFS on Symmetrix TimeFinder/Mirror

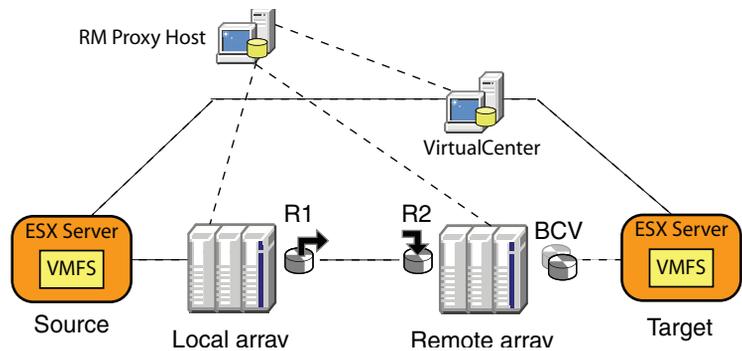


Figure 36 VMFS on Symmetrix TimeFinder/Mirror (remote)

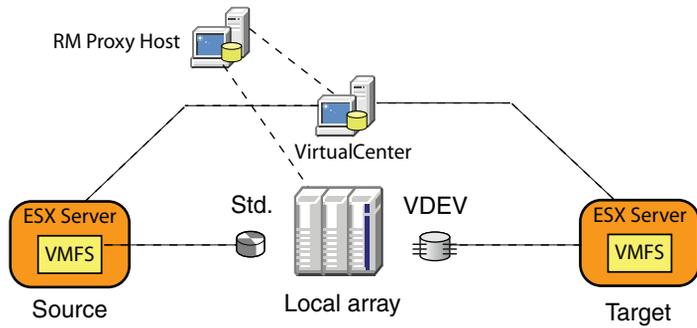


Figure 37 VMFS on Symmetrix TimeFinder/Snap

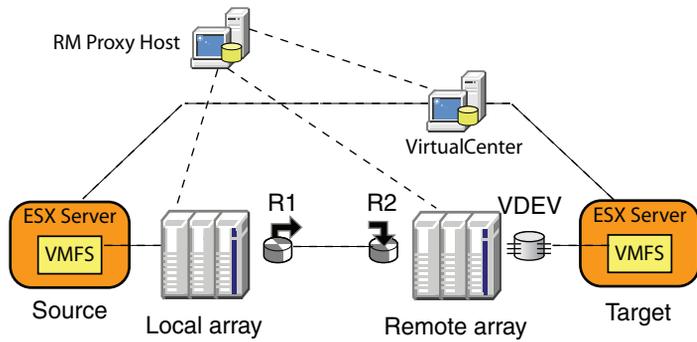


Figure 38 VMFS on Symmetrix TimeFinder/Snap (remote)

VMware virtual disk

Data protection of VMware virtual disk (in Windows environments) has the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services:
 - Celerra File Server (iSCSI) or VNXe (iSCSI)
 - CLARiiON storage array (including CLARiiON SnapView support, CLARiiON-to-CLARiiON Full SAN Copy support, or CLARiiON-to-CLARiiON Incremental SAN Copy support)
 - VNX storage array (including VNX Snapshot support, VNX-to-VNX Full SAN Copy support, or VNX-to-VNX Incremental SAN Copy support)
 - Symmetrix storage array
- ◆ **Source** — VMware VMFS within an ESX Server stored on one of the following source devices:
 - Celerra iSCSI or VNXe iSCSI LUNs (source for a Celerra SnapSure local replica)
 - Celerra SnapSure local snapshot (source for a Replicator remote copy replica)
 - CLARiiON LUNs, SnapView Snapshots, or MirrorView/S or MirrorView/A primary devices or MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
 - VNX LUNs, SnapView snapshots, or MirrorView/S or MirrorView/A primary devices or MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
 - Symmetrix Standards (STDs), TimeFinder/Mirror BCV (non-RAID5), TimeFinder/Clone BCV, R1 from SRDF/A or SRDF/S
- ◆ **Target** — VMware VMFS created on an ESX Server with one of the following target devices:
 - Celerra SnapSure local snapshot
 - Replicator remote snapshot
 - CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)

- VNX LUNs, VNX SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
- Symmetrix STDs, VDEVs, BCVs or BCVs of an R2 (depending upon the selected underlying replication technology)
- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply (see sections above). In addition to the following VMware requirements:
 - In SCSI environments, Replication Manager can only replicate environments with one unique SCSI target across all SCSI controllers on the virtual machine, that is the SCSI target of the virtual disk being replicated must not be used on other SCSI controllers.

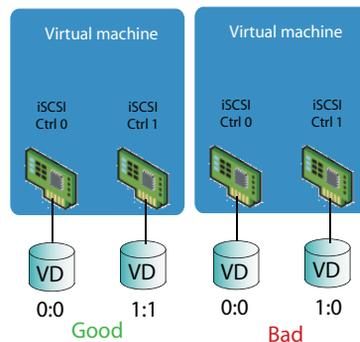


Figure 39 Unique target requirement across controllers

The number before the colon in [Figure 39 on page 93](#) is the controller; the number after the colon is the target. Target must be unique and not used on more than one controller as shown in the second configuration.

- VMware Tools must be installed in this environment. “Installing VMware Tools on the virtual machine” in the *EMC Replication Manager Administrator’s Guide* describes how to install VMware Tools.

One of the following requirements apply for remote CLARiiON or VNX replication:

- (For remote CLARiiON or VNX replication) at least one Replication Manager client, residing on a virtual machine, must have a virtual disk on a VMFS created using a LUN from the remote CLARiiON or VNX array.
- (For remote CLARiiON or VNX replication) at least one Replication Manager client, residing on a virtual machine, must have a RDM device configured from the remote CLARiiON or VNX array.

The following requirements apply for Symmetrix virtual disk environments:

- Each virtual machine must be zoned with access to a Symmetrix gatekeeper device.
- The gatekeeper should be available to the virtual machine as an RDM device in physical compatibility mode.
- ◆ **Mount and Recovery** — During mount of a virtual disk replica, choose a virtual machine as the **Mount host**. No proxy host is needed for virtual disk replicas.

The Keep LUNs Visible option is not supported if you are mounting snap replicas of the VMware environment.

For Celerra iSCSI or VNXe iSCSI environments, these additional requirements apply:

- The ESX Server used for the mount operation must be connected to the target iSCSI servers IP address using the dynamic discovery tab.
- The target IQN of the mount must have at least one placeholder LUN masked to the ESX Server.
- In Symmetrix environments, replica storage must be made visible to the ESX Servers that host the mount VM prior to running the mount job from Replication Manager.

Refer to [Figure 40 on page 95](#) through [Figure 48 on page 97](#) for graphical representations of these configurations.

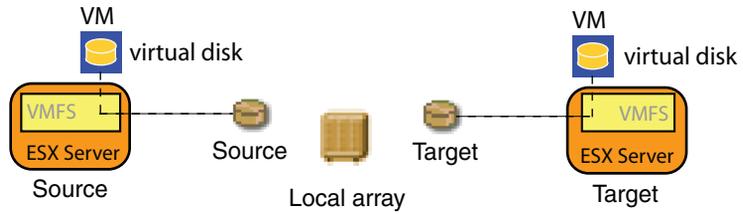


Figure 40 Virtual Disk on Celerra (local)

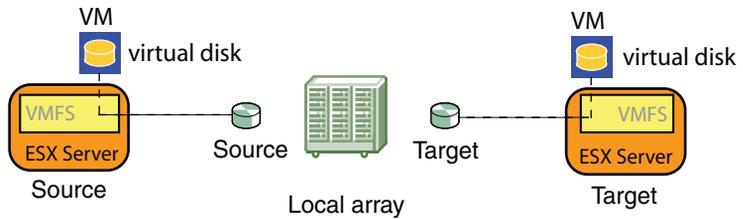


Figure 41 Virtual disk on CLARiiON or VNX local

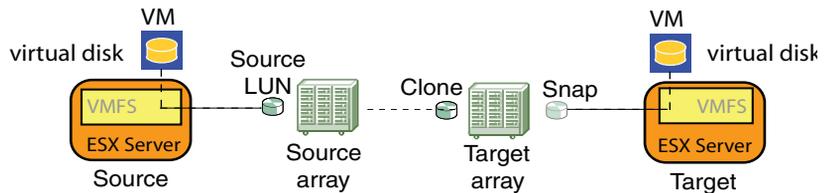


Figure 42 Virtual disk on CLARiiON or VNX using SAN Copy replication

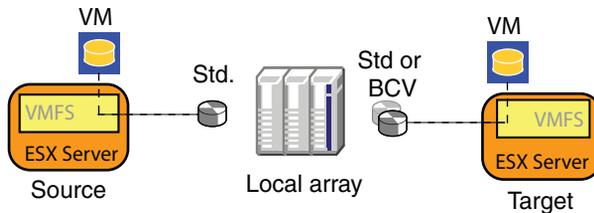


Figure 43 Virtual disk on Symmetrix TimeFinder/Clones

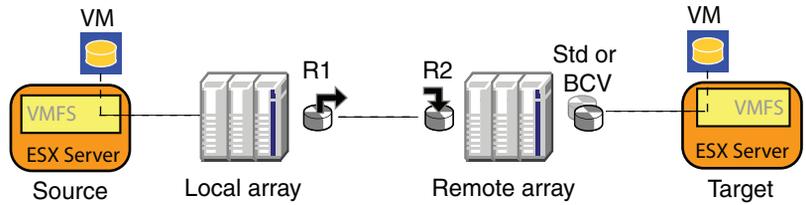


Figure 44 Virtual disk on Symmetrix TimeFinder/Clones (remote)

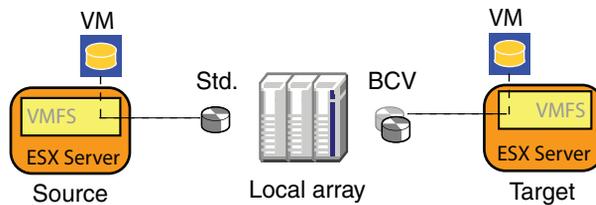


Figure 45 Virtual disk on Symmetrix TimeFinder/Mirror

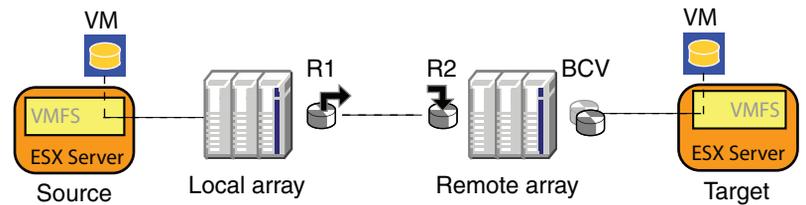


Figure 46 Virtual disk on Symmetrix TimeFinder/Mirror (remote)

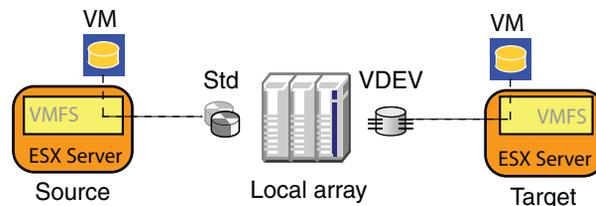


Figure 47 Virtual disk on Symmetrix TimeFinder/Snap (VDEVs)

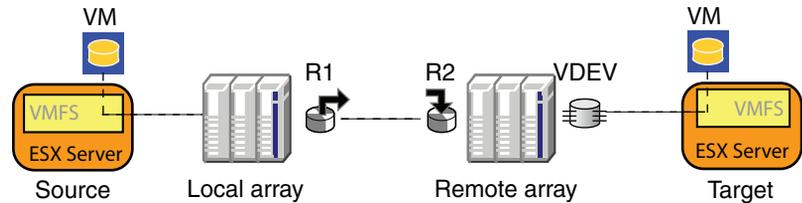


Figure 48 RDM on Symmetrix TimeFinder/Snap (VDEVs) remote

VMware Raw Device Mapping

Data protection of Raw Device Mapping (RDM) (in Windows environments) has the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services:
 - Celerra File Server
 - Celerra File Server to remote Celerra File Server
 - CLARiiON SnapView
 - CLARiiON-to-CLARiiON SAN Copy
 - CLARiiON-to-VNX SAN Copy
 - VNX SnapView
 - VNX-to-VNX SAN Copy
 - VNX-to-CLARiiON SAN Copy
 - Symmetrix-to-CLARiiON SAN Copy
 - Symmetrix TimeFinder/Clone
 - Symmetrix TimeFinder/Clone Remote
 - Symmetrix TimeFinder/Mirror
 - Symmetrix TimeFinder/Mirror Remote
 - Symmetrix TimeFinder/Snap (VDEVs)
 - Symmetrix TimeFinder/Snap (VDEVs) remote
- ◆ **Source** — VMware RDM disks exposed to the ESX Server stored on one of the following source devices:
 - Celerra LUNs (source for a Celerra SnapSure local replica)

- CLARiiON LUNs, SnapView Snapshots, MirrorView/S or MirrorView/A primary devices, MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
- VNX LUNs, SnapView snapshots, MirrorView/S or MirrorView/A primary devices, MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
- Standards (STDs) RAID5 and non-RAID5, TimeFinder/Clone, TimeFinder/Mirror, or R1s from SRDF/A or SRDF/S, or BCVs (RAID5 and non-RAID5). Replication Manager supports replication of thin devices with this technology.
- ◆ **Target** — VMware RDM disks exposed to the ESX Server created on one of the following source devices:
 - Celerra SnapSure local snapshot
 - CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - VNX LUNs, VNX SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - Symmetrix STDs, BCVs or BCV of an R2 (depending upon the selected underlying replication technology)
- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply. In addition to the following VMware requirements, you must use one of the following to connect: VMware iSCSI Initiator, Microsoft iSCSI Initiator, or fibre connections when replicating or mounting Raw Device Mapping (RDM). The *EMC Replication Manager Administrator's Guide* provides information on when to use each and guidelines to follow when installing iSCSI Initiators. See the chapter titled "VMware Setup."
- ◆ **Mount and Recovery** — Static mount of devices is required. No proxy host is needed for RDM replication.

The storage technologies supported with RDM disks are illustrated in [Figure 49 on page 99](#) through [Figure 59 on page 101](#).

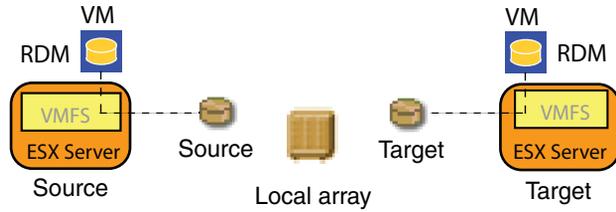


Figure 49 RDM on Celerra (local)

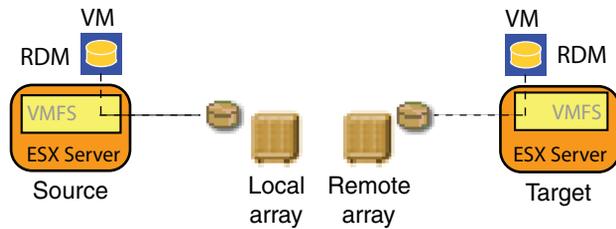


Figure 50 RDM on Celerra (remote)

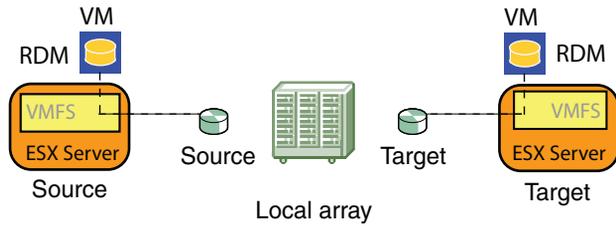


Figure 51 RDM on CLARiiON or VNX local

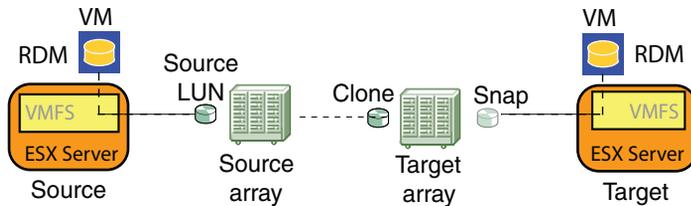


Figure 52 RDM on CLARiiON-to-CLARiiON SAN Copy

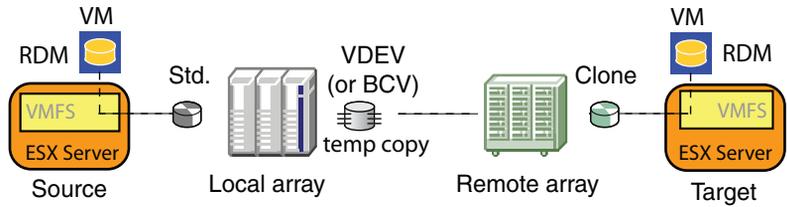


Figure 53 RDM on Symmetrix-to-CLARiiON SAN Copy

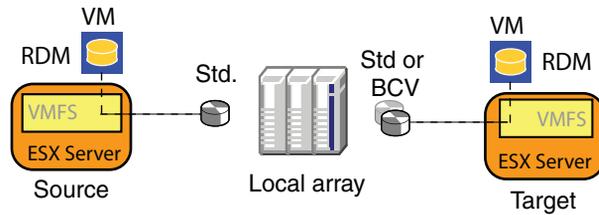


Figure 54 RDM on Symmetrix TimeFinder/Clones

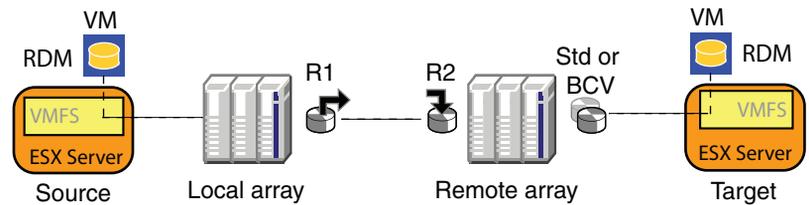


Figure 55 RDM on Symmetrix TimeFinder/Clones (remote)

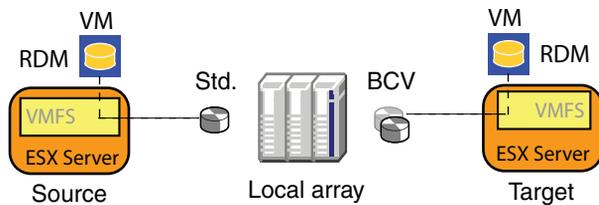


Figure 56 RDM on Symmetrix TimeFinder/Mirror

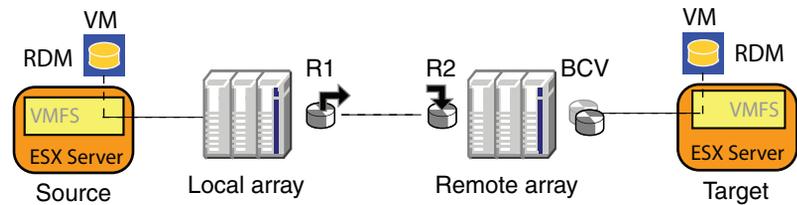


Figure 57 RDM on Symmetrix TimeFinder/Mirror (remote)

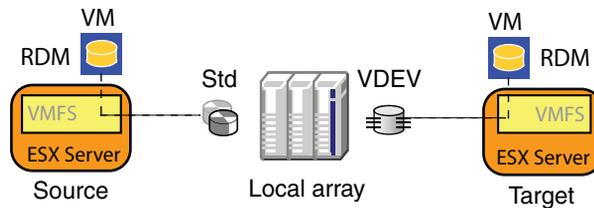


Figure 58 RDM on Symmetrix TimeFinder/Snap (VDEVs)

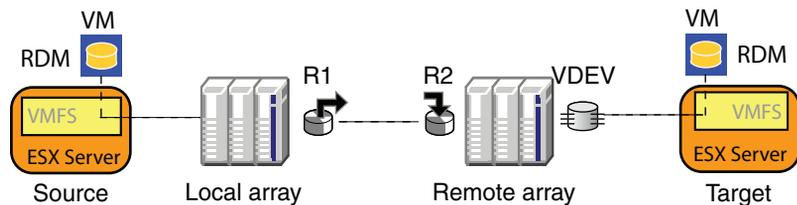


Figure 59 RDM on Symmetrix TimeFinder/Snap (VDEVs) remote

VMware Microsoft iSCSI Initiator discovered LUNs

Data protection of Microsoft iSCSI Initiator discovered LUNs (in Windows environments) has the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services. Refer to [Figure 60 on page 103](#) through [Figure 63 on page 103](#) for an illustration of these configurations:
 - Celerra File Server
 - CLARiiON storage array (including CLARiiON SnapView support, CLARiiON-to-CLARiiON Full SAN Copy support, or CLARiiON-to-CLARiiON Incremental SAN Copy support)

- VNX storage array (including VNX Snapshot support, VNX-to-VNX Full SAN Copy support, or VNX-to-VNX Incremental SAN Copy support)
- ◆ **Source** — iSCSI LUNs exposed directly to the virtual machine:
 - Celerra LUNs (source for a Celerra SnapSure local replica)
 - Celerra SnapSure local snapshot (source for a Replicator remote copy replica)
 - CLARiiON LUNs, SnapView Snapshots, MirrorView/S or MirrorView/A primary devices, MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
 - VNX LUNs, SnapView snapshots, MirrorView/S or MirrorView/A primary devices, MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
- ◆ **Target** — Any one of the following target devices:
 - Celerra SnapSure local snapshot
 - Replicator remote snapshot
 - CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - VNX LUNs, VNX SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply. In addition to the following VMware requirements:

You must use Microsoft iSCSI Initiator. The *EMC Replication Manager Administrator's Guide* provides information on when to use each and guidelines to follow when installing the Microsoft iSCSI Initiator. See the chapter titled "VMware Setup."

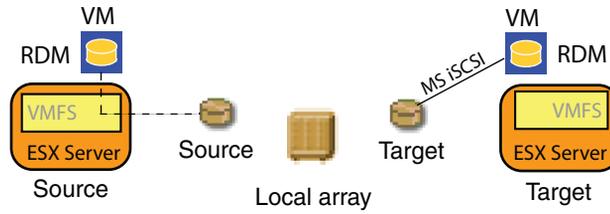


Figure 60 MS iSCSI Initiator discovered LUNs on Celerra (local)

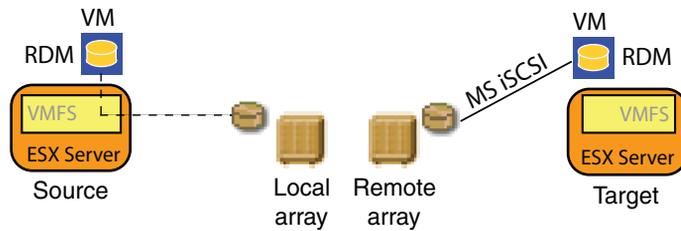


Figure 61 MS iSCSI Initiator discovered LUNs on Celerra (remote)

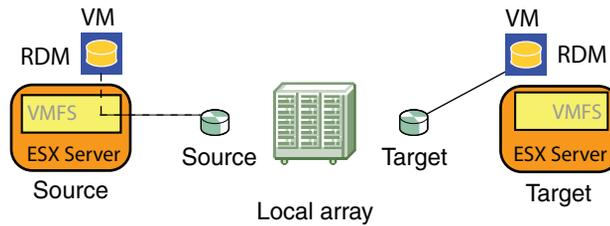


Figure 62 MS iSCSI Initiator discovered LUNs on CLARiiON or VNX local

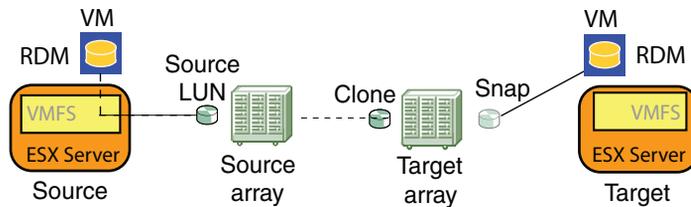


Figure 63 MS iSCSI Initiator discovered LUNs on CLARiiON-to-CLARiiON SAN Copy

More specifics of VMware setup requirements are described in the *EMC Replication Manager Administrator's Guide*. Refer to that document for more information on setup requirements.

Hyper-V support

Data protection of Microsoft iSCSI Initiator-discovered LUNs or pass-through LUNs in a Hyper-V environment has the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services. Refer to [Figure 64 on page 105](#) through [Figure 68 on page 106](#) for an illustration of these configurations:
 - Celerra File Server
 - CLARiiON storage array
 - VNX storage array
 - Symmetrix storage array
- ◆ **Source** — iSCSI LUNs exposed directly to the Hyper-V virtual machine, or pass-through LUNs exposed to the Hyper-V virtual machine:
 - Celerra LUNs (source for a Celerra SnapSure local replica)
 - Celerra SnapSure local snapshot (source for a Replicator remote copy replica)
 - CLARiiON or VNX LUNs, MirrorView/S or MirrorView/A primary devices, MirrorView/S or MirrorView/A secondary devices (however a single application set cannot contain both MirrorView/S and MirrorView/A primary devices)
 - Symmetrix pass-through standard devices or R1 devices
- ◆ **Target** — Any one of the following target devices:
 - Celerra SnapSure local snapshot
 - Replicator remote snapshot
 - CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - VNX LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - Symmetrix TimeFinder Clones, TimeFinder Snaps (depending upon the underlying technologies used)

- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply. Storage devices can be exposed to the guest OS as iSCSI or pass-through LUNs.

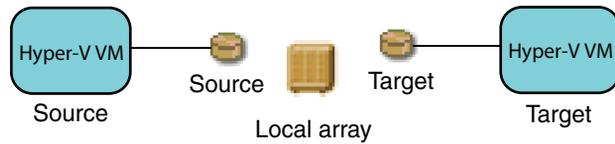


Figure 64 Hyper-V MS iSCSI LUNs on Celerra (local)

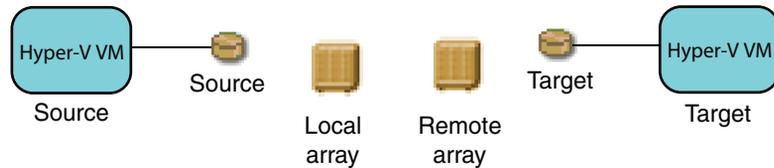


Figure 65 Hyper-V MS iSCSI LUNs on Celerra (remote)



Figure 66 Hyper-V MS iSCSI LUNs on CLARiON or VNX (local)

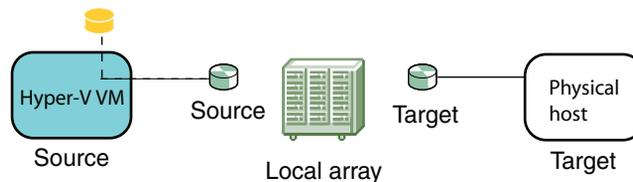


Figure 67 Hyper-V pass-through LUNs on CLARiON or VNX (local)

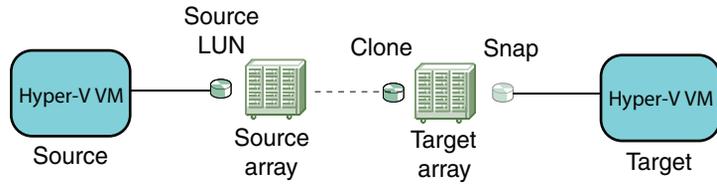


Figure 68 Hyper-V iSCSI LUNs on CLARiiON-to-CLARiiON SAN Copy (remote)

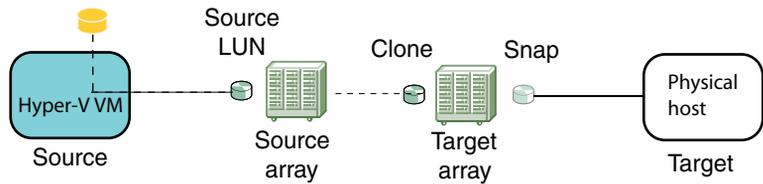


Figure 69 Hyper-V pass-through LUNs on CLARiiON-to-CLARiiON SAN Copy (remote)

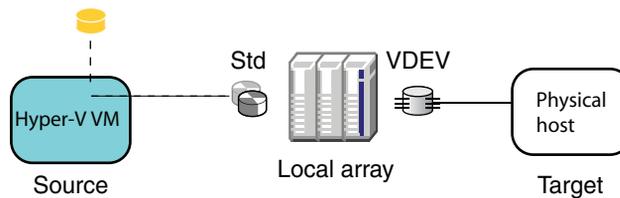


Figure 70 Hyper-V pass-through LUNs on Symmetrix (local)

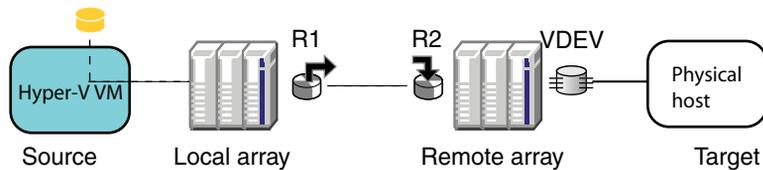


Figure 71 Hyper-V pass-through LUNs on Symmetrix (remote)

The *EMC Replication Manager Administrator's Guide* describes more specifics of Hyper-V setup requirements.

AIX VIO LPAR support

Data protection of Virtual I/O (VIO) Logical Partitions (LPARs) is supported in two possible ways:

- ◆ Physical HBAs connected to each LPAR
- ◆ NPIV configuration for each LPAR

IBM AIX systems can have a physical fibre connection to multiple unique worldwide port names (WWPN). The physical storage is mapped to logical units (LUNs) and the LUNs are mapped to the ports of physical Fibre Channel adapters.

This configuration has the following characteristics:

- ◆ **Storage Services** — Can include any of the following underlying storage services. Refer to [Figure 72 on page 108](#) through [Figure 77 on page 110](#) for illustrations of these configurations:
 - CLARiiON local clones
 - CLARiiON SnapView snaps
 - CLARiiON-to-CLARiiON SAN Copy
 - CLARiiON-to-VNX SAN Copy
 - VNX local clones
 - VNX SnapView snaps
 - VNX-to-VNX SAN Copy
 - VNX-to-CLARiiON SAN Copy
 - Symmetrix TimeFinder/Clone
 - Symmetrix TimeFinder/Mirror
 - Symmetrix TimeFinder/Mirror Remote
 - Symmetrix TimeFinder/Snap (VDEVs)
 - Symmetrix TimeFinder/Snap (VDEVs) remote
 - BCVs of an R2 (depending upon the selected underlying replication technology)
- ◆ **Source** — IBM AIX VIO LPARs with a physical fibre connection to multiple unique worldwide port names (WWPN). The physical storage is mapped to logical units (LUNs) and the LUNs are mapped to the ports of physical Fibre Channel adapters or to NPIV WWPNs. Acceptable storage includes:

- Symmetrix Standards (STDs), TimeFinder/Mirror BCV, TimeFinder/Clone BCV, R1 from SRDF/S
- CLARiiON LUNs, SnapView Snapshots
- VNX LUNs, SnapView snapshots
- ◆ **Target** — Any one of the following target devices:
 - Symmetrix STDs, BCVs, snapshots (VDEVs), remote snapshots, BCV of an R2, or SAN Copy targets (depending upon the selected underlying replication technology)
 - CLARiiON LUNs, SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
 - VNX LUNs, VNX SnapView Snaps, SnapView Clones, or remote LUNs (depending upon the underlying technologies used)
- ◆ **Storage Requirements** — All storage requirements imposed by the underlying technologies apply. In addition, the following requirements exist:
 - The virtual machine must be zoned and masked to at least four gatekeeper LUNs on the Symmetrix.
 - Gatekeepers must be exported in physical compatibility mode.
 - To support successful mounts, the target storage must be zoned to the VIO client or physical server used as the mount host.
 - To support mounts the target host must be an LPAR or physical AIX host.
 - In Symmetrix environments, target storage must be pre-exposed to the mount host.

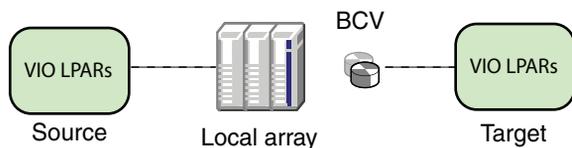


Figure 72 VIO LPARs on Symmetrix TimeFinder/Mirror (mirror remote and clone also available)

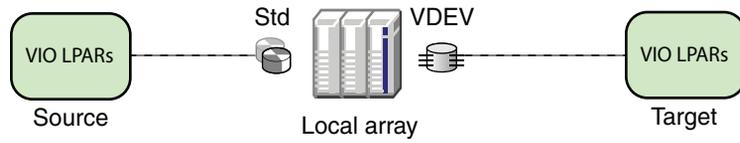


Figure 73 VIO LPARs on Symmetrix TimeFinder/Snap (VDEVs)

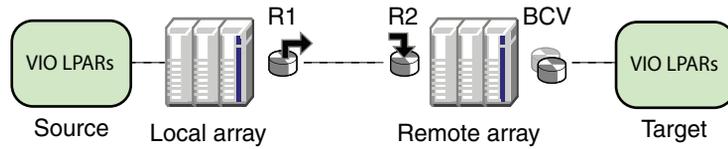


Figure 74 VIO LPARs on BCV of SRDF R2

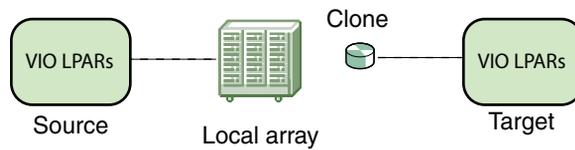


Figure 75 VIO LPARs on CLARiiON or VNX SnapView clones

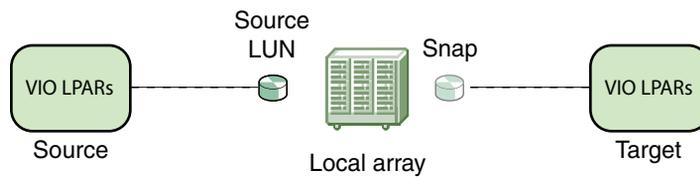


Figure 76 VIO LPARs on CLARiiON SnapView snaps

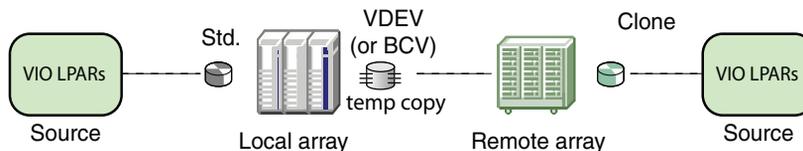


Figure 77 VIO LPARs SAN Copy

Number of replicas supported

Table 5 on page 110 lists the number of replicas per source device supported by each replication technology in Replication Manager.

Table 5 Number of replicas supported (1 of 2)

Replication technology	Number of replicas per source
TimeFinder/Snap	Up to 128 snaps per source with microcode releases that support 128 snaps (15 snaps with older microcode). Use of multiple replication technologies may reduce the number of replicas supported. Also, refer to the Symmetrix Array Setup chapter of the EMC Replication Manager Administrator's Guide for more information on enabling support for 128 TimeFinder/Snap sessions. Note that TimeFinder/Snap copy jobs must use the same snap pool as the source job.
TimeFinder/Duplicate Snap	Up to 128 snaps per source with microcode releases that support 128 snaps (15 snaps with older microcode). Use of multiple replication technologies may reduce the number of replicas supported. Also, refer to the Symmetrix Array Setup chapter of the EMC Replication Manager Administrator's Guide for more information on enabling support for 128 TimeFinder/Snap sessions.
TimeFinder/Clone	Up to 8 TimeFinder clones per source. Use of multiple replication technologies may reduce the number of replicas supported.
TimeFinder/Mirror	Up to 8 BCVs per source. Use of multiple replication technologies may reduce the number of replicas supported.
EMC Open Replicator	Up to 15 replicas per Symmetrix control device. Use of multiple replication technologies may reduce the number of replicas supported.

Table 5 Number of replicas supported (2 of 2)

Replication technology	Number of replicas per source
SnapView Snap	Maximum snaps per source depends upon several factors. Refer to Table 18 on page 251 for specifics. Maximums are inclusive of (SnapView snap + Incremental SAN Copy) per source.
SnapView Clone	Maximum clones per source depends upon several factors. Refer to Table 18 on page 251 for specifics.
Full SAN Copy	Limited only by number of available target devices.
Incremental SAN Copy	Up to 8 per source (including the SnapView snap plus the incremental SAN Copy).
Celerra SnapSure	2000 (1000 tested) snaps per LUN.
Celerra or VNX Replicator	iSCSI: 1024 configured, 128 concurrent replication sessions.
Celerra or VNX NFS	Depends upon available space in SavVol.
RecoverPoint	336 replicas per application set.

Managing application sets

Replication Manager uses application sets to define the following parameters:

- ◆ What physical or virtual hosts contain the information you want to replicate.
- ◆ What information will be replicated. You can choose data from one or more applications.
- ◆ Which users can have access to the application set, its jobs, and its replicas.

[Figure 78 on page 112](#) shows a graphical representation of the steps necessary to create an application set.

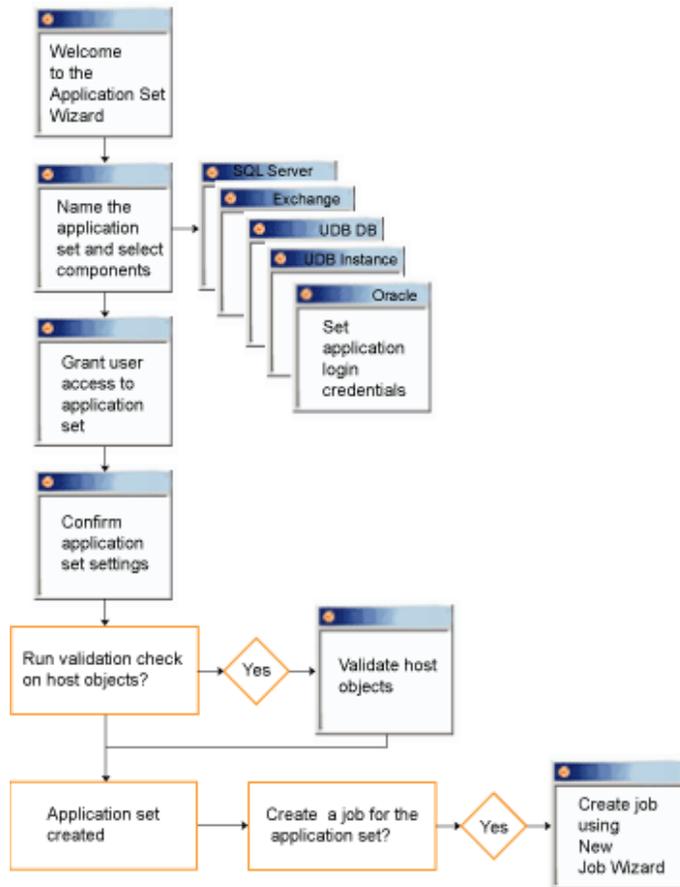


Figure 78 Application set creation process

Federated application set support

Replication Manager includes support for federated application sets. In other words, Replication Manager can support application sets that replicate data residing in multiple storage arrays of the same type, hosted on multiple agent hosts, with data from multiple applications. Replication Manager maintains consistency across multiple hosts, applications, and storage devices by leveraging EMC Engenuity™ Consistency Assist (consistent split) technologies.

For complete information about how to create a federated application set, refer to [Chapter 7, "Configuring Federated Data."](#)

Composite application set support

If you are creating an application set that contains multiple, possibly heterogeneous sets of application data residing on a single production host (also known as a composite application set), remember the following important considerations:

- ◆ If you include Exchange data in an application set with any other database object (except file systems) the replication will likely fail due to a VSS timeout. EMC does not recommend that combination.
- ◆ Multiple applications can mean data from two different applications (for example, Oracle and SQL Server) or data from two different instances of the same application (for example, two different instances of SQL Server).

Note: A single application set cannot support heterogeneous logical volume managers in the same application set. Refrain from creating application sets that span different logical volume managers.

RecoverPoint consistency groups

Devices found in RecoverPoint consistency groups must match the devices specified in the corresponding application sets and job options. If the devices do not match, an error is generated when the job is run.

After modifying a RecoverPoint consistency group, you should modify the corresponding application set and any job options that affect what is replicated (for example, Oracle flash recovery area and archive log directories), then run a job.

Avoid changing the name of a RecoverPoint consistency group while Replication Manager is managing the protection. Mount and restore failures can occur after a consistency group is renamed.

MirrorView/A and MirrorView/S consistency groups

An application set must consist of LUNs that are all part of a given consistency group or contain only individual LUNs that are not part of any consistency group. Additionally, the consistency group cannot contain any additional LUNs that are not part of the replication. You cannot mix some LUNs from a consistency group and some that are

individual in the same application set. An application set cannot contain LUNs from more than one consistency group.

If a MirrorView / A application set contains multiple LUNs, then EMC recommends that all LUNs that are part of the replica be placed together into a single consistency group with no additional LUNs in that group.

Application sets and VMware

When you are defining an application set to create a crash-consistent VMFS replica or a virtual disk replica in a VMware environment, there are additional tasks to perform before and during the application set configuration. This section outlines those additional tasks:

1. Before creating the application set, register the appropriate hosts:

- **When creating crash-consistent VMFS replicas on an ESX Server** — Register a physical or virtual proxy host that can communicate with VMware VirtualCenter software. VirtualCenter is the component of VMware software that manages the VMware-specific processing during a replication, mount, or restore operation. The proxy host allows Replication Manager to list every VMFS that is configured on all the ESX Servers managed by the VirtualCenter, regardless of whether the VMFS is on supported storage or not.

Note: ERM Administrators must know which VMFS resides on supported storage and choose only those for replication.

- **When replicating virtual disks on a virtual machine** — Register the virtual machine production host as well as any virtual machine you intend to use as a mount host. Again, these virtual machines must have visibility to VMware VirtualCenter software.

Note: In CLARiiON environments, ensure that the hostname you use to register an ESX Server matches the hostname in the storage array management software, CLARiiON Navisphere. Also, be sure to use fully-qualified host names for ESX Servers and not IP addresses.

[Figure 79 on page 115](#) illustrates the register new host screen where you can set the VirtualCenter credentials.

The screenshot shows a 'Register New Host' dialog box with the following fields and options:

- Host Name: RMProxy-14
- Port: 6542
- Options tab selected, containing:
 - Enable RecoverPoint
 - RPA management hostname: [empty field]
 - Enable VMware
 - Virtual center host name: MyVirtualCenter
 - Username: Administrator
 - Password: [masked with asterisks]
 - Require users to supply these credentials for VMFS replications
- Buttons: OK, Cancel, Help

Figure 79 Registering a VMFS Proxy host to enable VMware to create crash-consistent VMFS replicas

2. During registration of VMware hosts (either Replication Manager proxy hosts or virtual machines that are to be used as mount hosts), enter the VirtualCenter credentials to grant access to the VirtualCenter software. Replication Manager uses VirtualCenter to create, mount, or restore replicas in VMware environments.
3. Enter the hostname where VirtualCenter software is installed in the **Virtual center host name** field. If you are using a nonstandard port to connect to that host, add the port after a colon. For example, vhost:1234.

4. Enter the username. If the user is a domain user, specify the username with the following format:
domain/username
5. Enter the password for that account.
6. Decide whether to **Require users to supply the credentials for VMFS replications** when they create an application set.
7. If you choose to require VirtualCenter credentials to configure an application set, Replication Manager requests a username and password when you attempt to create a VMFS application set. These credentials are the VirtualCenter credentials as shown in [Figure 80 on page 116](#).



Figure 80 VirtualCenter credentials dialog (VMFS application sets only)

Note: For information about mounting and restoring VMware replicas, read [“VMware mount and restore” on page 200](#). For more information on VMware support and setup for VMware environments, read the chapter on [“VMware Setup”](#) in the *EMC Replication Manager Administrator’s Guide*.

Defining a new application set

To define a new application set:

1. Right-click **Application Sets** and select **New Application Set**. The **Application Set Wizard** appears.
2. Use the wizard to perform each of the steps outlined in [Figure 78 on page 112](#).

Note: For specific Application Set Wizard panel information, click **Help**.

To view information about existing application sets, click **Application Sets** in the tree panel. The content panel displays information about each application set.



CAUTION

EMC strongly recommends that you avoid mapping two different application sets to the same source hypervolume because of the potential for contention between the two application sets. If simultaneous jobs on two separate application sets try to access the same source hypervolume, corruption and loss of data could occur. You can run multiple jobs from the same application set.

Validating an application set

After you have defined a new application set, Replication Manager can validate the configuration of the application set. Replication Manager reveals potential problems in the agent configuration such as, but not limited to:

- ◆ Invalid or missing database configuration files
- ◆ Invalid database settings
- ◆ Invalid or nonexecutable support applications
- ◆ Incorrect agent configuration (file system, Oracle, IBM UDB, Microsoft SQL Server, Microsoft Exchange, or SharePoint)

As noted earlier, Replication Manager specifically validates configuration information related to the following agents: file system, Oracle, IBM UDB, Microsoft SQL Server, and Microsoft Exchange, and SharePoint. In all cases, the permissions and executables associated with these agents are validated. Depending on the type of agent, Replication Manager may also check other details of the configuration, specific to that agent type.

Many of these validation processes occur as soon as the information is entered. Others occur when you choose **Validate** on the last Application Set Wizard panel.

Replication Manager validation cannot detect hybrid storage (one application set with more than one type of storage array). Replication Manager application set validation will not show an error in the case of this type of application set, however, it is not legal to set up an application set that contains “hybrid storage.” For example, you

cannot configure an application set containing both CLARiiON and Symmetrix objects, except in the case of an Exchange 2007 Cluster Continuous Replication (CCR) environment with separate jobs accessing separate storage types for each physical node of the cluster.

On the application set validation panel you can choose whether to display properties and you can also choose whether to display the messages in Text View or Table View.

Replication Manager defaults to the Table View shown. Click the **Text View** button to change to the Text View.

Modifying an application set

To modify an application set:

1. Expand **Application Sets**.
2. Right-click a particular application set from the tree and select **Properties**. The **Application Set Properties** window appears.
3. Modify the information appropriately.
4. Click **OK** to complete the operation.
5. If you added or removed a database, remember to update the job information for both original jobs and copy jobs associated with the application set as follows:
 - Modify advanced replication options as appropriate.
 - Modify mount options (enable mount to do so).
 - For SQL Server jobs, change from by instance to by database. and verify that all of the databases are visible, then save the changes.
 - If mount was originally disabled, disable it again and save the changes again.

Note: For specific information about each wizard panel, click **Help**.

Deleting an application set

When you delete an application set, Replication Manager also deletes all related jobs and replicas as well. A warning message appears to warn you that this will occur.

To delete an application set:

1. Right-click the application set from the tree and select **Delete**.
2. Click **Yes** when asked if you are sure that you want to delete the application set and its related components.

Removing application objects from an application set

To delete just one application object from an application set with multiple application objects:

1. Modify the application set as described in [“Modifying an application set” on page 118](#).
2. Click the **Objects** tab.
3. Select the objects you want to remove and click **Remove**. Remember to modify the job if necessary to remove unneeded information about applications that are no longer part of the application set.
4. Clear the checkbox next to an object to remove it.
5. Click **OK**.

Federated application set

If the removal of an object will cause a federated application set (one containing data on more than one production host) to become nonfederated (containing data on only one host) you will not be allowed to complete the operation. If that is the case, delete the application set and rebuild it.

SharePoint application set

Individual objects cannot be added or removed from a SharePoint application set. You can update a SharePoint application set for a configuration that has changed since the last save by clicking the **Update** button in the application set properties' **Objects** tab.

Celerra jobs

If you remove an object from an application set that has existing remote Replicator replicas, Replication Manager marks those existing replicas as not restorable the next time a Celerra job runs to create a replica with the new set of objects.

About user access to application sets

The user who defines an application set owns that application set. The application set owner can grant or deny other users permission to access an application set. Once a user has been granted access, the level of control depends on the user role. See “[User roles](#)” on page 36 for a list of what actions a user can perform based on role.

Copy and link jobs

Access to certain application sets can impact a user’s ability to modify or run jobs that are linked across application sets. Ensure that a user who is running jobs that link across application sets has adequate access to all related application sets. For more information, refer to “[Understanding link and copy jobs](#)” on page 137.

Storage

The storage available to an application set is limited to the storage that is available to the user who created the application set. A user who has access to an application set, but not to its storage, has limited ability to perform actions related to the application set’s storage.

Users with the ERM Administrator role have access to all included storage; therefore, any application set owned by an ERM Administrator can use any included storage.

Granting or revoking access to an application set

To modify the list of users who have access to the application set:

1. Modify the application set as described earlier and click the **Access** tab.
2. Click **Grant**, select the users who you want to add to the list and click **OK**.
3. Select the users to remove from the list and click **Deny**.

Special application set considerations

When you are creating an application set, you must consider your environment and the organization of the data on the storage devices. Application data that shares devices with other unrelated information can cause undesirable consequences and may even result in lost data. The following section describes the issues that you must consider when configuring an application set.

Issues when multiple applications share volume groups or hypervolumes

When data on your production system unintentionally becomes part of a replica because of its proximity to the data you intend to protect, that is an issue that can cause data loss.



CAUTION

Shared storage may occur without your knowledge and if you inadvertently restore an application that shares the same storage with another application or file system, you may inadvertently lose data. In most cases, Replication Manager warns you that you are about to start a restore that might cause this issue.

You can prevent situations like this by planning your data layout based on replica granularity. The granularity of a replica depends on the environment. If the system is using a volume manager, the replica is taken at volume group granularity.



CAUTION

If you are restoring a Veritas volume group that has the same name as a logical volume that resides in a different volume group, the restore fails. Ensure that the volume group and logical volume names are unique in your system.

If the system is not using a volume manager, the replica granularity is at the hypervolume level on Symmetrix arrays, the LUN level on CLARiON, VNX, VNXe, and Celerra.

File system shared storage issues

Replication Manager always creates replicas of one or more entire hypervolume (on Symmetrix), LUN (on other CLARiON, VNX, VNXe, and Celerra iSCSI). A hypervolume in the storage array is a logical representation of a physical volume. From the perspective of the host system, the hypervolume is the same as a physical volume.

If you are using a volume manager, Replication Manager creates a copy of all the hypervolumes, LUNs, or virtual volumes that make up any part of the volume group.

For example, consider a hypervolume that contains:

- ◆ A Solaris partition (/data) mounted on /dev/dsk/c1t1d1s6.
- ◆ A raw partition containing an Oracle tablespace on /dev/rdisk/c1t1d1s5.

If you create a replica of the file system `/data`, Replication Manager creates a replica of the entire hypervolume, including the file system and the Oracle tablespace because they both reside on the same hypervolume. If you decide to restore the replica, the older version of the Oracle tablespace will also be restored. The Oracle tablespace gets restored because it shares storage with the file system, even though you did not intend to restore it. The Oracle tablespace was restored as a by-product of the hypervolume layout.

Therefore, you should plan your hypervolume layout to include only objects that should be replicated and restored as a single unit. Don't mix unrelated data objects in the same hypervolume.

Veritas volume group backup required maintenance

The `vxconfigbackupd` daemon detects changes in the Veritas environment, and makes backup copies of volume groups in the following directory:

```
/etc/vx/cbr/bk
```

These changes often occur as a result of using Replication Manager in a Veritas environment. Entries in this directory may accumulate over time. The storage administrator should monitor this directory to ensure that it does not fill up the root file system.

The `vxconfigbackupd` daemon is started automatically in Veritas environments.

Volume group affected entities

Plan your volume group layout to include only objects that you want to replicate and restore as a single unit. If you want to restore at the tablespace level, create one tablespace per volume group. If you want to restore at the database level, create one database per volume group.

When you are using a volume group manager, such as Veritas Volume Manager (VxVM), replicas contain the entire volume group.

For example, consider a Veritas Volume Manager installed on a Solaris system. Also consider that there is a volume group that includes three different hypervolumes, like the example shown in

Figure 81 on page 123.

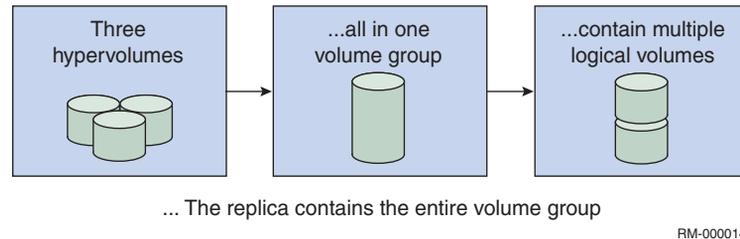


Figure 81 Veritas volume group example

In the example shown in [Figure 81 on page 123](#), Replication Manager replicates the entire volume group, even if you choose to replicate only a subset of what resides in the volume group. When you restore the replica, Replication Manager deports the volume group, restores the hypervolumes, and re-imports the volume group.

Veritas Volume Manager: Mixing CDS and non-CDS volume groups

In an environment using Veritas Volume Manager, Replication Manager can import non CDS volume groups (that is, with sliced disks) onto a host where CDS (Cross-Platform Data Sharing) is the default volume group mode. Replication Manager does not support mounting of volume groups from one host platform to another, for example, mounting Solaris replicas on non-Solaris hosts.

Application sets on arrays that support concurrent BCVs

Some Symmetrix storage arrays support Concurrent BCV technology, which allows two BCVs of the same standards to be established simultaneously. Replication Manager uses Concurrent BCV technology to create replicas of data that is already protected by other replication technology.

To ensure that Replication Manager does not infringe on the operation of other products, established BCV connections are not split unless they were originally established as part of a replica creation. If there are no available mirror positions, and Replication Manager cannot split any of the existing BCVs from the standard, the replication fails.

Managing jobs

Users with any role except Operator can specify how to replicate an application set by defining one or more jobs for that application set.

When you define a job, you determine:

- ◆ How long to save each replica (retention period) or the maximum number of replicas in a replica rotation (the maximum number of replicas kept for the job, not for the application set).
- ◆ Whether to use consistent-split technology to quiesce the data.

Note: If you are using consistent-split technology, refer to [Chapter 6, "Using Consistent Split,"](#) for more information and considerations associated with mounting applications using that technology. Remember that for federated application sets, you must use Consistent Split.

- ◆ Whether to start optional user-defined pre- and post-replication scripts before or after each part of the job.
- ◆ What user credentials to use when running a script (if the Replication Manager Agent (IRCCD) was started in secure mode).
- ◆ How to select storage for replicas using storage pools. You can choose from a wide variety of storage options, including TimeFinder/Mirrors, TimeFinder/Clones, TimeFinder/Snaps, or Full SAN Copy (on the Symmetrix storage array); SnapView clones, SnapView snaps, Full SAN Copy, or Incremental SAN Copy (on the CLARiON or VNX storage arrays). In addition, there are options that include virtual environments such as VMware or Hyper-V.
- ◆ Whether and where to mount or back up the system after replication. You can choose to mount to alternate hosts (or mount to multiple alternate hosts if this is a federated application set) or to the production host.

Figure 82 on page 125 illustrates the job creation process.

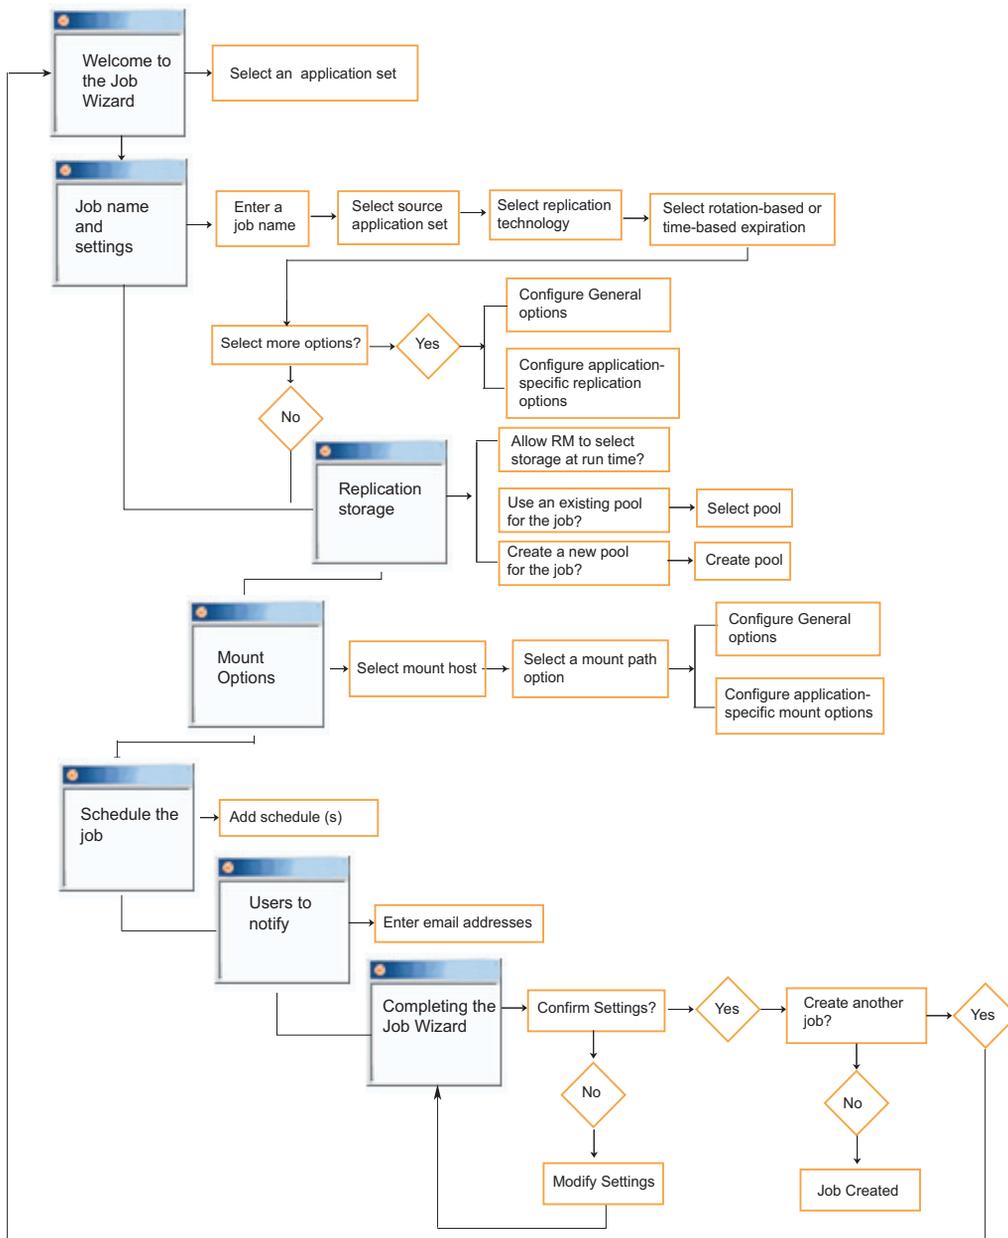


Figure 82 Job creation process

Defining a new job

To define a new job:

1. Right-click **Jobs** in the tree panel.
2. Select **New Job** from the context menu. The Job Wizard steps you through the process to create a job.

Note: For specific information about each panel, click the **Help** button on the bottom of the wizard panel.

After you create a job, you can right-click the job on the tree panel and choose **Properties** from the context menu to view job properties.



CAUTION

When you define or modify a job that mounts a replica to an alternate host, you must choose the host to which the replica will be mounted. Ensure that you choose a host to which the owner of the application set has been granted access. If the owner of the application set does not have access to the selected mount host, the job will fail.

Job name and settings panel

The Job Name and Settings panel of the Job Wizard helps Replication Manager select the appropriate storage for a replica and also helps you decide when and how the replica will be automatically expired and deleted.

The panel differs depending on the type of storage on which your source resides.

To complete this panel for all application sets *except* those that include Celerra network file systems:

1. Enter a name for the job. The job name cannot contain any of the following characters / : * ? " < > , ;
2. Select a source from the **Replication Source** field. The source can be an application set or a job. If the source is set to a job, then your job becomes a copy job that will use a replica as its source.

For more information on copy jobs, refer to [“Understanding link and copy jobs” on page 137](#).

3. Choose a replication technology depending upon:
 - The type of storage on which your source resides
 - The types of storage available in your environment
 - Whether you are creating a copy job

For more information on the available replication technologies, refer to [“Available replication technologies” on page 58](#).

4. Choose how to limit the replica count, you can:
 - Create a rotation
or
 - Set a replica retention period

Rotation and retention options are not available when RecoverPoint is the selected replication technology.

To complete this panel for application sets that include Celerra network file systems:

1. Enter a name for the job. The job name cannot contain any of the following characters / : * ? “ < > , ;
2. Choose an included Celerra or VNX Network Server from the **Snap Destination (Celerra)** field. The selection you make here, in conjunction with the current location of the source allows Replication Manager to determine the replication technology to use in order to complete the job as shown in [Table 6 on page 127](#).

Table 6 Celerra replication technologies used in various configurations

Configuration	Replication technology
Source and destination on the same Celerra Network Server	Celerra SnapSure
Source and destination on different Celerra Network Servers	Celerra or VNX Replicator

This method of specifying the snap destination is failover tolerant because if the Celerra or VNX environment fails over, causing the necessary replication technology to change, Replication Manager automatically adjusts to use the correct replication technology in the current state.

It is possible to configure your jobs to run only before or after a failover has occurred if desired. In the case of Celerra Network File System job, there is an additional checkbox under the Advanced link within General Replication Options titled **Only**

run job if in-frame replication. See [Figure 84 on page 130](#) for an illustration of this panel. Selecting this checkbox configures the job to run only when source and destination are on the same Celerra.

Note: When you select Only run job if in-frame replication Replication Manager runs the job regardless of whether it is local or remote. The job fails if it is not an in-frame replication, however Replication Manager then cleans up and deletes the job, so the effect is the same as if the job had never run.

For more information on the available replication technologies, refer to [“Available replication technologies” on page 58](#).

3. Select a **Snap type**, either read-only or read/write to specify the desired state of the destination replica.

Read/write snapshots can be mounted read/write and changes will persist upon unmount and or restore. Read-only snapshots can also be mounted read/write, however any changes made to a read-only replica mounted read/write will be lost on unmount.

4. Choose how to limit the replica count, you can:
 - Create a rotation (by choosing a replica count)
 - or*
 - Set a replica retention period

5. Click **Advanced** to set the advanced replication settings shown in Figure 83 on page 129.

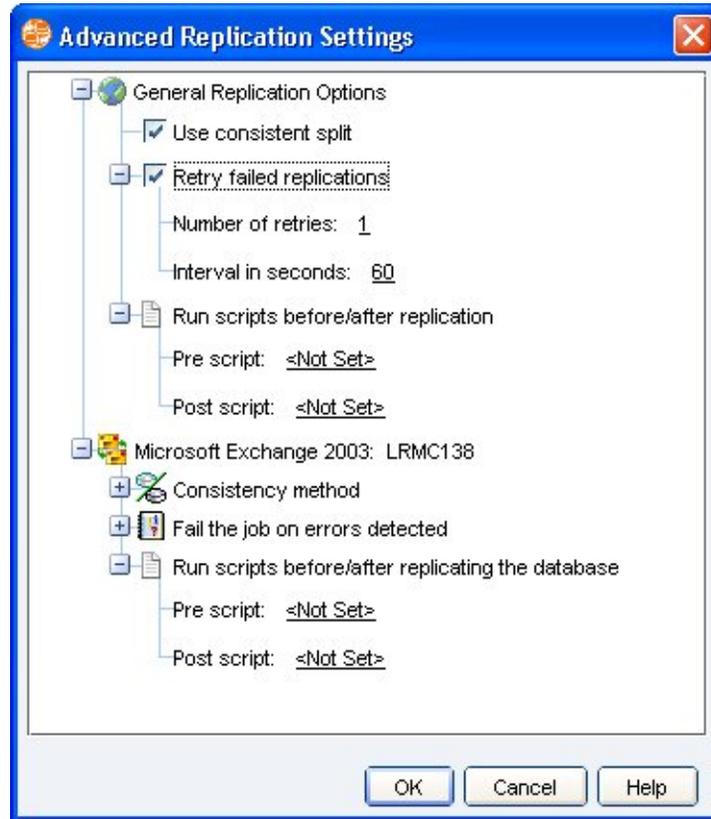


Figure 83 Advanced replication settings (Exchange)

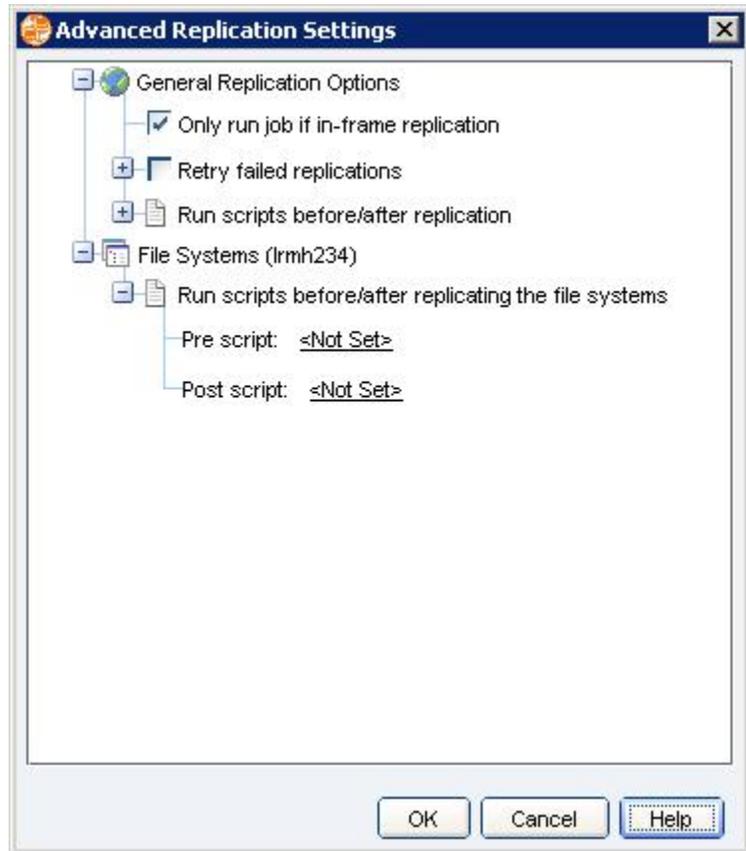


Figure 84 Advanced replication settings (Celerra or VNX NFS)

Note: For more information about each section of the advanced settings, click the **Help** button on the bottom of the panel.

6. You can choose to click **Topology** at any time while you are creating the job. Refer to “[Topology view](#)” on page 136 for more information about the topology view.

Note: The Topology view is not available in NFS environments.

7. Click **OK** to exit the advanced panel and **Next** to access the Replication Storage panel of the Job Wizard.

Replication Storage panel

The Replication Storage panel, shown in [Figure 85](#) on page 131, allows you to define how Replication Manager should choose replication storage for the replicas created by this job.

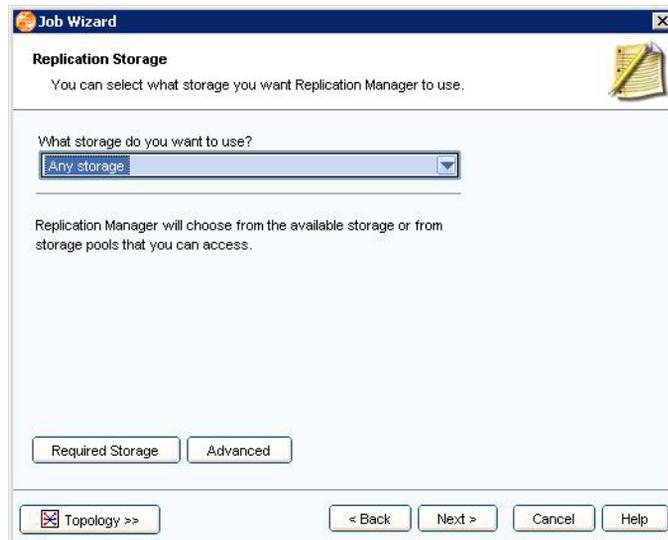


Figure 85 Replication Storage panel

Using storage pools

You can allow Replication Manager to select storage for you, use an existing storage pool, or create a new storage pool on the Replication Storage panel of the wizard. When you use a storage pool, Replication Manager selects devices from that pool to accommodate the storage needs of the replica:

1. Choose the storage you want to use (any storage, existing pool, or new pool).
2. If you need information about what storage type, device size, and the number of devices that you need, click **Required Storage**.
3. Click **Advanced** to view the advanced storage options and select a Save Pool.
4. Click **Next** to move to the next screen of the Job Wizard.

Note: Replication Manager does not support storage pools with Celerra iSCSI, VNXe iSCSI, Celerra NFS or VNX NFS or RecoverPoint. Storage pools are required for some SAN Copy operations. Use Required Storage to review storage needs prior to setting up a pool.

Replication Manager uses pools differently based on the storage technology that you chose in that earlier panel. Refer to [“Available replication technologies” on page 58](#) for more information about each replication technology.

Mount Options panel

The Mount Options panel [Figure 86 on page 132](#) allows you to set the mount options.

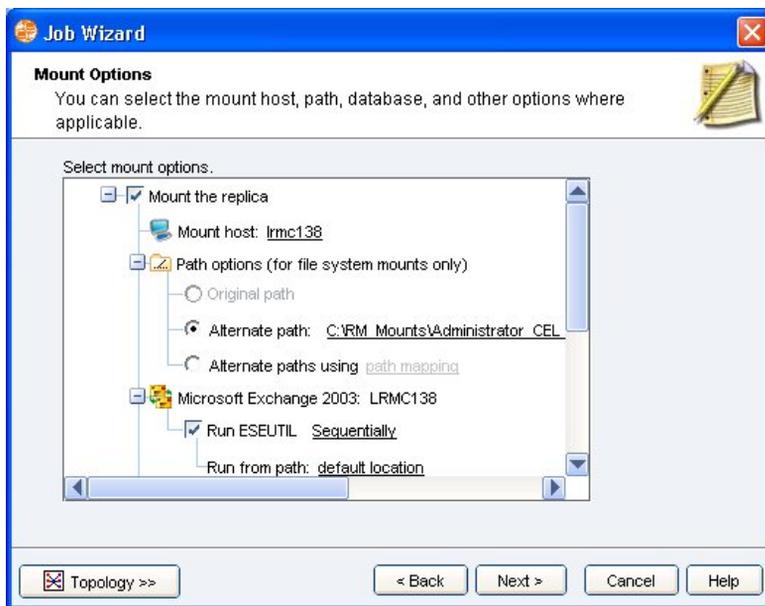


Figure 86 Mount Options panel

Note: EMC recommends that you choose to mount all Windows-based replicas in the job that creates those replicas.

For more information about setting mount options, refer to [“Mounting replicas” on page 162](#). Once you have set the appropriate mount options, click **Next** to move to the next panel of the wizard.

Starting the Job panel

The Starting the Job panel, shown in [Figure 87 on page 133](#), allows you to define how a job gets invoked. Your options include:

- ◆ Starting the job manually or with a third-party scheduler.
- ◆ Creating a schedule on which this job will run.
- ◆ Starting the job after another job completes. The job selected can be associated with the same application set or with another application set.



Figure 87 Starting the Job panel

To start the job panel:

1. Select the radio button that corresponds to your selected action. If you choose **Schedule the job**, the Schedule fields appear as shown above.
2. Click **Add** from the panel and configure a schedule. Refer to “Managing job schedules” on page 152 for more information.

Notification panel

The Notification panel, [Figure 88 on page 134](#), allows you to set up automatic email notification of the status of a job when it runs. Use this panel to set up email notifications and decide whether to notify individuals whenever the job runs, or only when the job fails.

Figure 88 Notification panel

To configure notifications:

1. Enter the email addresses of the people who should be notified of job status.
2. Select the **Send email notification on failures only** checkbox to restrict email notifications to failure conditions.
3. Click **Next** to move to the next panel of the wizard.

Completing the Job Wizard panel

Replication Manager displays a summary screen with all the information about the job in a tree.



Figure 89 Completing the Job Wizard panel

Perform the following tasks:

1. Review the information in the tree shown in [Figure 89 on page 135](#) to ensure that it is correct. If information is incorrect, use the **Back** button to make changes.
2. If the information is correct, do one of the following:
 - Click **Finish** to complete the job and save the information.
 - Click **Another Job >>** to start the Job Wizard at the beginning and create a new copy job.

Topology view

The Topology view displays the relationship between jobs. To access this view, click the **Topology >>** button in the Job Wizard.

Note: The Topology view is not available in Celerra r VNX NFS environments.

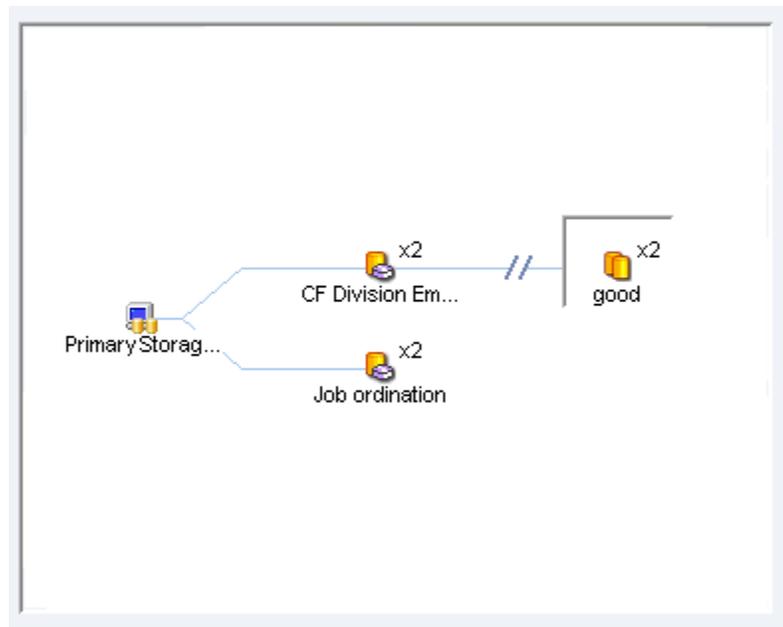


Figure 90 Topology view

The topology view has certain key features illustrated in [Figure 90 on page 136](#):

- ◆ The left most node is always the source (primary) storage or job. A node consists of a replica icon, the name of the job that creates that replica, and the number of replicas in the rotation (x1, x2). When the expiration method is time-based, nothing is displayed.
- ◆ The replication job being configured is shown as selected (with a white background). Refer to the replica created by the job titled *good* in [Figure 90 on page 136](#).
- ◆ A line between two nodes indicates that the left most node is the data source for the right most node.

- ◆ A double slanted line (//) indicates a remote replication.
- ◆ Source storage or a replica created by a job can be the source for multiple other replicas.

Understanding link and copy jobs

Replication Manager allows you to create specialized jobs, known as link jobs and copy jobs:

- ◆ **Link job** — Allow a specified job to start as soon as another job has finished. The wizard gives you the ability to choose the application set and then the job associated with that application set. This allows you to choose any job to link to the current job.

Note: If the job you are trying to link to another source job is already a linked copy job you cannot link it to another job.

Figure 91 on page 137 illustrates a link job.

If the first job fails, the second job of a linked set still runs.

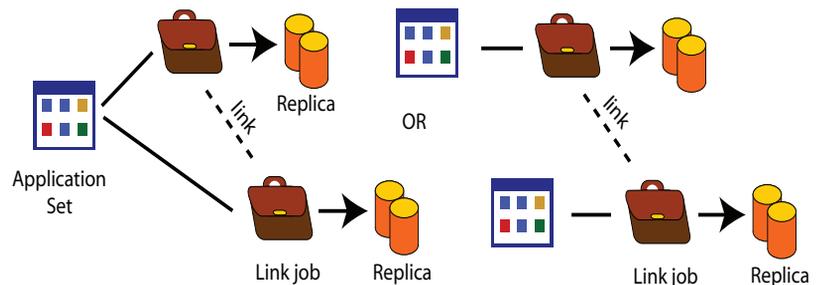


Figure 91 Link job

- ◆ **Copy job** — Allow a specified job to use a replica created by another job as its source. The copy job cannot run unless its source replica was created successfully. [Figure 92 on page 138](#) illustrates a copy job.

Note: When the copy job creates a snap of an existing clone replica, performing a restore of that snap replica restores to the original source of the clone (the STD) not to the clone itself. Also, copy jobs can only be restored if the storage is CLARiiON or VNX. Symmetrix-based or Open Replicator copy jobs cannot be restored.

Copy jobs are not available with Celerra NFS replications.

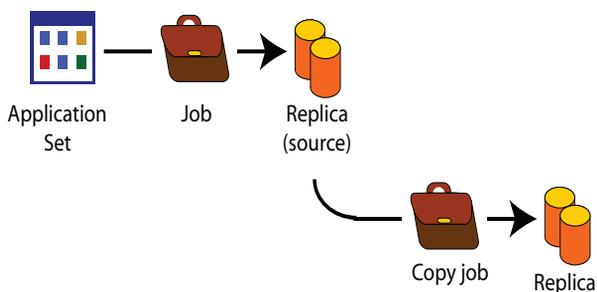


Figure 92 Copy job

- ◆ **Linked copy jobs** — Combines linking and copying so that one job runs, creates a replica, then a second job starts as soon as the first job completes. The second job uses the replica created by the first job as its source. [Figure 93 on page 138](#) illustrates a linked copy job.

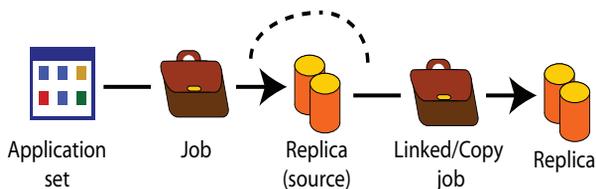


Figure 93 Linked copy job

Copy jobs can copy a replica using SAN Copy, EMC SnapView, TimeFinder/Clone or Replicator to create a copy of a replica.

Once a job exists for an application set, its replicas can be selected as the replication source for a new job (using the Source field in the Job Wizard).

When you link a copy job, Replication Manager will start the linked copy job, immediately after the source job completes successfully, creating a new replica using the first replica as its source.

Replication Manager gives you the ability to link a copy job to a source job. For example, you can easily configure a set of linked jobs that will create a local clone, incrementally copy that clone to a remote array, and then perform a snap of the remote clone.

Configuring a copy job

To configure a job as a copy job:

1. On the **Job Name and Settings** panel, set **Source** to the job whose replicas you want to use as a source.

Note: Setting another job as a source reduces the options you have when selecting a replication technology. Only replication technologies that can be used for copy jobs are listed. Some source replicas reside on storage that prevents you from creating a copy job. In that case, you will not have the option to select the job as a source.

2. On the **Advanced Replication Settings** panel, set the **Replica to be copied** field to one of the following settings:
 - Choose **Most recent** to create a copy of the newest replica created.
 - Choose **Oldest** to create a copy of the oldest replica created.
 - You can also choose a specific replica that is neither the most recent nor the oldest. See [“Choosing a specific replica for a copy job.”](#)

Note: If the replica that meets the criteria of most recent or oldest is mounted read-write when the copy job is run, that replica is not available. In that case, Replication Manager chooses the next best match of the criteria (such as, second most recent or second oldest). If there are no unmounted replicas available then the job fails. An exception is with RecoverPoint: If a bookmark replica of a CDP or CRR copy is mounted read-write, the copy job fails; no attempt is made to find the next suitable time. If there is a *read-only* mount, the replica will succeed, provided the read-only mounted replica meets the criteria for oldest or most recent.

Choosing a specific replica for a copy job

To select a specific copy job:

1. To select a specific copy job that is neither the most recent nor the oldest, on the **Advanced Replication Settings** panel, expand **Replica to be copied**.
2. Check the **Prompt for version for on-demand run** option.
This enables you to choose a specific replica when a copy job is run manually.
3. To choose a replica, right-click on a copy job and select **Run**.
4. In the Run Copy Job window, choose the replica from the list of all the replicas created by the source job.

The selected replica will be used as the source for the copy job.

This option is available only when:

- ◆ the copy job is run manually. If the copy job is scheduled to run automatically, or linked to another job, the source for replication will be the replica that is Oldest or Most recent, depending on what the user has chosen in the copy job properties.
- ◆ the copy job has been created using the TimeFinder/Snap.
- ◆ the source job is a TimeFinder Clone/ TimeFinder Snap.
- ◆ the storage being used is Symmetrix Microcode 5875 or greater.

Configuring a link job

To configure a link job, on the **Starting the Job** panel, choose the application set and job from select lists labelled **Start after job completes**.

There is also a checkbox, **Run this job only if source job succeeds**, that allows you to run this link job conditionally only if the source job succeeds.

Set the remaining fields of the Job Wizard as you see fit. Link jobs and copy jobs are not run when you simulate the source job.

Creating replicas of RecoverPoint targets

Step-by-step

This section describes the two methods for creating a CLARiiON or VNX SnapView replica or a Symmetrix TimeFinder replica of a RecoverPoint target.

Two-phase job

To create a SnapView replica or a Symmetrix TimeFinder replica of a RecoverPoint target using a two-phase job:

1. In the job wizard, select RecoverPoint CDP or RecoverPoint CRR as the replication source.
2. Select SnapView Snap or SnapView Clone as the replication technology.
3. Complete the steps in the job wizard.

Copy job

To create a SnapView replica or a Symmetrix TimeFinder replica of a RecoverPoint target using a copy job:

1. Create a job RecoverPoint CDP, RecoverPoint CRR, or RecoverPoint CLR as the replication source.
2. Create a second job using the first job as the replication source. If the source job is a CLR job, you need to select the CDP copy or the CRR copy.

Managing existing jobs

This section describes how to manage existing jobs. It includes information on how to do the following tasks:

- ◆ Simulate jobs
- ◆ Run jobs on demand
- ◆ Change job schedules
- ◆ Modify jobs
- ◆ Change notifications

Each of these tasks is described in the sections below.

Simulating a job

A simulation is a preliminary execution of a job that helps to identify potential problems without actually creating a replica.

Note: RecoverPoint jobs and copies of RecoverPoint targets cannot be simulated.

During job simulation, Replication Manager tests the job and reveals potential problems including, but not limited to:

- ◆ Software and operating system revisions
- ◆ Missing or nonexecutable scripts
- ◆ Insufficient resources for replica
- ◆ Application set on nonreplication storage
- ◆ Application set on more than one type of storage array
- ◆ Specific storage-array configuration problems
- ◆ Specific database configuration problems

Because job simulations never actually *create* a replica, they allow you to validate application sets and their related jobs without interfering with normal day-to-day storage operations or production data. As opposed to running the job, simulating a job will not:

- ◆ Execute scripts (such as, mount scripts, pre- and post-replication scripts, callout scripts, and so on).
- ◆ Move applications into backup mode.
- ◆ Manipulate mirrors.
- ◆ Back up datafiles from the application database.
- ◆ Run any copy jobs or link jobs associated with the simulated source job.

Running a job simulation

To run a job simulation:

1. Select **Jobs** in the tree panel. Individual jobs are displayed in the content panel.
2. Right-click a job in the content panel and select **Simulate** from the context menu. A progress panel appears and displays validation information. Refer to [Figure 94 on page 143](#) for an example.
3. Click **Details>>** to see the log details generated during the job simulation. Replication Manager simulates the job and provides you with detailed information about the success or failure that you can expect if you attempt to run the job.

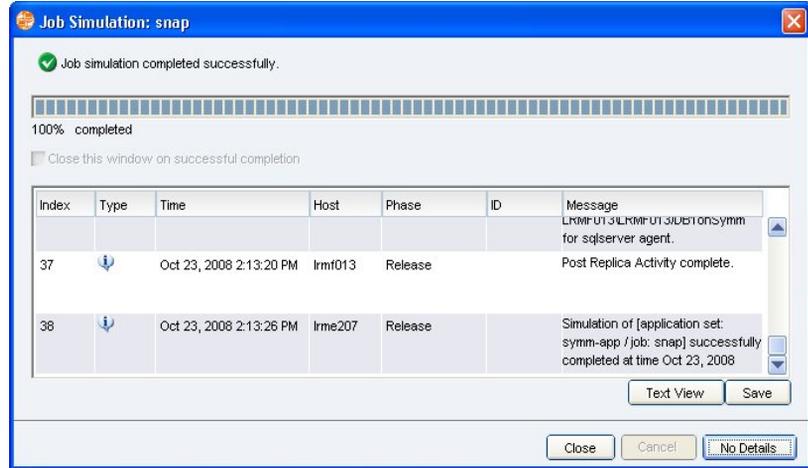


Figure 94 Progress update for job simulation

Note: For specific information about each wizard panel, click **Help**.

Simulation results

In most cases, when the simulation process discovers a failure, the process continues to run and collect troubleshooting information. However, critical failures that would normally cause the replication to fail will also cause the simulation to fail.

In addition to the on-screen simulation results, Replication Manager also stores the results of the simulation in log files.

Running a job on demand

For information about how to run a job on demand, refer to [“Running an existing job on demand”](#) on page 157.

Changing job schedules

Replication Manager allows you to create schedules to control when each job runs.

Note: When a user belonging to an active directory domain, creates a schedule, Replication Manager will run the scheduled job through a local user with the ERM Administrator role.

Note: If your job is BRbackup compliant, do not schedule it here. A BRbackup job can only be run from the command line (using the brbackup commands) on the backup host.

To schedule a job:

1. Select **Jobs** in the tree panel. Individual jobs are displayed in the content panel.
2. Right-click a job in the content panel and select **Schedule** from the context menu. The screen shown in [Figure 95 on page 145](#) allows you to create a new schedule for the job. For more information, refer to the online help.

Note: When you schedule jobs keep in mind that your schedules should ensure that host maintenance such as virus scanning should be planned outside of the window of time when Replication Manager jobs will be running. Additionally, try to schedule jobs so that mount operations do not overlap too much, since contention among these tasks have been known to cause performance issues or failures under certain circumstances.

In addition, in Celerra NAS environments, do not schedule jobs to run while the NASDB Backup process is running. The NASDB Backup process locks the NAS database and prevents Replication Manager from completing required tasks.

New Schedule

Application Set:

Scheduled Job:

Schedule Name:

Daily Frequency

Occurs once at:

Occurs every: hour(s) Starting at: Ending at:

Recurrence Pattern

Once On only

Daily

Weekly

Monthly

Range of Recurrence

Start Date:

End Date:

Create Close Help

Figure 95 Scheduling a job

Modifying a job

To modify a job:

1. Select **Jobs** in the tree panel. Individual jobs are displayed in the content panel.
2. Right-click a job in the content panel and select **Properties** from the context menu. The **Job Properties** window appears.
3. Make the appropriate changes to the job. Click each tab that contains information you want to change and update that information.
4. Click **OK** to complete the operation.

Note: For specific information about each tab, click **Help**.

Changing notifications

Users can instruct Replication Manager to send an email notification whenever a job runs to create a replica. The email can be sent every time the job runs or only when the job fails to run to completion.

In federated environments, more than one email will be sent, one for each production host to give specifics of success or failure on that specific part of the replica associated with each host as well as a general email for status of the entire federated job.

Note: In Windows environments, the email functionality depends on the Internet Information Service (IIS), which should be installed and running before installing the Replication Manager Server in order for the email functionality to work. In Windows 2003, Windows 2008, or Windows 2012, IIS is not automatically installed. The user should verify that IIS is installed on the system.



Figure 96 Select users to notify panel

The panel shown in [Figure 96 on page 147](#) accepts email addresses of individuals who should receive notification when a job runs. Separate email addresses with a semicolon. At the bottom of the screen, you can choose whether Replication Manager should **Send email notifications on failures only** by selecting the checkbox. Clearing the checkbox sends notifications whenever the job runs.

Replication Manager sends an email with the following information:

- ◆ The date and time the job ran
- ◆ Whether it was a success or failure
- ◆ Application set, job, and hostname
- ◆ History of the job

Example of success notification

The following is an example of a success notification email:

```
Automated Email for Successful Replica
Replica Date       : 2008 05 05 11:21:17
Application Set Name : lrmg068-e2k7
Job Name           : job exch
Host Name          : lrmg068
-----
----
2008 05 05 11:22:41 lrmc209 001595 WARNING:
2008 05 05 11:22:41 lrmc209 INFO:Replica 2008 05 05
11:22:41 created from application set lrmg068-e2k7, job
job exch by Administrator.
2008 05 05 11:22:41 lrmc209 INFO:Starting RecoverPoint
checkpoint of [application set:lrmg068-e2k7 / job:job
exch] at time 2008 05 05 11:22:41.
2008 05 05 11:22:41 lrmc209 INFO:This operation can take
a long time. Please be patient.
2008 05 05 11:23:54 lrmc209 102074 INFO:
2008 05 05 11:23:54 lrmc209 INFO:Creating solution for
RecoverPoint replica on RPA RPA@lrmh045.lss.emc.com using
event name lrmc209-em-5547 and RecoverPoint consistency
group lrmg068-EXCH2k7.
2008 05 05 11:23:55 lrmg068 024173 INFO:
2008 05 05 11:24:04 lrmg068 024172 INFO:
2008 05 05 11:24:04 lrmg068 024139 INFO:
2008 05 05 11:24:04 lrmc209 INFO:Verifying software
packages and versions. This may take a few moments ...

... (more details of the processing omitted) ...
```

Example of failure notification

The following is an example of a failure notification email:

```
Automated Email for Failed Replica
Replica Date       : 2008 05 05 11:26:37
Application Set Name : lrmg068-e2k7
Job Name           : crr exch job
Host Name          : lrmg068
-----
----
```

```

2008 05 05 11:26:37 lrmc209 001595 WARNING:
2008 05 05 11:26:37 lrmc209  INFO:Replica 2008 05 05
11:26:37 created from application set lrmg068-e2k7, job
crr exch job by Administrator.
2008 05 05 11:26:37 lrmc209  INFO:Starting RecoverPoint
checkpoint of [application set:lrmg068-e2k7 / job:crr
exch job] at time 2008 05 05 11:26:37.
2008 05 05 11:26:37 lrmc209  INFO:This operation can take
a long time. Please be patient.
2008 05 05 11:27:54 lrmc209 102074 INFO:
2008 05 05 11:27:54 lrmc209  INFO:Creating solution for
RecoverPoint replica on RPA RPA@lrmh045.lss.emc.com using
event name lrmc209-em-5568 and RecoverPoint consistency
group lrmg068-EXCH2k7.
2008 05 05 11:27:56 lrmg068 024173 INFO:
2008 05 05 11:28:04 lrmg068 024172 INFO:
2008 05 05 11:28:04 lrmg068 024139 INFO:
2008 05 05 11:28:05 lrmc209  INFO:Verifying software
packages and versions. This may take a few moments ...

... (more details of the processing and errors omitted)
...

```

Restrictions on simultaneous replications

When determining how often to run a job, how to schedule jobs, creating rotations and similar topics, remember that there are certain restrictions that prevent you from running more than one job on the same application set simultaneously. These restrictions apply:

- ◆ If a job creates a TimeFinder/Snap of a TimeFinder/Clone or TimeFinder/Mirror replica, no other Symmetrix job can run simultaneously from the same application set.
- ◆ If a job creates a TimeFinder/Clone of a TimeFinder/Clone, no other Symmetrix job can run simultaneously from the same application set.

Deleting a job

Before you can delete a job, all schedules related to that job must be deleted. After the schedules have been deleted, you can delete the job by following these steps:

1. Select **Jobs** in the tree panel. Individual jobs are displayed in the content panel.
2. Right-click a job in the content panel and select **Delete** from the context menu.

3. Click **Yes** when Replication Manager asks **Are you sure you want to delete job *jobname*.**

Effect of failed job on clone synchronization

Replication Manager clone LUNs remain synchronized if replication failed after the sync for reasons such as application freeze failure. The next run of the job will select the synchronized pair to create a replica.

Managing replica rotations

When you define a replica rotation, you specify the maximum number of replicas that can exist simultaneously. Replication Manager uses a schedule to determine how often replicas get created. Replication Manager creates the replicas at a specified frequency, and can maintain no more than the maximum number of replicas. When Replication Manager has used the maximum number of replicas, older replicas are unmounted (if necessary) and deleted to free disk space for the next replica in the rotation. Replication Manager does not consider failed or simulation replicas when calculating whether the rotational set has reached its maximum number of replicas. You can manage replica rotations from the Job Wizard panel.

For example, if you want to create a replica rotation that provides one replica for each of the last three days, you could define the rotation as follows:

- ◆ Replica Count: 3
- ◆ Scheduled Replica Frequency: Once a day

The replica rotation would work as shown in [Figure 97](#) on page 151.

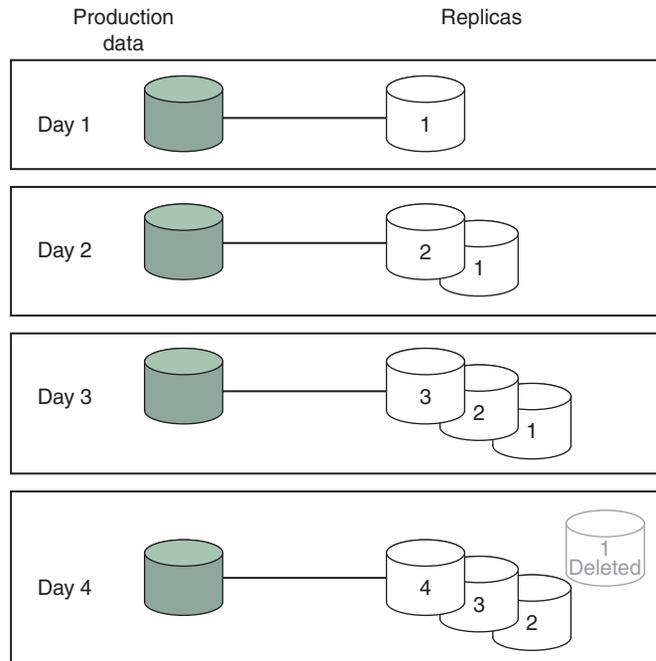


Figure 97 Replica rotations

Note: If you specify three as the maximum number of replicas, there will be times during the standard rotation period when only two valid replicas exist (during the replica creation process). If you want to ensure that you have at least three valid replicas at all times after the rotation is established, you should specify a maximum of four, not three.

To define a replica rotation:

- ◆ When creating a job, specify the rotation definition in the **Job Name and Settings** panel of the **Job Wizard**.
- ◆ When modifying an existing job, select the **Replication** tab under **Job Properties**.

Note: For specific information about the replica rotation, click **Help**.

For information on modifying a replica rotation, refer to [“Modifying replica rotation”](#) on page 302.

Managing job schedules

You can specify when to perform each job by creating a schedule for existing jobs.

A schedule defines:

- ◆ The date and time when the job should first occur
- ◆ How often the job should recur

Scheduling a job

To schedule a job:

1. Select **Jobs** in the tree panel. Individual jobs appear in the content panel.
2. Right-click a job in the content panel and select **Schedule** from the context menu. The **New Schedule** screen allows you to specify when to perform each replication.

Special considerations when scheduling on Windows 2008 Hosts

On Windows 2008 hosts, modifications are needed before you can create a job schedule or scheduled task. For example, you may encounter the following error when scheduling a task:

```
026432 ERROR:An internal error occurred on the
Replication Manager host as follows: Unable to
create/update schedule: test Please look in the logs
for more details.
```

In the event of this error, you must disable the Network access: Do not allow storage of credentials or .NET Passports for network authentication policy. See EMC knowledgebase article emc239567 for details.

Modifying a schedule

To modify a scheduled job:

1. Select **Jobs** in the tree panel. Individual jobs appear in the content panel.
2. Right-click a job in the content panel and select **Properties**. The **Job Properties** window appears.
3. Click the **Startup** tab. Existing schedules are listed.
4. Select the schedule you want to modify and click **Modify**.
5. Make the appropriate modifications to the schedule.

6. Click **OK** to complete the operation.

Note: You cannot modify startup options for a job if you are using SAP BRbackup compliant replicas.

Deleting a schedule

To delete a schedule:

1. Select **Jobs** in the tree panel. Individual jobs appear in the content panel.
2. Right-click a job in the content panel and select **Properties**. The **Job Properties** window appears.
3. Click the **Startup** tab. Existing schedules are listed.
4. Select the schedule you want to delete and click **Remove**.
5. Click **OK** to complete the operation.

Stopping a schedule

To stop a scheduled task:

1. Select **Jobs** in the tree panel. Individual jobs appear in the content panel.
2. Right-click a job in the content panel and select **Properties**. The **Job Properties** window appears.
3. Click the **Startup** tab. Existing schedules are listed.
4. Clear the checkbox in the **Enabled** column next to the job you want to disable.
5. Click **OK** to prevent the schedule from running jobs.

Restarting a schedule

To restart a disabled schedule:

1. Select **Jobs** in the tree panel. Individual jobs appears in the content panel.
2. Right-click a job in the content panel and select **Properties**. The **Job Properties** window appears.
3. Click the **Startup** tab. Existing schedules are listed.
4. Select the checkbox in the **Enabled** column next to the job you want to enable.
5. Click **OK** to reenable the schedule.

All users can perform daily operations on the application sets to which they have access. This chapter describes how to perform daily operations in the following sections:

- ◆ Introduction 156
- ◆ Running an existing job on demand 157
- ◆ Storage arrays and devices 158
- ◆ Unmounting replicas on demand 160

Introduction

Users can run a previously configured job on an application set to which they have been granted access. A Power User, Database Administrator, Power DBA, or ERM Administrator configures the job in advance and any user (including an Operator) can run that job when the need arises.

Running an existing job on demand

To run a previously configured job on demand:

1. Select **Jobs** in the tree panel. Individual jobs appear in the content panel.
2. Right-click a job in the content panel and select **Run**. The job starts processing after you confirm that you want to run the job.

Note: Although the job has been configured previously, any user with any role except Operator, who has been granted specific access to the job can modify that job for an on-demand run. Operators cannot modify a job.

Monitoring running tasks

To monitor running tasks (including jobs and other background tasks), select **Active Tasks**. Details about individual tasks appear in the content panel.

Note: Sometimes Replication Manager adds background tasks to the task list. These tasks are listed under the user Internal Administrator. For example, an Internal Administrator job appears during an automated discover storage process. Please ignore these tasks.

Canceling a task

To cancel a running task:

1. Select **Active Tasks**. Individual tasks appear in the content panel.
2. Right-click the task in the content panel and select **Cancel Task**.
3. Click **OK** to confirm the cancellation.
4. Click **OK** to accept the cancellation successful message.

Note: You can also cancel a running job from the progress window associated with that job. Remember that sometimes a job has progressed too far to be canceled.

Storage arrays and devices

You can examine the storage array devices available to Replication Manager for replications. To do so:

1. Expand **Storage Services**. The tree lists each kind of storage array as follows:
 - **Celerra iSCSI** — Listed using the Celerra IQNs. Replication Manager can discover Celerras if they are zoned and connected to a registered host.
 - **VNXe iSCSI** — Listed using the VNXe IQNs. Replication Manager can discover VNXe arrays if they are zoned and connected to a registered host.
 - **Celerra or VNX File NFS** — Listed using the hostname of the Celerra; For VNX, it is listed under the VNX serial number. Control stations are manually entered by the storage administrator. Right-click **Storage Services**, and select **Add Storage**, then add each Celerra or VNX Control Station.
 - **CLARiiON storage arrays** — Listed using the CLARiiON serial number for the CLARiiON array. Replication Manager can discover CLARiiON storage arrays if they are zoned and connected to a registered host.
 - **VNX storage arrays** — Listed using the VNX serial number for the VNX array. Replication Manager can discover VNX storage arrays if they are zoned and connected to a registered host.
 - **Symmetrix storage arrays** — Listed using their Symmetrix serial number. Symmetrix storage arrays are discovered based on which hosts have been added to Replication Manager and on the Add Storage Wizard.
 - **RecoverPoint** — Listed by the fully-qualified DNS name or IP address of the RPA.

Note: For more information on how to prepare storage arrays and services for use with Replication Manager, refer to the *EMC Replication Manager Administrator's Guide*.

Storage services are listed as shown in [Figure 98 on page 159](#).



Figure 98 Storage services

2. Select a storage service on the tree shown in [Figure 98 on page 159](#). The content panel displays the following information about each device in the array:
 - **Device** — Device name
 - **State** — State of the device (for example, In Use, Not In Use, Excluded)
 - **Size** — Size of the device in GB
 - **Type** — Type of storage device (for example, Local Clone, Local snapshot, and so on)
 - **Visible to Hosts** — List of hosts that can see the device
 - **Capabilities** — Replication technologies in which this device is used

Unmounting replicas on demand

Unmounting a replica on demand requires that you select an existing replica that is currently mounted in the main window and explicitly unmount that replica.

To unmount a replica:

1. Expand **Application Sets**.
2. Select the application set with the replica you want to unmount. Replicas are displayed in the content panel.
3. Verify that the replica is mounted.
4. Right-click the replica and select **Unmount**.



CAUTION

The unmount process begins immediately without further input from the user.

Note: Replicas that are not currently mounted do not offer the **Unmount** option and have a different icon.

Replication Manager helps you mount, restore, and recover information. This chapter describes the following tasks:

- ◆ Mounting replicas 162
- ◆ Mounting replicas to an alternate host 168
- ◆ Mounting replicas to the production host 182
- ◆ Mounting multiple replicas simultaneously 184
- ◆ Using mount and backup scripts 186
- ◆ CLARiiON or VNX static mount 189
- ◆ Mounting replicas to an alternate MSCS cluster 191
- ◆ Mounting replicas to a production MSCS cluster 192
- ◆ Mount and restore on UNIX Clusters 193
- ◆ VMware mount and restore 200
- ◆ RecoverPoint mount and restore 215
- ◆ Unmounting a replica 220
- ◆ Replica mount performance guidelines 222
- ◆ Restoring from a replica 224
- ◆ Using application callout scripts 235
- ◆ Troubleshooting mount failures on the array 245

Mounting replicas

Replication Manager can mount snaps, clones, or bookmark replicas across various EMC storage arrays. Each replica type has specific features and limitations related to mount operations.

Figure 99 on page 162 shows a graphical representation of the steps necessary to mount a replica.

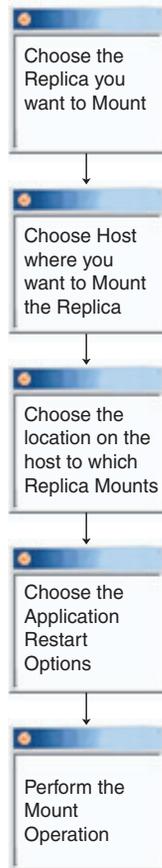


Figure 99 Mount process

Mount restrictions and limitations

Note the following mount restrictions:

- ◆ Remember when mounting a replica of a federated application set, it is necessary to complete a separate set of mount options for each host that is part of the federated replica. Multiple hosts can be mounted to a single mount host or different mount hosts. The user configuring the mount options is responsible for preventing device conflicts when performing such a mount.
- ◆ Mounts and restores may fail when the application set contains nested mount points. For instance, if some of the files in the application set are on L:\ and other files are on L:\SG1DBMP (where SG1DBMP is a mount point), mounts and restores of the associated replicas may fail.
- ◆ The **vxconfigbackupd** daemon detects changes in the Veritas environment, and makes backup copies of volume groups in `/etc/vx/cbr/bk`. As a result of mounting replicas in a Veritas environment, entries in this directory may accumulate over time. The administrator should monitor this directory to ensure that it does not fill up the root file system.

Host-based limitations

In general, you can mount a replica onto mount hosts that have identical environments as the production host. There are a few exceptions as follows:

- ◆ Replicas of Windows 2003 and Windows Server 2008 data can be cross-mounted between 32-bit and 64-bit systems as long as those systems are otherwise supported.
- ◆ File system replicas of Windows 2003 data can only be mounted on Windows 2003 mount hosts.
- ◆ File system replicas of Windows Server 2008 data can be mounted on Windows Server 2008 *and* Windows Server 2012 mount hosts.
- ◆ File system replicas of Windows Server 2012 data can only be mounted on Windows Server 2012 mount hosts.
- ◆ File system replicas of Solaris 2.8 data can be mounted onto Solaris 2.8 and Solaris 2.9 mount hosts.
- ◆ File system replicas of Solaris 2.9 data can only be mounted on Solaris 2.9 mount hosts.

Table 7 on page 164 describes the features and limitations available when mounting various replica types.

Table 7 Mount features and limitations by technology (page 1 of 3)

Technology	Mount features	Mount limitations
TimeFinder/Mirror	Mountable to hosts on which that BCV has been made visible. On Windows 2003, Windows 2008, and Windows 2012 hosts, can mount replicas read-only to facilitate backup.	Visibility depends on zoning and bin file settings.
TimeFinder/Mirror (remote)	Mountable to hosts on which the BCVs have been made visible. In a Solaris environment, LUN visibility can be enabled by setting an environment variable. See “LUN visibility for Symmetrix devices in a Solaris environment” on page 167 for more details.	Visibility depends on zoning and bin file settings. Can mount the replica but the replica cannot be restored.
TimeFinder/Snap	Mountable to hosts on which the VDEV device has been made visible. In a Solaris environment, LUN visibility can be enabled by setting an environment variable. See “LUN visibility for Symmetrix devices in a Solaris environment” on page 167 for more details.	Visibility depends on zoning and bin file settings.

Table 7 Mount features and limitations by technology (page 2 of 3)

Technology	Mount features	Mount limitations
TimeFinder/Clone	Mountable to hosts on which that clone has been made visible. On Windows 2003, Windows 2008, and Windows 2012 hosts, can mount replicas read-only to facilitate backup. In a Solaris environment, LUN visibility can be enabled by setting an environment variable. See “LUN visibility for Symmetrix devices in a Solaris environment” on page 167 for more details.	Visibility depends on zoning and bin file settings.
Celerra or VNXe iSCSI SnapSure Replicas	Mountable on hosts that are logged into the IQN(s) of the production LUNs on the source Celerra or VNXe.	
Celerra or VNX NFS SnapSure Replicas	Mountable on hosts that have IP access to the control station of the production Celerra or VNX.	
Celerra or VNXe iSCSI Replicator Replicas	Mountable on hosts that are logged into the IQN(s) of the target LUNs on the target Celerra or VNXe.	
Celerra or VNX NFS SnapSure Replicas	Mountable on hosts that have IP access to the control station of the production Celerra or VNX	
SnapView Clone	Mountable on hosts that have access through the appropriate storage processor (SP A or SP B) and storage group. Choose Create and mount a snap of the replica to cause changes you make to the mounted replica to be discarded on unmount.	Some configurations require the use of static LUN visibility to mount in CLARiiON or VNX environments.

Table 7 Mount features and limitations by technology (page 3 of 3)

Technology	Mount features	Mount limitations
SnapView Clone (remote)	<p>Mountable on hosts that have access through the appropriate storage processor (SP A or SP B) and storage group.</p> <p>Choose Create and mount a snap of the replica to cause changes you make to the mounted replica to be discarded on unmount.</p>	Some configurations require the use of static LUN visibility to mount in CLARiiON or VNX environments.
SnapView Snap	Mountable on hosts that have access to the CLARiiON or VNX.	Some environments require static snapshot visibility including a separate snapshot device for each replica.
SnapView Snap (remote)	Mountable on hosts that have access to the CLARiiON or VNX.	Some environments require static snapshot visibility including a separate snapshot device for each replica.
SAN Copy	Mountable on hosts that have access to the target CLARiiON or VNX.	Multiple array visibility to the source and target hosts is necessary.
Open Replicator	Mountable on hosts that have access to the target CLARiiON or VNX.	Only read-only mounts or Create and mount a snap of the replica (which discards changes on unmount) are supported.
RecoverPoint	Mountable to hosts on which the target volumes have been made visible.	<p>Visibility depends on zoning.</p> <p>Changes made to a RecoverPoint replica while it is mounted are not persistent; they are lost when you unmount the replica.</p> <p>You can mount only one RecoverPoint replica at a time per application set per site.</p>

Difference between on-demand and job mounts

Note the differences in behavior between an on-demand mount and a mount as part of a job:

- ◆ On-demand mount: If you try to mount a replica to a location where another replica from the same application set is already mounted, the mount fails. You should mount to a different location or unmount the existing replica first.
- ◆ Mount as part of a job: If the mount is to a location where another replica from the same application set is already mounted, the original replica is unmounted, then the new replica is mounted.

LUN visibility for Symmetrix devices in a Solaris environment

In a Solaris environment, Symmetrix devices need not be statically visible to a mount host in order to be a replication target. To enable visibility for Symmetrix devices in a Solaris environment, you must set an environment variable on the mount host to indicate the Masking View under which the devices need to be added. Once the masking view is specified, Replication Manager can choose any appropriate Symmetrix device for replication irrespective of its visibility to the mount host.

This feature is supported with Mirocode 5874 and above.

Note: EMC recommends you create a Replication Manager storage pool for the devices to be selected from.

Enabling dynamic LUN visibility

- Check whether the Initiator group and port group are configured correctly.

1. On the client machine, navigate to **Mount Host > Properties > Advanced**

The environment variable is `EMC_ERM_MASKING_VIEW_NAME_SID` where `SID` is the associated Symmetrix ID.

2. In the **Value** field, specify the masking view.

Symmetrix devices are added to this view regardless of their visibility to the mount host.

Note: Once the environment variable is set on the mount host, on an unmount, the selected devices are unmasked if they belong to the specified masking view.

Specify hosts for storage discovery

During storage discovery, instead of allowing Replication Manager to discover storage through all the hosts attached to the array, you can select the host on which to perform the discovery. You can select the host of your preference for each array.

In the Add Storage wizard, after you select the arrays, you can see a listing of all the hosts attached to the selected arrays. Select a host for each array to discover storage.

Storage visibility validation during mount operations

Before doing a mount (on-demand or on-job), a mount failure could occur because of insufficient devices visible to the mount host. To avoid the error, you can validate the visibility of storage on the selected mount host.

In case of on-demand mount, validate the visibility of replica devices on the selected mount host before performing the actual mount operation.

In case of on-job mount, validate the visibility of adequate devices on the selected mount host before running the activity. For on-job mount, visibility validation can be done only after creation of a job.

Mounting replicas to an alternate host

When you mount a replica, you mount to an alternate host or to the original production host. (For SAP, mounting to an alternate host is forced). You can also choose (when creating the job) whether a mount failure should cause the entire replication to fail or not.

This section describes how to mount to an alternate host:

- ◆ In the same location as the production host
- ◆ In an alternate location using alternate root path
- ◆ In an alternate location using path mapping

Note: If a Replication Manager mount fails, the replica creation process may still succeed if you select the option during the job creation that allows that to occur. This results in the creation of an unmounted replica.

Before describing the different mount options, consider these general alternate mount issues.

Alternate mount considerations

When mounting to an alternate host, remember:

- ◆ Mount host SCSI-based file system versions must be identical to the production host SCSI-based file system versions unless an explicit exception is listed in [“Host-based limitations” on page 163](#).
- ◆ Mount host logical volume manager versions must be identical to the production host logical volume manager versions.
- ◆ Any changes made to the mounted replica are persistent, except in the following cases:
 - Replica with source data on a CLARiON or VNX array that you have specifically selected not to retain changes during the mount.
 - Read-write mounts of a read-only replica of a Celerra network file system.
 - Changes made to a mounted Celerra or VNXe iSCSI replica are discarded when the replica is unmounted.
 - Celerra NFS allows you to mount the replica read-only (regardless of whether the replica itself is read-write or read-only). In this case, the file system is not writable at all.

Persistent mount changes show up if the replica is unmounted and restored or remounted somewhere else. Changes that are not persistent are lost upon unmount.

- ◆ When mounting logical volumes, all of the devices in the volume group are imported.
- ◆ When mounting an NFS, if there are multiple mount paths within a file system (exported using different export names), including only one of those mount paths in an application set causes Replication Manager to replicate the entire file system but only mount the selected mount path.
- ◆ Restoring a replica of a given mount path (as described in the bullet above) restores the entire file system.
- ◆ It is possible to mount a replica from one MSCS cluster to a single node of an alternate cluster. Special setup is required for this operation. Refer to the [“Mounting replicas to an alternate MSCS cluster” on page 191](#) for more information.
- ◆ Mounting an Oracle instance that was created from an OPS cluster creates a standard Oracle instance, not a separate cluster.

- ◆ Choose the appropriate database mode for mounts and restores. Depending on the application, different mount options may be provided for alternate mounts compared to mounts to the production host.
- ◆ All hosts of the same operating-system type as the production host (to which you have access) appear on the list as mount host choices.
- ◆ The user performing the mount must have access to all storage associated with the mount or the mount operation fails.
- ◆ The primary group for the OS username provided needs to be the same on both production and mount host (GIDs must match). The username/UIDs can be different.

Mounts to an alternate mount host with Windows

Replication Manager allows VSS (Volume ShadowCopy Services) replicas taken on Windows Server 2003 SP2, Windows Server 2008, and Windows Server 2012 systems to be mounted on alternate Windows hosts. When the VSS mount is between different CLARiiON arrays, ensure that the FLARE version is the same between the arrays or VSS mounts can fail.

Mounts to an alternate mount host with VxVM

When mounting a replica on a mount host that uses Veritas Volume Manager, the production and mount hosts are not required to match with respect to Veritas enclosure-based naming. Replication Manager can convert to and from enclosure-based naming when mounting to an alternate host (as needed).

Considerations for VxVM volumes on Windows 2003 in an MSCS environment

- ◆ To restore VxVM volume in a Windows 2003 MSCS environment, delete volume group resource from the resource group.
- ◆ When restoring VxVM volumes on Windows 2003, make sure all the volumes in the disk group are selected.
- ◆ If the VxVM volume group is a cluster resource, the VxVM volume cannot be restored.
- ◆ Extended maintenance mode is not supported for a VxVM volume group resource.

Configuring mount hosts to work with storage arrays

Before you can mount a replica to an alternate mount host, the mount host must be configured to work with the storage array where the replica resides. For more information, refer to the *EMC Replication Manager Administrator's Guide* for information about how to configure mount hosts to work with your storage arrays. The following considerations apply:

- ◆ The host must have access to both the source and target array when using the SAN Copy feature to create replicas of data from one array onto a separate array.
- ◆ When mounting a federated application set all hosts that are part of the application set must have visibility to all data that is part of the federated application set or the mount fails.

Mounting Windows replicas to an alternate host

When you first mount a LUN from a CLARiiON or VNX storage array to a Windows mount host, the mount host may require you to reboot to rebuild operating system information.

Note: If a Windows host requests that you reboot, you should reboot the host and run the mount operation a second time. A reboot is required only with the first mount to that alternate mount host and is not necessary during subsequent mounts.

Mounting data in the same location as production

You can choose to mount the replica of the production data to an alternate mount host in the same location that it occupied in the production host. [Figure 100 on page 172](#) shows a graphical representation.

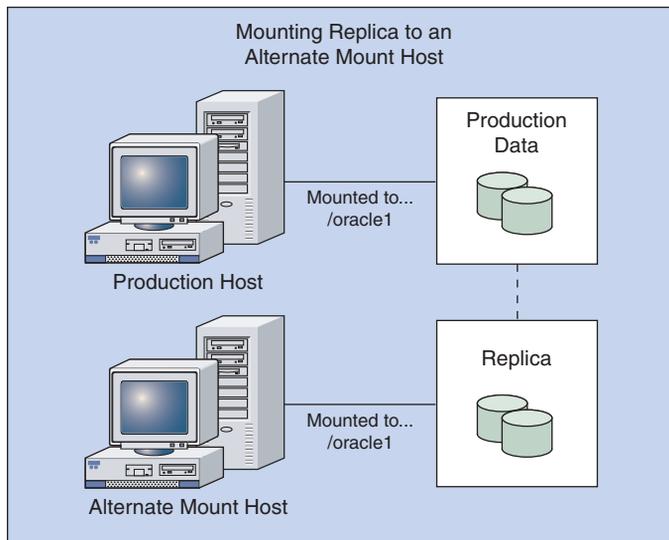


Figure 100 Mounting to an alternate host at the same location as the production host

To perform a mount on an alternate host to the same location as the production host:

1. Expand **Application Sets**.
2. Select the application set with the replica you want to mount. Replicas are displayed in the content panel.
3. Right-click the replica you want to mount and select **Mount**. The mount replica wizard appears and asks which replica you want to mount and to which host you want to mount the replica.
4. Select a host other than the production host and expand the **Mount host** part of the tree as shown in [Figure 101](#) on page 173.

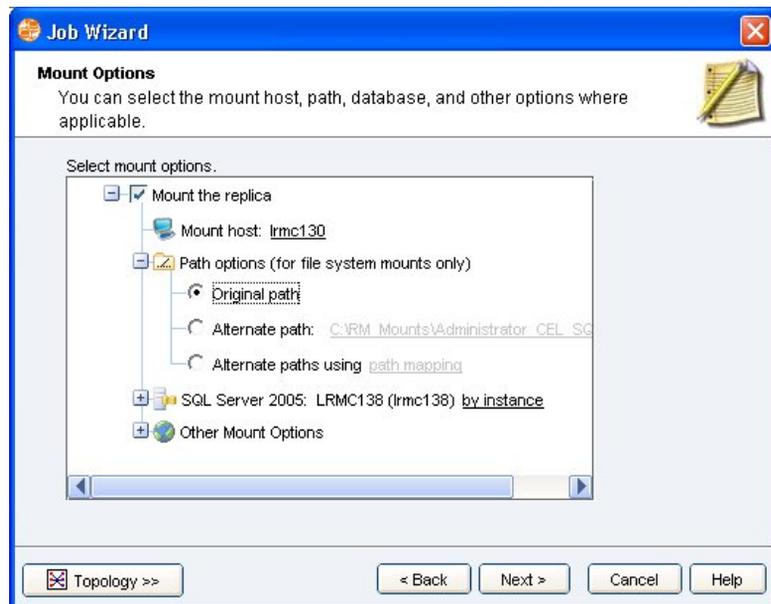


Figure 101 Mounting to the original path

5. Select **Original Path** to mount to the original location on the alternate mount host.

Note: This selection is forced when using SAP BRbackup compliant replicas.

6. For information on how to complete the remaining information in the tree, access the appropriate help screens.

Note: If you specify multiple sets of application data in a single application set, Replication Manager requires that you specify mount settings for each set of application data that you mount.

Mounting using alternate path

You can mount the replica to an alternate host or production host; however, the option also changes the path where the data is located on that host by prepending a single alternate path to the existing paths on the system. Refer to [Figure 102 on page 175](#) for a graphical representation.

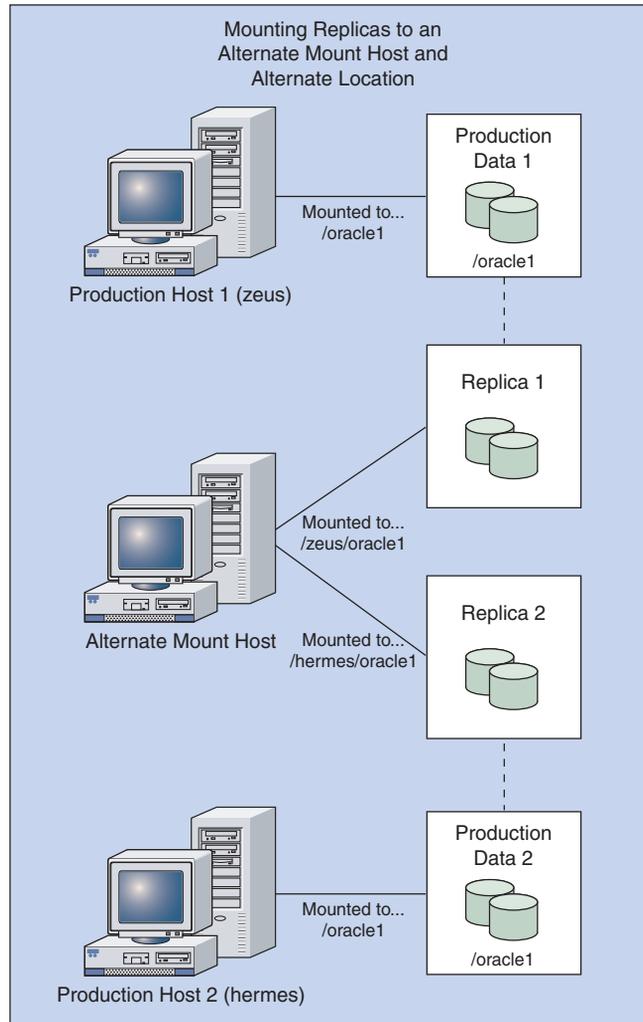


Figure 102 Mounting by using an alternate path

The alternate path is a partial path that is prepended to the original path where the data resides on the production host. Mounting by using an alternate path preserves the original path structure as a subdirectory under the alternate path.

To perform a mount using alternate root paths:

1. Expand **Application Sets**.
2. Select the application set with the replica you want to mount. Replicas are displayed in the content panel.
3. Right-click the replica you want to mount and select **Mount**.
4. Select a host other than the production host.
5. Select **Alternate Path** to mount using an alternate path and enter the alternate path to prepend in the available field, as shown in [Figure 103 on page 176](#). For examples of alternate paths, see the help file.

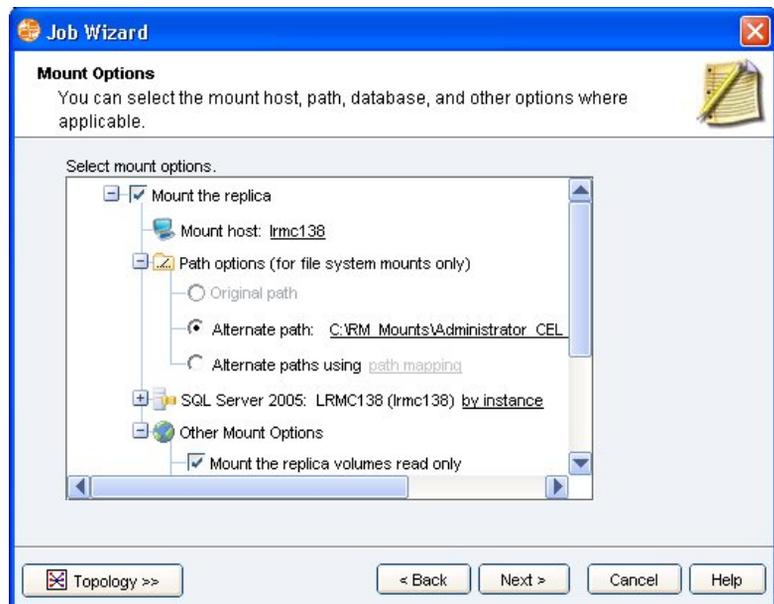


Figure 103 Entering the alternate path

Note: In Windows, the drive specified in an alternate path must already exist on the mount host. Do not include a semicolon (;) or forward slash (/) in any Windows alternate path. UNIX paths must begin with a forward slash.

Access the appropriate help screens for specific help on each wizard screen.

Note: If you specify multiple sets of application data in a single application set, Replication Manager requires that you specify mount settings for each set of application data that you mount.

Mounting using an alternate path: Windows example

When you specify an alternate mount path in a Windows environment, Replication Manager must use some convention to deal with drive letters that are part of the path. Here is an example to describe how this is handled.

Consider a replica that contains three volumes that are mounted to the production host using the following drive letters:

- ◆ H: \
- ◆ J: \
- ◆ L: \

If you use an alternate path to mount this replica (either back to the production host or to an alternate host), you might specify something like the following as the alternate path:

```
C:\MyAlternateMount
```

In this case, the resulting drive structure after the mount will be as follows:

- ◆ H: \ will be mounted to C:\MyAlternateMount\Hdrive
- ◆ J: \ will be mounted to C:\MyAlternateMount\Jdrive
- ◆ L: \ will be mounted to C:\MyAlternateMount\Ldrive

To avoid the addition of the Hdrive, Jdrive, and Ldrive convention used in the Windows alternate path implementation, consider using path mapping. See the example in [Table 9 on page 178](#). Using a path map instead of an alternate path gives you more control over Windows mounts because you can specify the exact drive and pathname where you want your data mounted.

Mounting using path mapping

The path mapping option mounts the replica to an alternate host using a path mapping table to set the alternate locations. When you use a path mapping table, you have more control over where data is located.

You can specify a substitute path for each mount point that exists on the production host, which can replace that mount point with a different pathname.

Note: If you choose to replace only part of a pathname with a path mapping table, the substitute pathname must start at the beginning of the original path. For example, if the original path is `/prod1/acct/data/` you can create a substitution pathname for `/prod1/acct/` but you cannot create a substitution pathname for `/acct/data/` unless you substitute the entire pathname.

Entries in the path mapping table must meet the following guidelines to be considered valid:

- ◆ Filenames should not be included in the path mapping table.
- ◆ Subdirectories under mount points cannot be specified.
- ◆ On Windows, existing mapped network drives should not be used. The mount operation will succeed but the specified drive letter will be masked by the network drive and will not be available until the network drive is unmapped.

Consider the sample path mapping tables for UNIX in [Table 8 on page 178](#).

Table 8 Sample UNIX path mapping table

Substitute	For
<code>/data/emc/</code>	<code>/myEMCData/</code>
<code>/controlfiles/</code>	<code>/alpha/cntrlFiles/</code>

Also consider the next sample of a path mapping table in a Windows environment, shown in [Table 9 on page 178](#).

Table 9 Sample Windows path mapping table drives to other drives

Substitute	For
<code>H:\</code>	<code>I:\</code>
<code>J:\</code>	<code>K:\</code>
<code>L:\</code>	<code>M:\</code>

In the example path map in [Table 9 on page 178](#), volumes originally mapped to `H:\`, `J:\`, and `L:\` will be remapped to `I:\`, `K:\`, and `M:\`.

I:\, K:\, and M:\ should not exist on the mount host before you run the job, Replication Manager will create them and map the data to them as part of the mount operation.

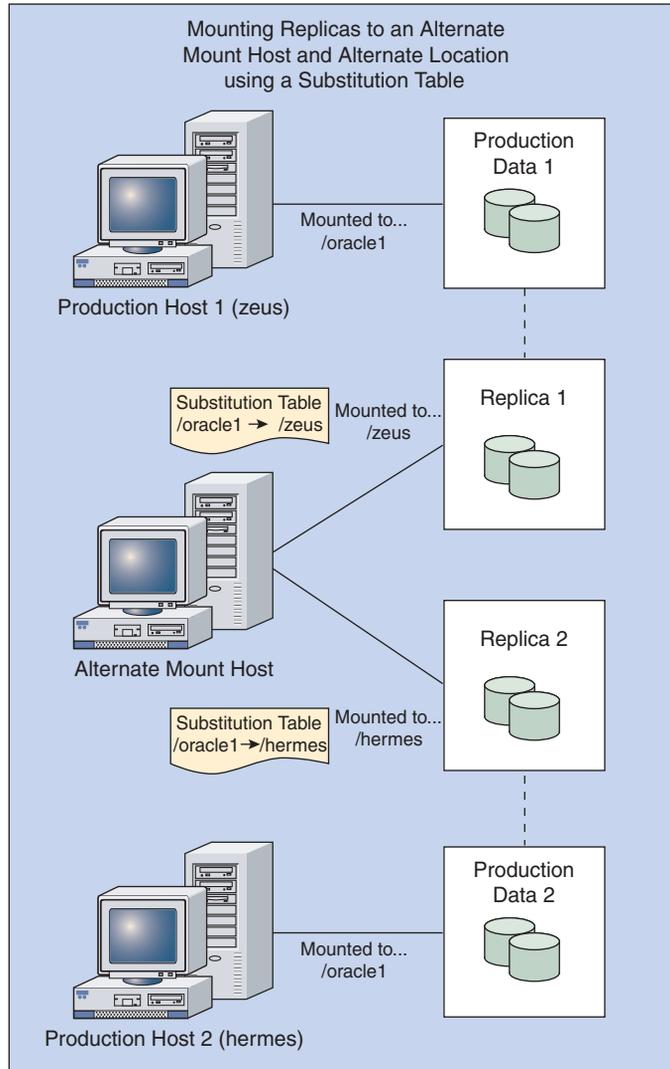
Table 10 Sample Windows path mapping table drives to mountpoints

Substitute	For
H:\	C:\MyMount\data
J:\	C:\MyMount/logs
L:\	C:\MyMount\ctrlFiles

In the example pathmap in [Table 10 on page 179](#), data originally mapped to H:\, J:\, and L:\ will be mapped to three separate mount points (C:\MyMount\data, C:\MyMount/logs, and C:\MyMount\ctrlFiles.) There will be no Hdrive, Jdrive, or Ldrive mountpoint added as there was in the alternate path example. Please note that the drive letter of the destination must be mapped on the mount host, but the mountpoints should not exist prior to the mount.

If an entry in the path mapping table is not valid at the time of the mount (such as a table entry that ends in a filename), the mount will continue as if that entry did not exist. In that case, Replication Manager attempts to mount the data to the same location as it occupied on the production host. If that location is not free, the mount operation fails.

Refer to [Figure 104 on page 180](#) for a graphical representation of a UNIX alternate mount using a path map.



RM-00027

Figure 104 Mounting two production hosts to the same alternate host using path mapping tables

To perform two different mounts from different application sets that both mount to the same alternate mount host:

1. Expand **Application Sets**.
2. Select the particular application set that contains the replica you want to mount.
3. Right-click the replica you want to mount and select **Mount**.
4. Select a host other than the production host.
5. Select **Alternate paths using path mapping** to mount using pathname substitution. Enter the original location and the new location in the substitution table in [Figure 105 on page 181](#).

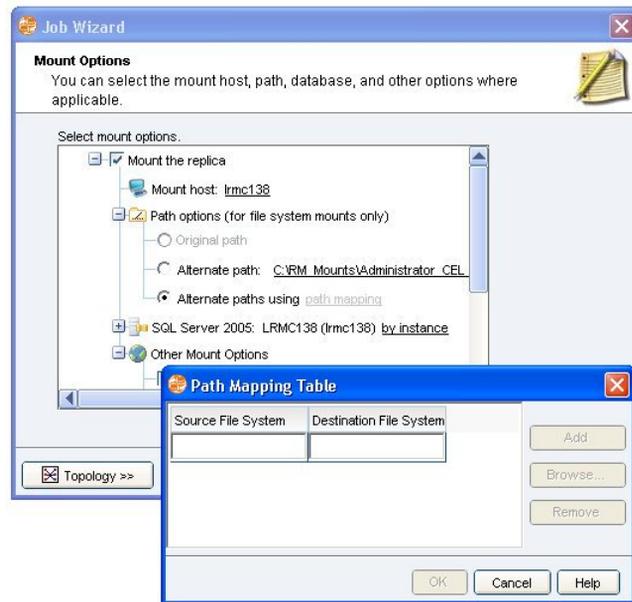


Figure 105 Alternate paths using path mapping table

Access the appropriate help screens for specific help on the mount screen.

Note: If you specify multiple sets of application data in a single application set, Replication Manager requires that you specify mount settings for each set of application data that you mount.

Mounting replicas to the production host

When you choose to mount a replica to the original production host, you must select an alternate path where that replica data can reside. Doing so ensures that the mount does not fail, because the production data is already mounted on that host in the original location.

Production mount considerations

The following additional restrictions apply when you mount to the original production host:

- ◆ (With Oracle) When mounting to the production host with the production database running, you must choose an alternate SID and mount the database without recovering it.
- ◆ (With Oracle) Bring down the production database and rename the database you are mounting when mounting to the production host in read-only or read/write mode. Oracle does not allow two copies of the same database to run simultaneously on the same host (even if the Oracle SIDs are different).
- ◆ (With UDB and SQL Server) Choose an alternate instance when mounting to the production host (you must create the alternate instance manually, in advance of the mount operation).

Note: If you are using consistent-split technology, refer to [Chapter 6, "Using Consistent Split,"](#) for more information and considerations associated with mounting applications using that technology.

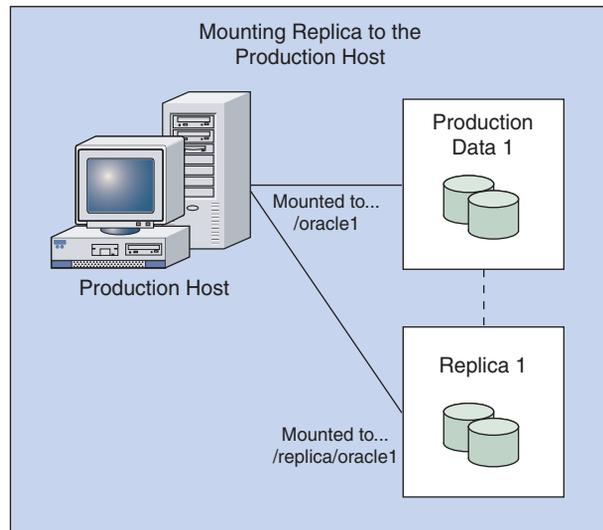
If you follow these guidelines, Replication Manager changes the appropriate application data to allow you to mount a separate copy of the application data to a new location without affecting the production instance.

The ability to mount data onto the same production server can reduce the overall number of servers needed to review data. No extra mount server is necessary. Data can be mounted to an alternate location using alternate path or a path mapping table, similar to mounts to an alternate mount host.

Other modifications to the mounted replica may be necessary (depending on which application is associated with the data that you are mounting). For more information about specific applications, refer to the appendixes at the end of this guide.

Mounting to production using alternate path

You can mount the replica to the original production host, and change the path where the data is located on that host based on a single alternate path that gets prepended to the existing paths. Refer to [Figure 106 on page 183](#) for a graphical representation.



RM-000028

Figure 106 Mounting to a production host using an alternate path

To mount a replica on demand to the original production host:

1. Expand **Application Sets**.
2. Select the particular application set that contains the replica you want to mount.
3. Right-click the replica you want to mount and select **Mount**.
4. Select the production host.
5. Select **Alternate path** and enter the alternate path in the available field.

Note: If you specify multiple sets of application data in a single application set, Replication Manager requires that you specify mount settings for each set of application data that you mount.

Mounting multiple replicas simultaneously

You can also choose to simultaneously mount more than one replica from the same application set. In other words, you can mount a replica of the information that was taken at 2:00 p.m. and also mount a replica that was created at 3:00 p.m. Further, you can mount these replicas onto the same host (production or alternate mount host). In some circumstances, you may have to use another instance of the application to access the data that was mounted. [Table 11 on page 184](#) describes the rules.

Note: You can mount only one RecoverPoint replica at a time per application set, per site.

Table 11 Rules for mounting multiple replicas simultaneously (page 1 of 2)

Application/file system	Applicable rules
Oracle	Only one instance of a given database name can be running at one time. You must rename the SID and database to mount two replicas of the same Oracle database simultaneously. If you rename the SID but do not rename the database, one database instance can be started and the other mounted, however, the second database cannot be recovered or restarted.
UDB	You can create an additional instance (manually) and use that additional instance to access two replicas of the same database simultaneously.
Microsoft SQL Server	You can create additional instances and use those instances to access two replicas of the same database simultaneously. If you want to recover two replicas of the same database in one SQL Server instance, you must rename the database.
Microsoft Exchange	The only reasons you should mount an Exchange replica to an alternate location is for backup or consistency checking purposes or to recover individual mailboxes using a recovery storage group or third-party tool. You should not try to recover an Exchange replica in an alternate location.
SharePoint	You can create an additional SQL Server instance and use that additional instance to access two replicas of the same component database simultaneously.

Table 11 Rules for mounting multiple replicas simultaneously (page 2 of 2)

Application/file system	Applicable rules
UNIX file systems	<p>There are no special rules associated with mounting straight file-system data to an alternate location. Remember that application data or applications themselves stored on those file systems may not work properly if restored to an alternate location unless they are one of the supported applications (Oracle or UDB).</p> <p>File systems made on logical volumes will cause Replication Manager to create unique temporary names for the volume group before it is imported.</p>
File systems	<p>There are no special rules associated with mounting straight file system data to an alternate location. Remember that application data or applications themselves stored on those file systems may not work properly if restored to an alternate location unless they are one of the supported applications (Microsoft Exchange, Microsoft SQL Server, SharePoint, or Oracle).</p>
Network file systems	<p>Replicas of network file systems are created as a Celerra snapshots and each new replica is assigned an alternate name using the following convention <code><servername>-<source file system name>-<unique id></code></p>

Using mount and backup scripts

Replication Manager offers two types of scripts associated with mounting a replica, namely: mount and backup scripts. This section describes these two types of scripts and how to specify them. A main difference between post-mount scripts and backup scripts is that backup scripts can be run on any host, while post-mount scripts will run only on the mount host.

Mount scripts

If you specify a mount script, Replication Manager runs the specified script on the mount host after successfully mounting the replica. To use mount scripts:

1. Right-click an existing Job and select **Properties** from the context menu.
2. Click the **Mount** tab as shown in [Figure 107 on page 187](#).
3. Expand the **General Mount Options** tree and select the **Run a post-mount script** checkbox and enter the path and filename of the script you want to run or click **Browse** to find an appropriate post-mount script.

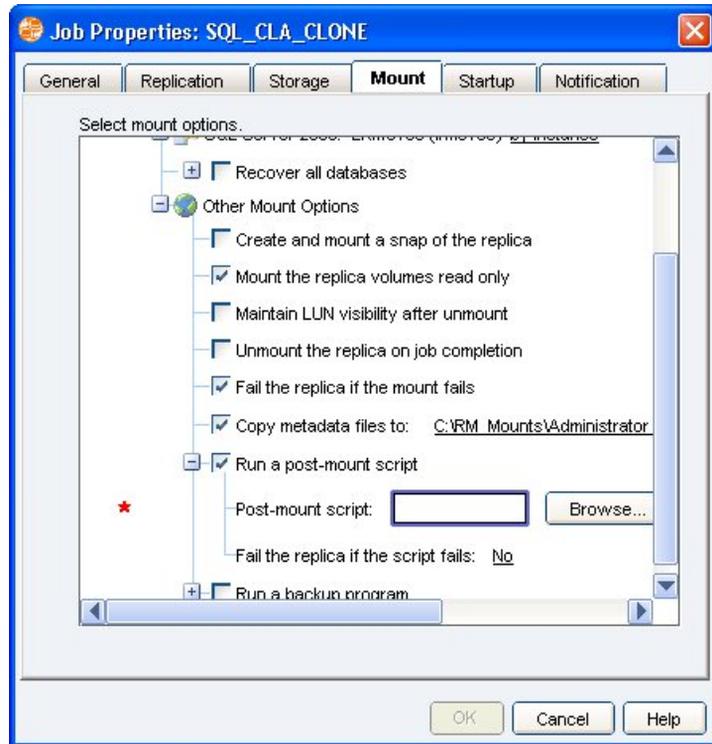


Figure 107 Choosing mount scripts

Backup scripts

If you specify a backup script, Replication Manager runs the specified script on the host of your choice after successfully mounting the replica. To use backup scripts:

1. Right-click an existing Job and select **Properties** from the context menu.
2. Click the **Mount** tab.
3. Expand the **General Mount Option** tree and select **Run a Backup Program** checkbox.
4. Choose a backup host.
5. Enter the path and filename of the backup script or click **Browse** to find an appropriate backup script.

Note: The host you choose must have the Replication Manager Agent software installed. If the production host is a UNIX machine, then the backup host must also be UNIX. Likewise, If the production host is a Windows machine, then the backup host must be running Windows. This parity also applies to application types. For example, if you have an Exchange agent on the production host, then the Exchange agent needs to be installed on the Windows host.

CLARiiON or VNX static mount

Replication Manager enables more reliable mounts of CLARiiON or VNX snaps, clones, and snaps of clones as described in the following sections.

Clones

If mount is part of the job, Replication Manager queries the mount host for visible devices so that it can intelligently select devices for replication. Replication Manager selects visible devices for replication over the devices that require dynamic visibility.

Replication Manager moves the replication clone LUNs to the mount host storage group only if they were not visible to the mount host during mount; that is, it performs a dynamic mount.

The job option **Maintain LUN visibility after unmount** instructs Replication Manager to leave the replication clone LUNs from the mount host storage group visible during and after unmount.

Snaps

Replication Manager creates new snap device(s) during the replication phase. During a mount operation, if the mount host has snap devices that are both deactivated and visible, then:

- ◆ Those devices are used to activate the snap session.
- ◆ Those devices are renamed.
- ◆ The replication snap device will be deleted.

If no such devices are found, the mount operation uses the snap devices created during replication.

In the case of snaps, the job option **Maintain LUN visibility after unmount** instructs Replication Manager to leave the replication snap devices from the mount host storage group visible during unmount.

During expire/purge phase, Replication Manager only destroys the snap session, but does not delete the snap devices if the devices are visible to the host.

Snaps of clones

During a mount operation, if the mount host has snap devices for the clone LUN that are both deactivated and visible, then those devices are used to activate the snap session. If no such devices are found, the mount operation creates new devices.

For snaps of clones, when you enable the job option **Maintain LUN visibility after unmount**, Replication Manager leaves the replication snap devices from the mount host storage group visible during and after unmount, but removes the snap session.

Mounting replicas to an alternate MSCS cluster

Replication Manager running in an MSCS (Windows 2003), or Fail Over clustering (Windows 2008 and Windows 2012) environments (either 32-bit or 64-bit systems) supports the mounting of replicas created in one cluster to a single node of an alternate cluster as a non-clustered resource.

It is not possible to cross-mount a Windows 2003 replica to a Windows 2008 mount host (or vice versa). Attempting to perform this would cause your job to fail.

Replicas mounted using alternate cluster mount functionality are primarily used as the source for backups. The goal is to facilitate the use of passive nodes on existing alternate cluster to perform tasks such as offline backup to maximize resource return on investment for those passive cluster nodes.

Implementing mount to an alternate cluster functionality requires some special configuration steps. The steps to configure your environment is dependent upon what type of storage arrays you are using.

Mounting to an alternate MSCS cluster is supported in environments using Symmetrix or CLARiiON or VNX storage arrays. The *Replication Manager Administrator's Guide* describes the required storage configuration steps.

Unmounting from an alternate cluster after failover

If you mount a replica to a virtual node of an alternate cluster and make the replica a clustered resource, and that resource fails over to another physical node, you will not be allowed to unmount that replica until you perform a manual failback to the physical node that was active when the original mount occurred. This is true because Replication Manager interacts with the underlying devices on the storage array and those devices are different after a failover has occurred.

Mounting replicas to a production MSCS cluster

Replication Manager supports the mount of file system or SQL databases (2008 and above) to a host in a production MSCS cluster environment. The mount can be to an active or passive node of the production cluster.

You can mount a replica on a host (active or passive) belonging to the production MSCS cluster.

Note: The production mount is allowed only for the GPT type devices.

- ◆ To mount SQL databases with the *Recover Databases* option, the production cluster should have another SQL instance running, other than production SQL instance.
- ◆ Restore of a mounted replica is not allowed when the replica is mounted to any of the production cluster nodes.
- ◆ Configuring VNX/CLARiiON/SYMMETRIX steps to alternate MSCS cluster for mount is applicable to production cluster mount as well.
- ◆ The mount host must have a separate storage group other than the cluster storage group for VNX/CLARiiON.
- ◆ In case of Symmetrix devices, the static visibility of devices should be done for mount host alone.
- ◆ Registering cluster disks as target devices under storage services is not recommended. Instead, it is recommended to create a Replication Manager storage pool where the devices of choice are added. This will ensure that the clustered disks are not selected as target devices while doing the replication.
- ◆ If the production host is part of a cluster, and a host from the production cluster is selected for mount, then *Original path* mount is disabled. You must select an alternate path for the mount.
- ◆ It is not possible to cross-mount a Windows 2003 replica to a Windows 2008 or Windows 2012 mount host (or vice versa). Attempting to perform this would cause your job to fail.

Mount and restore on UNIX Clusters

Replication Manager supports mounting of file system and Oracle replicas to cluster nodes running supported UNIX cluster software. Replication Manager takes care of the post-mount steps of identifying the physical volumes on each cluster node, determining valid IDs for the volumes, and propagating them to all the nodes in the cluster.

This feature is supported for HP Serviceguard and IBM HACMP.

This section provides an overview and describes general procedures and requirements. [Appendix A, “Oracle Procedures,”](#) has detailed information for mounting and restoring from Oracle replicas on a UNIX cluster.

Overview

This section provides an overview of the UNIX cluster mount feature.

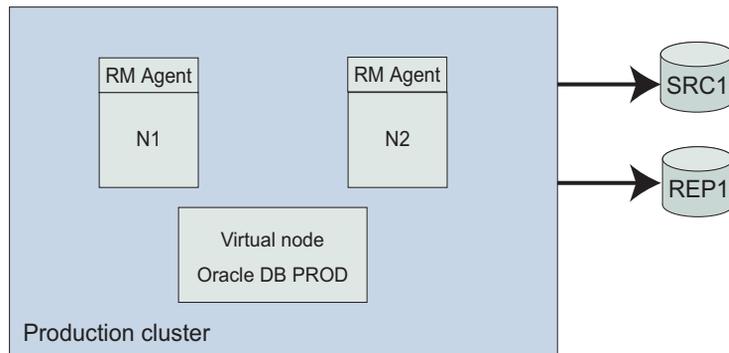


Figure 108 Example of UNIX cluster

Figure 108 on page 193 illustrates a two-node cluster on which Oracle database PROD is installed. PROD runs on the virtual node. Replication Manager agent runs on both physical nodes. SRC1 is the production database device. It must be visible to both nodes for failover to succeed. REP1 is the replica of SRC1.

Node N1 is active and is the database node. When the Oracle replica was mounted to the passive node N2, the mount option **Clustered volume group import** was enabled.

When this mount option is enabled, Replication Manager performs these actions:

1. Makes REP1 device visible to N2.
2. Imports the volume group on N2 and a unique volume group name is generated.

At this point, for HACMP, Replication Manager:

- Runs the **rm_hacmp_pvidupdt.pl** script on N2 to identify physical volumes on each cluster node
- Discovers all the nodes of the cluster
- Updates the ODM definitions with the new physical volume ID
- Identifies available major numbers on all nodes of the cluster and selects a unique major number across all nodes
- Imports the volume group on all nodes of the cluster

For HP Serviceguard, Replication Manager:

- Runs **rm_serviceguard_vgidupdt.pl** to create a map file for the imported volume group
- Identifies all the nodes of the cluster
- Copies them to all other nodes
- Creates a volume group-specific character device on all nodes
- Imports the volume group using the map file
- Activates the volume group, depending on the activation option on the source node

3. Creates mount points.
4. Mounts the file system on N2.

After mount, the administrator needs to create a cluster resource group (using cluster management software) containing the volume groups, file systems, database, and virtual IP that was used to mount the database.

Note: The **Clustered volume group import** option applies to mount hosts running at least Replication Manager version 5.1; otherwise this setting is ignored.

Prerequisites

The following prerequisites must be present on the production and mount hosts for UNIX cluster mount support.

Required by all supported UNIX clusters:

- ◆ rsh enabled between cluster nodes.
- ◆ Cluster nodes must have the Perl interpreter version 5.8.0 or higher, located in `/usr/bin/perl`.
- ◆ For a CLARiiON or VNX array, cluster nodes must be in the same storage group.
- ◆ For a Symmetrix array, verify that all cluster nodes can see the replication devices.

Procedures for mounting a replica (UNIX cluster)

To mount a replica to an HP Serviceguard or IBM HACMP node, run the Replication Manager mount wizard. This performs a mount, imports volume groups on all nodes, and starts Oracle or file systems on one node.

In case of mount failure

Failure to mount the replica to the UNIX cluster could occur for one of these reasons:

- ◆ A cluster node is down.
- ◆ Devices are not visible.
- ◆ An OS command to import a volume group fails.

If a failure occurs:

1. Resolve the issue that caused the failure.
2. Run the Replication Manager script specified in the mount progress details.

Mounting a replica multiple times

If you need to mount the same replica multiple times to the same cluster, note that Replication Manager does not retain the same volume group names for subsequent mounts. You must either create or modify the resource group for the associated replica, or you can recreate the mount points by refreshing the mount node.

Post-mount steps

After the Replication Manager mount wizard successfully completes, do the following if you want to place the mounted replica under cluster control:

1. Create a resource group using the appropriate cluster software with objects specified in replica progress log:
 - Volume groups
 - File systems
 - Database
 - Virtual IP that was used to mount the database (if necessary)
2. Create mount points on all nodes (other than mounted node).
3. Shut down mounted database.
4. Unmount corresponding file systems.
5. Deactivate volume groups on all nodes:
 - Serviceguard: `vgchange -a n volumegroup`
 - HACMP: `varyoffvg volumegroup`
6. Bring the resource group online using cluster software.

rm_hacmp_pvidupdt.pl command

This utility is used to synchronize the PVID (physical volume ID) and major device across multiple nodes in an HACMP cluster. It is run by Replication Manager as part of mount and restore operations, but it can also be run at the command line by the administrator if a node is down during mount or restore, or if you changed something on the production host and want to propagate the changes.

Location The utility is located at `/opt/emc/rm/client/bin/rm_hacmp_pvidupdt.pl` on AIX systems running a Replication Manager agent.

Syntax The command syntax is:

```
rm_hacmp_pvidupdt.pl volume_group_name [-log]
```

Prerequisites Before running this command, verify that:

- ◆ The volume group is in the varied-on state.
- ◆ (If in a cluster resource group) the resource group is not currently online. (Script must take the volume group offline and re-import.)
- ◆ File systems on the volume group are not mounted.

Description The command contacts all nodes in the cluster and updates the PVID known to the node to match that on the devices in the volume group. The major device number of the volume group on the current system will be propagated to all the other nodes, if it is available on all nodes. If the major device number of the volume group is not available, the lowest available major number will be used.

Options *volume_group_name* specifies the name of the volume group to be propagated.

`-log` causes all diagnostic output to be directed to stdout rather than to `/tmp/vgclus_pid.log`.

rm_serviceguard_vgidupdt.pl command syntax

This utility is used to initiate a mount or restore of a volume group on an HP Serviceguard node. It is run by Replication Manager during mount and restore, but it can also be run at the command line prompt by the administrator if a node is down during mount or restore, or if you changed something on the production host and want to propagate the change.

Location The utility is located at `/opt/emc/rm/client/bin/rm_serviceguard_vgidupdt.pl` on HP-UX systems running a Replication Manager agent.

Syntax The command syntax is:

```
rm_serviceguard_vgidupdt.pl {-mount | -restore | -unmount}
vgname
[-minor minor_number]
[-mapfile mapfile_name]
[-importoptions "vgimport options"]
[-changeoptions "vgchange options"]
[-activationoptions "vgchange -a options"]
[-log]
```

Options The following options are supported:

-restore

Initiates a restore.

-mount

Initiates a mount.

-unmount

Initiates an unmount.

Options for -restore and -mount*vgname*

Volume group name to operate on, in `/dev/vgname` format.

-minor *minor_number*

Specifies the minor number of the group node.

-mapfile *mapfile_name*

Specifies the volume-group specific mapfile name to be generated.

-changeoptions "*vgchange options*"

Specifies options (enclosed in quote marks ("")) to be passed to the HP-UX command `vgchange`.

-activationoptions "*vgchange -a options*"

Specifies options (enclosed in quote marks ("")) to be passed to the HP-UX command `vgchange -a`.

-log

Sends logs to stdout.

Options for -unmount*vgname*

Volume group name to operate on in `/dev/vgname` format.

-log

Sends logs to stdout.

Examples

The following example creates the volume group specific nodes, and updates the volume group ID and volume group minor number on the cluster nodes on mount:

```
rm_serviceguard_vgidupdt.pl -mount -log /dev/oravg1
```

The following example cleans up the volume group on the cluster nodes on unmount:

```
rm_serviceguard_vgidupdt.pl -unmount -log /dev/oravg1
```

The following example update the VGID and minor number of the volume group to other cluster nodes:

```
rm_serviceguard_vgidupdt.pl -restore -log /dev/oravg1
```

Considerations for UNIX cluster restore

Disable cluster resource groups

Replication Manager does not interact with the state of the cluster resource groups when performing a restore. It is essential that the cluster resource groups for the items being restored be disabled so that the cluster software does not attempt a failover restart of the item being restored when the instance on the active node is shut down.

HP Serviceguard troubleshooting tip

If Replication Manager is used to replicate a volume group which is in an "available" state when running `vgdisplay`, you may need to run the VGID propagation script (`rm_serviceguard_vgidupdt.pl`) manually after the restore completes.

Running the script will ensure that all the cluster nodes know about the restored volume group. This is necessary because with this volume group state, Replication Manager cannot determine whether the volume group is part of the cluster resource group or not.

To run the script manually use the command:

```
./rm_serviceguard_vgidupdt.pl -log -restore vg_name
```

If the replica is mounted at least once using the **Cluster Import** option, Replication Manager will know that it is part of the cluster resource group and will run the propagation script automatically on restore.

An alternative is to set the environment variable `ERM_HP_CLUSTERMOUNT` in the Replication Manager console under Host properties, Advanced tab. When this variable is set, Replication Manager will automatically update the cluster nodes, even if the source volume group is set to "available."

VMware mount and restore

This section describes special considerations for mount and restore of VMware replicas.

General mount considerations for VMware mounts

Note the following before mounting a VMware replica:

- ◆ **When mounting a VMware crash-consistent VMFS or NFS datastore replica** — Type the name of an ESX Server as the **Mount host**. Also, the mount options panel includes a **Mount using proxy host** field as shown in [Figure 109 on page 201](#). In this field, choose a physical or virtual host that has a Replication Manager agent installed and has access to the VirtualCenter software. Replication Manager uses this proxy host, which can either be a Windows or a Linux machine, to execute the mount of the crash-consistent VMFS replica to the ESX Server selected as the mount host.

You can enable resignaturing, which allows VMware to write a new signature to the LUNs that are being mounted. If you are using ESX Server 4.0 and above, you can set the VMFS resignaturing switch through the Replication Manager console, as shown in [Figure 109 on page 201](#). For ESX Server 3.5 or earlier, resignaturing can be enabled through LVM settings.

- ◆ VMware NFS replicas can also be mounted to a Linux host as an NFS mountpoint. In this case, the proxy host in the mount panel is ignored as it is not needed for the NFS mount.

Note: EMC recommends that this proxy host be dedicated for VMFS mount operations and not be used to perform other Replication Manager tasks.

- ◆ **When mounting a VMware virtual disk replica** — Choose a virtual machine as the **Mount host**. No proxy host is needed for virtual disk replicas. Production mount of virtual disk replicas is not supported. Before you can use a virtual machine as a mount host for CLARiiON or VNX Snaps or clones, there are manual steps that must be performed prior to Replication Manager running any jobs. "[Mount considerations when using CLARiiON or VNX snaps or clones](#)" on [page 202](#) provides more information.

- ◆ **When mounting a VMware RDM disk replica** — Choose either a virtual or a physical machine as a mount host. No proxy host is needed for RDM replicas.

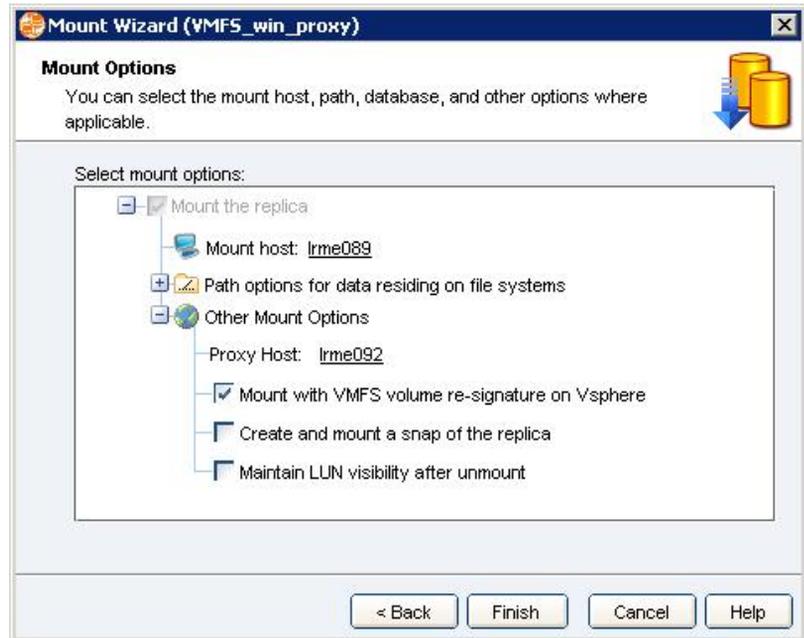


Figure 109 VMware VMFS mount options in the mount wizard

Licensing considerations when mounting VMware replicas

When you create and mount VMware replicas, the licensing requirements differ depending upon the type of datastore you are using. The guidelines are as follows:

Table 12 License requirements for VMware replicas

Datastore replica type	To create replicas you need ...	To mount replicas you need ...
VMFS datastore	RM-VM-Proxy license only	RM-VM-Proxy license only ^a
NFS datastore (mounted to an ESX Server)	RM-VM-Proxy license only	RM-VM-Proxy license only ^a
NFS datastore (mounted to a physical Linux mount host)	RM-VM-Proxy license only	RM Agent license (for physical mount host).

a. There is no Replication Manager Agent running on the ESX Server so no Agent license is required. If you pick a separate proxy host for mount (different than the one you used for replication) then you need another RM-VM-Proxy license (This is the use-case which customer uses proxy A to take replica but uses proxy B to mount it).

Mount considerations when using CLARiiON or VNX snaps or clones

If the target storage will be made visible to the mount host as RDMs then the storage must be pre-exposed to the mount host.

To pre-expose storage to a virtual machine so that it may be used as a mount host for CLARiiON or VNX snaps or clones, follow these general steps:

1. Use Navisphere Manager or EMC Unisphere to create a snap or clone of each source LUN.
2. In Navisphere or Unisphere, add the LUNs to the storage group with the storage for the ESX Server of the mount host's virtual machine.
3. In VMware vSphere or VMware Infrastructure client rescan the HBAs of the ESX Server.
4. Edit the virtual machine settings to add each of those LUNs as a Raw Device Mapped (RDM) disk.

5. In Windows Disk Management screen on the virtual machine, rescan the disks to expose them to the virtual machine.
6. In the mount options of the job, be sure to select **Maintain LUN visibility after unmount**.

Deactivating a CLARiiON or VNX Snapshot session on unmount

To deactivate a CLARiiON or VNX Snapshot session upon unmount, set the environment variable `EMC_ERM_GOLDEN_SNAPSHOTS` on the mount host. This environment variable is intended for use on a Windows VMware system for static mounts of RDM devices. When this environment variable is set, the behavior of a snapshot unmount is changed to deactivate the snapshot session upon unmount. This allows the snapshot session to be reused for mounts of other replicas, but also means that changes made to the snapshot replica are lost upon unmount. In addition, the expiration of a replica deletes the replica snapshot if it is not in a storage group, even if **Maintain visibility on unmount** is enabled.

In addition to setting the `EMC_ERM_GOLDEN_SNAPSHOTS` variable on the mount host, set the environment variable `EMC_ERM_USE_GOLDEN_SNAPSHOTS` on the production host. Setting the variable allows Replication Manager to create only a snapshot session, avoiding the creation of a snap device. This improves the performance of the snapview snapshot during mount and deletion operations. `EMC_ERM_USE_GOLDEN_SNAPSHOTS` is intended for use in the same environment as `EMC_ERM_GOLDEN_SNAPSHOTS`, but on the production host (unlike `EMC_ERM_GOLDEN_SNAPSHOTS` which is set on mount hosts).

Mount considerations on Celerra or VNX NFS

Note the following considerations for mounting replicas in a Celerra environment:

- ◆ If your replica is stored on a virtual disk on Celerra iSCSI, the ESX Server that is the target of the mounted replica must have the IP address and port of the remote Data Mover stored in its ESX iSCSI dynamic discovery panel shown in [Figure 110 on page 204](#). Make sure the appropriate ports are open to communicate between the mount host (ESX Server) and the Celerra target or mount operations will fail.
- ◆ If Replication Manager was upgraded from version 5.2.2 and the Celerra NFS environment variable `IR_ALT_HOST_IP_hostname` was used, EMC recommends that you unset the variable and

instead use the **NFS network host interface** option to specify the alternate IP address or hostname. The option is located under Job properties, Mount tab.

Mount considerations on Symmetrix

In Symmetrix environments, virtual disk replica storage must be made visible to the ESX Servers that host the mount VM prior to running the mount job from Replication Manager.

Note: Mount of a virtual disk replica to a production virtual machine is not supported.

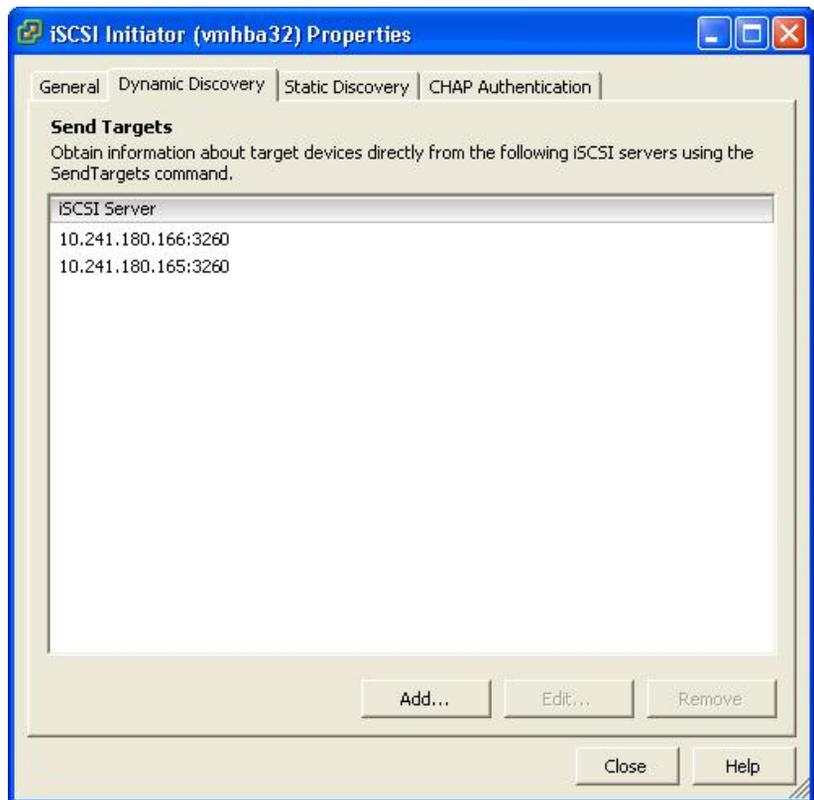


Figure 110 ESX Server Dynamic Discovery panel

Deploying virtual machines from a mounted VMFS

When you use Replication Manager to mount a crash-consistent VMFS replica, there are additional steps that you must perform to deploy the virtual machines once the Replication Manager mount is complete:

1. Add all the virtual machines that you want to deploy to the VirtualCenter inventory panel.

Note: Deploy the virtual machines from the VMFS replica to the ESX inventory only using VirtualCenter. Failure to use VirtualCenter can result in stale entries in the ESX upon unmount.

2. Power on the virtual machines in VirtualCenter. Your virtual machines will now be available for whatever processing you wish to perform on the virtual machines.

Note: If the source virtual machine powered on while you are trying to deploy and power on the replicated virtual machine from a mounted VMFS replica, the power on may fail if the source virtual machine has disks visible to it other than system disk that are share visibility with the source virtual machine.

In this case, remove the non-system disks from the replicated virtual machine before deploying and powering on the replicated virtual machine.

Before unmounting a deployed VMFS

Once you have performed the steps above to deploy virtual machines in a mounted crash-consistent VMFS replica, you must complete the following steps prior to unmounting that replica:

1. Power off the virtual machines in VirtualCenter.
2. Remove the virtual machines from the VirtualCenter inventory.

Restoring VMware replicas

Replication Manager restores fail with errors in the log if you attempt to perform them while the production data is being moved using VMware VMotion technology. If this occurs, retry the restore procedure once the VMotion operation is complete.

Restoring a VMFS replica

Before restoring a crash-consistent VMFS replica you should perform the following steps:

1. Power off the virtual machines that:
 - Are hosted within the VMFS
or
 - Own any virtual disks within the VMFS
2. Remove those virtual machines from the VirtualCenter inventory.
3. Restore the replica from Replication Manager.
4. Once the restore is complete, add the virtual machines into the VirtualCenter inventory.
5. Manually power on each virtual machine.

Note: For information about registering Replication Manager proxy hosts or creating application sets for VMware, read “[Application sets and VMware](#)” on [page 114](#). For more information on VMware support and setup for VMware environments, read the chapter on “VMware Setup” in the *EMC Replication Manager Administrator’s Guide*.

About the restored VMFS environment

The results of a VMFS restore depends upon the state of each virtual machine at the time the replica was taken. When Replication Manager takes a replica of an VMFS it first creates VMware snapshots of each virtual machine that remains powered on within the source VMFS at the time of the replication. Refer to [Figure 111 on page 207](#).

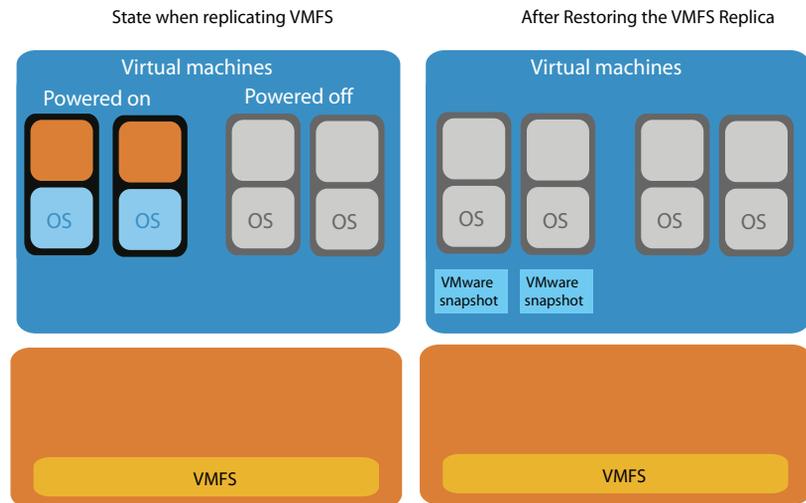


Figure 111 Virtual machines left powered on during VMFS replication generate a VMware snapshot

The VMware snapshot flushes the operating system information from memory into disk before the replication occurs to ensure that no operating system information is lost. These VMware snapshots are stored within the replicated VMFS.

After a restore of a VMFS in which virtual machines were powered on during the replication, you may do any of the following with the existing VMware snapshot(s):

- ◆ Delete the VMware snapshot(s).
- ◆ Revert to the VMware snapshot(s) to obtain an operating system consistent replica.

Continue operation with the VMware snapshot(s) in place. If you choose to continue operation the VMware snapshots continue to grow in order to track and allow you to revert to the specific point in time that the replica was taken.

Restoring a virtual disk replica

Restoring a virtual disk replica is no different than restores of disks in a physical environment. Steps like those outlined above for VMFS replicas are not necessary. Remember the following guidelines when restoring a virtual disk replica:

- ◆ When using Replication Manager, EMC recommends that you use independent virtual disks on each virtual machine instead of sharing disks across virtual machines.
- ◆ If the virtual machines that you replicate and plan to restore have shared virtual disks the virtual machines may fail to power on after a restore.

If this happens remove the shared virtual disks from the virtual machine and add them back into the virtual machine. This operation makes it possible to power on the shared virtual machines that were restored.

File level restore with VMFS and NFS Datastores

Files/Directories stored on virtual disks on a virtual machine can be restored by specifying the location for mounting virtual disks that are stored in VMFS and NFS datastores. The virtual disks will be mounted to a virtual machine which is specified by the user as the proxy host. Virtual disks stored inside the VMFS/NFS datastores which are replicated can be used for file level restore. The Replication Manager administrator can perform a file or directory level recovery by manually copying the files or folders from the location where the disk is mounted, to any other location.

File level restore works in the following environment:

- ◆ The Replication Manager proxy has to be a Windows 2003, Windows 2008, or Windows 2012 virtual machine.
- ◆ The proxy should be configured with the VC credentials, similar to a RM proxy machine.
- ◆ Same proxy can be used for VMFS/NFS datastores file level restore.
- ◆ Only files/folders of the Windows virtual machine residing inside the VMware datastore (both VMFS and NFS) can be recovered using this proxy.
- ◆ File level recovery of Linux VM not supported.
- ◆ Datastores supported are VMFS and NFS.

- ◆ Virtual disks belonging to Windows 2003 machines should be mounted on to Windows 2003 RM proxy and virtual disks belonging to Windows 2008 or Windows 2012 machines should be mounted on to Windows 2008 or Windows 2012 RM proxy, respectively.
- ◆ Make sure there are enough free unique SCSI target IDs available across the controllers, while mounting virtual disks to the proxy.
- ◆ For VMWare VMFS/NFS datastores, Replication Manager requires the virtual machine names and virtual disk paths to be ASCII characters only. File level recovery will not be enabled on virtual machines or virtual disk path with non-ASCII characters.

Enable the file level restore by selecting the datastores and virtual disks in the Mount Wizard. Select one or more virtual disks from one or more virtual machines. [Figure 112 on page 209](#) displays the selection to be made.

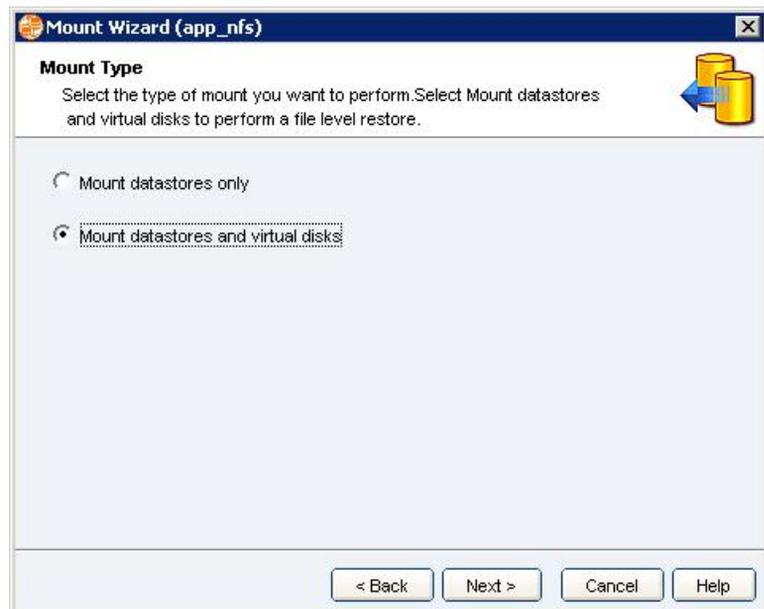


Figure 112 File level restore selection panel

After making the above choice, select the proxy host and the location on the proxy host where the virtual disks has to be mounted.

File level restore is a two phase process - in the first phase, the snapshot datastores are mounted to the ESX server on which the RM

proxy client is running. In the second phase, the selected virtual disks are mounted from the snapshot datastores to the Replication Manager proxy client.

Virtual disks can be unmounted by the administrator.

File level restore is not possible on dynamic disks. Also, file level restore cannot be enabled on virtual disks with multiple partitions. When a disk contains multiple partitions, only the first partition is mounted.

Upgrade considerations for file level restore

File level restore can be performed only on replicas that have been created in version 5.3.1 and above.

In case of an upgrade to RM 5.3.1 and above, new replicas from existing jobs can be used to perform file level restore.

Restoring a single virtual machine

Replication Manager gives you the ability to restore a single virtual machine from a datastore replica. When restoring a datastore replica, you have three options: full datastore restore, virtual machine, or file using mount wizard.

Prerequisites

In order to perform the restore:

- ◆ The full image of the virtual machine should be contained in the replica.
- ◆ The virtual machine should be powered off or powered off and removed from the inventory for the restore to complete successfully.
- ◆ There should be no RDM devices attached to the virtual machine that are part of the replicated datastore.
- ◆ Virtual machine folder/directory inside the datastore containing the virtual machine configuration, snapshot, and disk files should not contain any other virtual machine configuration and disk files; otherwise, they will be overwritten after the restore.
- ◆ For virtual machine restore, Replication Manager requires the virtual machine names and virtual disk paths to be ASCII characters only. Virtual machine restore is not supported on virtual machines with non-ASCII names.
- ◆ In case the virtual machine is deleted from the inventory, be sure to enter the ESX host where the source datastore is mounted in the restore wizard as Replication manager also registers the virtual machine to this ESX after restore is completed.

Support Restoring single virtual machines is supported on:

- ◆ The Replication Manager proxy
- ◆ Windows proxy host only
- ◆ CLARiON and VNX storage platforms
- ◆ ESX Servers and virtual center versions 4.1 and above.

Upgrade issues Old replicas cannot be used to perform single virtual machine restore operations. After an upgrade, you must create a new replica to take advantage of the single virtual machine restore option.

Restore procedure To restore a virtual machine:

1. Select the application set, right-click, and select **Restore a Replica**.
2. In the **Replica to Restore** dialog box, select the replica you wish to restore and click **Next**. The Restore Type dialog displays, [Figure 113 on page 211](#).

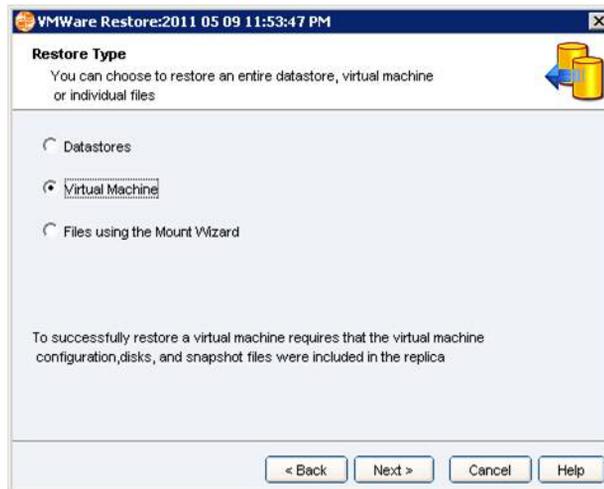


Figure 113 Restore Type

3. From the Restore Type dialog box, you have three options. For this example, select **Virtual Machine**. The Select Virtual Machine dialog displays, [Figure 114 on page 212](#).

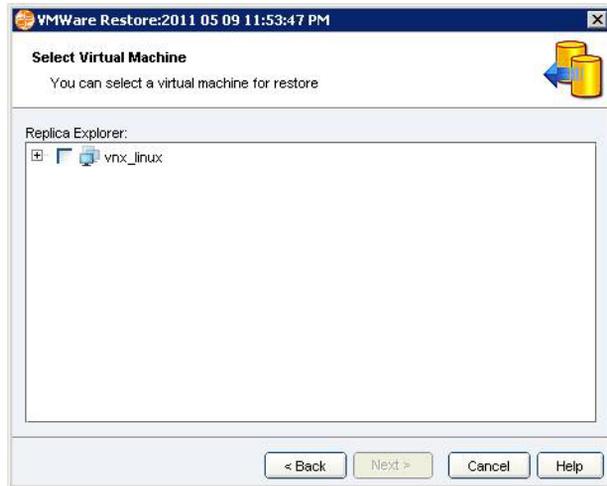


Figure 114 Select Virtual Machine

4. Select one or more virtual machines. All of the virtual machines in the datastore within the replica are displayed.
5. Click **Next**. If the replica is mounted to an ESX server, the Select ESX Server dialog displays, [Figure 115 on page 212](#).

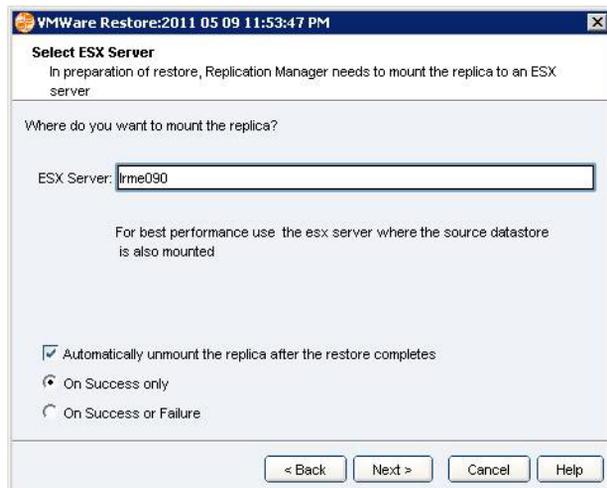


Figure 115 Select ESX Server

- In the **ESX Server** text box, enter the name of the ESX server where you want to mount the replica.

Note: Better performance will be achieved if you mount the replica onto the same ESX server as the source datasource.

- You have the option of keeping the replica mounted if the restore fails. In this case, check **Automatically unmount the replica after the restore completes** and **On Success** only. Otherwise, check **On success or failure**.
- Click **Next**. The Select Restore Proxy dialog displays, [Figure 116 on page 213](#).

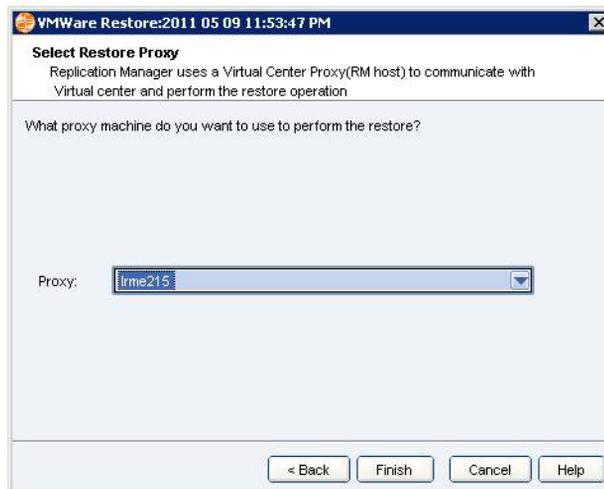


Figure 116 Select Restore Proxy

- Select the proxy host that will communicate with the Virtual Center and perform the restore.
- Click **Finish**.
- Click **Yes** in the **Confirm** dialog which displays the virtual machines selected for restore.

The datastore replica is mounted to ESX and the virtual machine is restored.

The datastore replica is mounted to the ESX server and the virtual machine restore is performed by copying virtual machines configuration and disk files from destination to source datastores.

Replication Manager also restores virtual machine snapshots if they were created during replication, see [Figure 117 on page 214](#). You can revert back to this snapshot after Replication Manager virtual machine restore is finished by logging into VMware infrastructure client and browsing the virtual machine snapshot manager. After reverting back the snapshot or after successful restore, you should delete this snapshot.

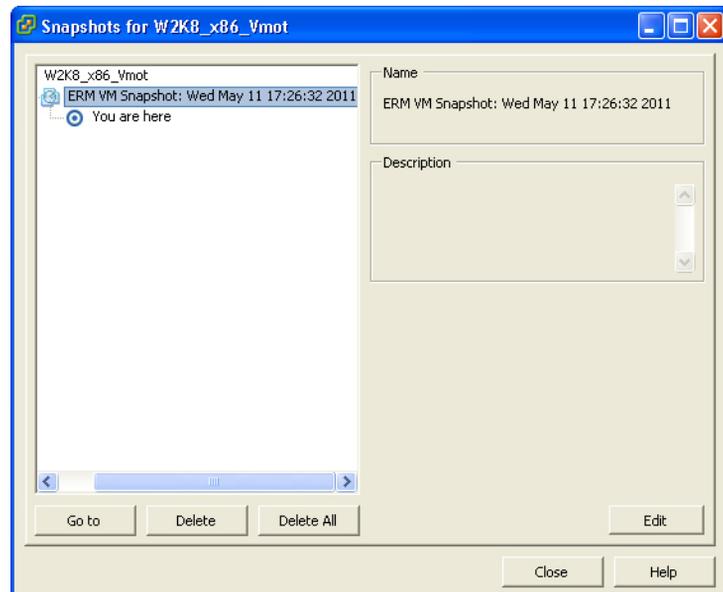


Figure 117 Restored virtual machine snapshot

RecoverPoint mount and restore

This section describes special considerations for mount and restore of RecoverPoint replicas.

General mount considerations for RecoverPoint

Note the following before mounting a RecoverPoint replica:

- ◆ You can mount only one RecoverPoint replica at a time, per application set, per site. In other words, for an application set, only one RecoverPoint replica can be mounted on the local site (CDP), and one on the remote site (CRR).
- ◆ After an application set modification, especially after adding or removing volumes from the application set, run a RecoverPoint job before attempting a crash-consistent (any-point-in-time) mount or restore. This gives Replication Manager the information it needs to mount to the correct drive letters, attach to the database, and so on.
- ◆ Changes made to a RecoverPoint replica while it is mounted are not persistent; they are lost when you unmount the replica. This behavior is different from mounted replicas using some other replication technologies in Replication Manager. (Changes made to a mounted replica can be restored, however.)
- ◆ You can restore from a read-write mounted RecoverPoint replica (unlike other replication technologies); changes made to a mounted RecoverPoint replica are propagated to the production volumes during a restore. When the replica is unmounted, changes made to the mounted replica are lost.
- ◆ Target volumes must be visible to the mount host.
- ◆ If you mount a point in time that is outside of the protection window, Replication Manager will mount the nearest available time, but the console will display the time you requested. The replica history log will display the actual time mounted.
- ◆ In the event that your production data is unavailable, due to failure of the local RPA or array for example, Replication Manager permits you to mount the remote copy. The consistency group copies you wish to mount must be available.

RecoverPoint replicas in mount and restore wizards

When you mount or restore a RecoverPoint replica, you have the option to select from a list of existing replicas (the same as for other replication technologies), or to create a replica from a point in time and mount or restore that replica.

If you select **Select a point in time to mount (or restore from)**, choose a time to create the replica from, and name the replica.

For CLR-configured jobs, you additionally specify which side to mount or restore from, CDP or CRR.

Figure 118 on page 216 shows a sample panel in the Mount Wizard.

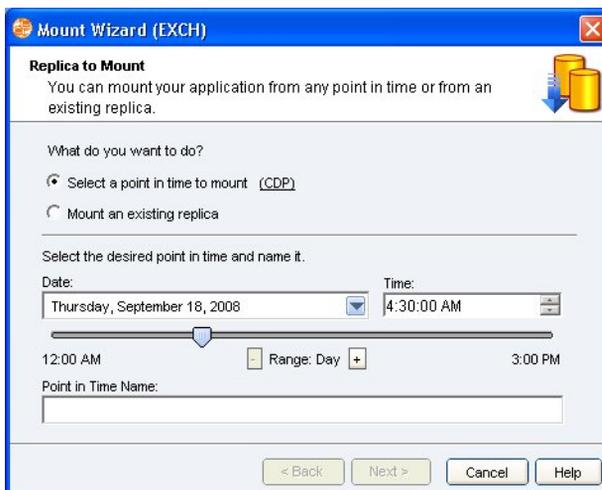


Figure 118 Mount Wizard, creating replica from point in time

RecoverPoint image access

In RecoverPoint, *image access* refers to providing a target-side host application the opportunity to write data to the target-side replication volumes, while still keeping track of source changes.

Image access can be physical (also known as logged), which provides access to the actual physical volumes, or virtual, with rapid access to a virtual image of the same volumes.

Note: If the underlying storage is VPLEX, only physical is supported.

When you mount a RecoverPoint replica in Replication Manager, the following mount options related to image access are available:

- ◆ Physical target access
Use this mount option if the integrity check entails the scanning of large areas of the replicated volumes. This is the only option available when you mount to the production host.
- ◆ Virtual target access
Provides nearly instant access to the image; it is not intended for heavy processing.
- ◆ Virtual target with roll
Provides nearly instant access to the image, but also updates the replicated volume in the background. When the replicated volumes are at the requested point in time, the RPA transparently switches to direct replica volume access, allowing heavy processing.

RecoverPoint replicas under SRM or CE management

Mount and restore

This section describes Replication Manager's interaction with RecoverPoint in a VMware SRM or RecoverPoint/CE environment.

In RecoverPoint 3.2 SP2 you cannot manage bookmarks while a consistency group is under VMware SRM management or RecoverPoint/CE. The default effect in Replication Manager is that you cannot mount or restore from a RecoverPoint replica when a consistency group is under SRM or CE management:

- ◆ For mount, Replication Manager provides an option (**Disable SRM management or RecoverPoint/CE while mounted**) to take the RecoverPoint consistency group out of SRM or CE management mode at mount time, and return it when the replica is unmounted.
- ◆ For restore, you must manually disable SRM or CE in the RecoverPoint Management Application GUI before the restore. There is no restore option in Replication Manager for this.
- ◆ Note that SRM or CE management operations will fail until Replication Manager or the RecoverPoint administrator turns on SRM or CE management.

In RecoverPoint versions before 3.2 SP2, you can manage bookmarks without Replication Manager's intervention.

Failover behavior

After a CE or SRM failover, Replication Manager continues to run RecoverPoint jobs without modification to the job or application set except in the following cases:

- ◆ Replication Manager jobs that contain a mount may require changes to use a different RecoverPoint copy (CDP to CRR or vice versa) or mount host.
- ◆ RecoverPoint CLR consistency group jobs that fail over may no longer have both local and remote targets. Both local and remote targets are required for a Replication Manager CLR job to create replicas.

The disaster recovery chapter of the *Replication Manager Administrator's Guide* describes in detail the steps related to failover.

Restore considerations for RecoverPoint

Note the following before restoring from a RecoverPoint replica:

- ◆ After an application set modification, especially after adding or removing volumes from the application set, immediately run a RecoverPoint job. This is especially important if you are planning any crash-consistent (any-point-in-time) mounts or restores. Running the job gives Replication Manager the information it needs to mount to the correct drive letters, attach to the database, and so on.
- ◆ If you change the property of a RecoverPoint consistency group whose replicas are managed by Replication Manager, immediately run a RecoverPoint job.
- ◆ Before restoring, verify that:
 - RecoverPoint CDP or CRR is running (not paused, stopped, or disabled) for volumes contained in the application set.
 - No job is running for the application set.
 - Target production volumes are mounted.

**CAUTION**

Verify that the data you are about to restore is what you want to restore. Mount and examine the data before restoring to a production system.

Non-RecoverPoint restores

If you are performing an array-based, non-RecoverPoint restore for an application set that is also configured for RecoverPoint, use the RecoverPoint management application to pause transfer before restoring.

Troubleshooting mount to a Windows 2008 or Windows 2012 host

Mounting a RecoverPoint replica to a Windows Server 2008 SP2 host or Windows 2012 host may fail due to the automount feature of the operating system. By default Windows Server 2008 SP2 and Windows 2012 automatically mounts new volumes and assigns drive letters to them. This can cause the mount to fail, especially in the case of a static mount and when mounting to the original path.

The workaround is to manually disable automount on the mount host. Run `diskpart` at a command prompt and enter **automount disable** at the `DISKPART>` prompt.

Unmounting a replica

Replication Manager allows the user to unmount a mounted replica, removing it from the mount host so that it is no longer visible to that host. To unmount a replica:

1. Expand the application set that contains the replica in the tree panel, a list of replicas appears in the content panel.
2. Right-click the replica you want to unmount and select **Unmount** from the context menu.
3. Replication Manager attempts to unmount the replica.

Note: If something prevents Replication Manager from unmounting the replica, the warning in [Figure 119](#) on page 220 appears.

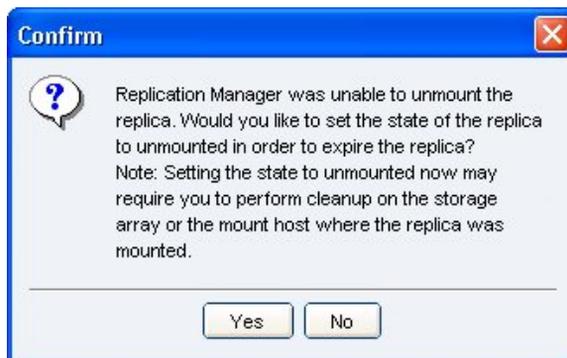


Figure 119 Force unmount warning

- a. Click **Yes** if you want to set the state of the replica to unmounted (even though it might continue to be mounted on the mount host). This requires you to perform additional cleanup on the mount host and the storage array to complete the unmount operation.
- b. Click **No** to leave the replica status as mounted.

Note: Unmount may take a long time if multiple application objects (databases, file systems, etc.) reside on a single LUN. The best practice is to ensure that each application object resides on a dedicated LUN as that configuration will avoid this performance issue.

Unmounting from a UNIX cluster node

To unmount a replica from an HP Serviceguard or IBM HACMP cluster node:

1. Take offline all resources associated with the mount copy.
2. Use the cluster software to shut down the resource group.
3. If mount was done using the virtual IP address, enable the virtual IP manually.
4. Unmount the replica using as described in [“Unmounting a replica” on page 220](#).
5. If the unmount fails (for example, because a cluster node is down or an invalid OS command), resolve the issue causing the failure, then run the script specified in the unmount progress details.

Replica mount performance guidelines

There are many factors that can affect the amount of time it takes to mount replicas on a mount host. These factors include:

- ◆ Mount host physical configuration (memory, CPU, HBA, etc.)
- ◆ Storage array utilization
- ◆ Network utilization
- ◆ Number of storage arrays zoned to mount host

Refer to this section for capacity planning of Replication Manager mount hosts.

Mount locks

When calculating the time it takes to perform the entire mount process, it is important to understand the way Replication Manager utilizes mount locks. Since only one mount operation can happen at one time, Replication Manager uses locks to allow only one mount or unmount to run. When a mount or unmount is started, the job will immediately check for the mount lock; if it is unavailable, it will poll for the mount lock in short intervals.

Because of the short intervals on checking for the lock, it is possible that a mount that was started later can obtain the lock before a mount that is already waiting for the lock. Due to this functionality, it is possible that the time spent waiting for the mount lock can increase the overall time of the mount job by several hours.

This time waiting is not considered part of the actual mount time. The actual mount time is determined by how long it takes to perform. The reason the wait time is not considered is because the wait time has little effect on the overall performance of Replication Manager. It is only during the time when the mount lock is obtained that will affect other jobs.

For example, if you mounted two replicas at the same time and they each take 30 minutes to mount, the first one to obtain the mount lock will complete in 30 minutes, and the second will take 1 hour, since it had to wait for 30 minutes to obtain the lock. Even though the second mount took a total of 1 hour, it is still considered to be only a 30-minute mount.

It is recommended that you configure mount operations so that they are staggered, to give you more control over the order in which replicas are mounted.

Serialization of mounts

The Replication Manager client on the mount host will lock a portion of a replica mount or unmount, which means that if two replicas are being mounted at the same time, there will be a period of time in which only one of the replica mounts will be able to make progress. The active mount acquires a lock, which prevents the other replica mount from proceeding. As a rule, the lock is held for about 90 percent of the total time it takes to mount or unmount a replica, so a mount that takes 5 total minutes would hold the lock for about 4 minutes and 30 seconds.

LUN surfacing and mounting

For CLARiiON, VNX, Celerra iSCSI, VNXe iSCSI, and Symmetrix (on Solaris only), the act of mounting a replica involves making the LUNs for the replica visible to the mount host, rescanning the SCSI bus, and then mounting the LUN to a drive or mount point. On UNIX, a mount could also involve Veritas volume importing and `fsck`. On average, one could assume an average time of 1 minute per LUN to surface the LUN and mount the LUN to a drive, so a six-LUN replica could take up to 6 minutes to mount. On Windows Server 2003, Windows 2008, or Windows 2012, a first mount of a replica uses VSS to import the replica devices.

LUN unmounting and submerging

This becomes an important consideration if the administrator chooses to leave replicas mounted until the next replica. Before a replica can be mounted, the previous replica of the same application set must be unmounted. This involves unmounting the drive or mount point and performing the necessary task to remove the LUN visibility from the host, and performing a SCSI rescan.

Summary

Careful planning should go into determining mount host requirements. EMC recommends careful testing and analysis of a mount host before completing the site implementation. If a mount host begins to exceed the desired backup window, consideration should be given to adding additional mount hosts.

Restoring from a replica

When a database or a file system that you are protecting with Replication Manager gets damaged or destroyed, you can restore the data from a valid replica back to the original production instance. If you are restoring a database, apply the logs to recover the database after the restore.

In Windows 2012 environments, when doing a restore, the data on LUNs is overwritten even if the volume is in use. This differs from other Windows platforms in which Replication Manager displays a warning if the LUN is in use. Since restores will overwrite everything, be sure that there is no other data on that volume and the volume is not in use.

Restore restrictions and limitations

The following list describes considerations that apply when you restore an original production instance. By definition, the restore process overwrites the production devices with the contents of the devices on which the replica resides:

- ◆ Restoring a replica of a source disk that is part of a CCR Exchange Cluster or a Windows Server 2008 or Windows Server 2012 MSCS cluster requires you to first mount the replica to another server using VSS to "import" it so internal flags on the disk can be set properly.
- ◆ At Symmetrix Microcode levels of 5874 or greater, restore of a TimeFinder/Clone that is part of a cascade or multi-hop causes Replication Manager to terminate the second hop (clone of clone) sessions of all other replicas. All original TimeFinder/Clone sessions remain intact.
- ◆ If a specific Symmetrix device (at Symmetrix Microcode levels less than 5874) has TimeFinder/Clone and TimeFinder/Mirror based replicas, restoring from one of the TimeFinder/Clone or TimeFinder/Mirror replicas terminates the TimeFinder/Clone replicas of the same source and requires a full copy on the next replication of the clones.
- ◆ TimeFinder clone replicas created using the No Copy job option can neither be restored nor be used to create a copy of a job.

- ◆ Restore from a TimeFinder clone replica is not supported when TimeFinder NoCopy sessions exist for the source. In order to perform such a restore, terminate all TimeFinder Clone NoCopy sessions first and then perform the restore operation.
- ◆ If a replica of RAID5 or RAID6 BCVs (TimeFinder/Mirrors) exists for a certain Symmetrix source, they will prevent restore of TimeFinder/Clone replicas to that same source.
- ◆ Remote replicas created using TimeFinder/Clone or TimeFinder/Snap of an R2 device cannot be restored by Replication Manager. To restore from these replicas requires manual steps that are beyond the scope of this guide.
- ◆ Restore from a TimeFinder/Mirror replica is not supported when a copy session exists between a TimeFinder/Mirror device and TimeFinder/Snap (VDEV) device. In order to perform such a restore, terminate all copy sessions between the TimeFinder/Mirror and the VDEV.
- ◆ Remote replicas created using SnapView/Clone or SnapView/Snap of a MirrorView/A or MirrorView/S secondary device cannot be restored by Replication Manager. To restore from these replicas requires manual steps that are beyond the scope of this guide.
- ◆ Mounts and restores may fail when the application set contains nested mount points. For instance, if some of the files in the application set are on L:\ and other files are on L:\SG1DBMP (where SG1DBMP is a mount point), mounts and restores of the associated replicas may fail.
- ◆ Open Replicator replicas and their copy jobs cannot be restored.
- ◆ Run a replication on the current production application set corresponding to the replica before running a restore operation as a protective measure. The replica allows you to roll back from the restore should the restore fail or if the successfully restored data is not what you want.
- ◆ If you restore striped data to a re-created stripe set that has different members, the restore will succeed, but the resulting configuration may be different from the original configuration.
- ◆ Do not restore from a replica if the items being restored are in an active resource group; this will cause cluster problems in UNIX clustered environments only.

- ◆ Restore is not supported from TimeFinder copy job replicas, unless the source job is TimeFinder Snap or TimeFinder Clone.
- ◆ Windows-based replicas that are mounted to the production host, mounted read-write, or mounted with the **Create and mount a snap of this replica** option cannot be restored.
- ◆ UNIX-based replicas that are mounted cannot be restored.
- ◆ Restores might fail if the source/production volumes have been shared, either manually by the user or automatically by operating system, such as Windows Server 2003, for administrative purposes. These shares would have to be disabled for a successful restore.

Additionally, administrative shares automatically enabled by the operating system might be re-enabled even after they are disabled upon a reboot. In that case, they would have to be disabled again.

Database restores

To restore from a replica:

1. Expand **Application Sets**.
2. Select the particular application set that contains the replica you want to restore.
3. Right-click the replica you want to restore and select **Restore**.
4. Select the parts of the database or file system that you want to restore.

Note: If you choose to restore a partial replica and the parts you choose share volumes with parts that you did not choose, Replication Manager automatically selects those additional parts of the replica, since the restore granularity is at the volume level.

When restoring databases, all of the devices holding the database are restored.

File system restores

When restoring replicas whose datafiles are built on file systems, Replication Manager automatically performs the following steps:

1. Unmounts the targeted production file systems from the production host.
2. Departs the affected volume groups.
3. Restores the file system.

4. Imports volume groups.
5. Repairs the file system using a utility such as **fsck**.
6. Mounts the recovered file system.

After restoring the replica data to the file systems, the software runs a file system check to check and recover the file systems that were busy when the replica was taken. Then, the software remounts the file systems on the production host.

To restore from a replica, follow the steps found in [“Database restores” on page 226](#).

Note: If you choose to restore a partial replica and the parts you choose share volumes with parts that you did not choose, Replication Manager automatically selects those additional parts of the replica, since the restore granularity is at the volume level.

When restoring file systems, all of the devices containing the file system are restored.



CAUTION

When you perform a restore of a partitioned Celerra or VNX NFS file system, the entire contents of the Celerra or VNX file system are restored, not just the mounted partition on the production host. If there are hosts which mount the other partitions of the Celerra file system, data on those hosts might be affected by the restore.



CAUTION

When you perform a volume-level restore from storage with logical volumes, all of the devices in the volume group get restored. If there is unrelated data in the volume group, other file systems or tablespaces on the production host might be affected by the restore.

For specific information about each wizard screen, click **Help**.

Application restore issues

For information about restore issues for specific applications, refer to the application-specific appendixes at the end of this guide. Each appendix has a section devoted to restore issues.

Restrictions on restoring to synchronized LUNs

It is not possible to successfully restore to a source LUN if that LUN is currently synchronized with another LUN. If you want to restore to a synchronized LUN, you must first split the clone from the source LUN using Navisphere.

Restoring a CLARiiON or VNX protected restore environment

When you restore from a clone replica on a storage array that has Protected Restore enabled, that Replication Manager restore automatically takes advantage of the Protected Restore feature.

Because storage arrays with Protected Restore capability allow writes to the source LUNs immediately, while the reverse synchronization is taking place, the Protected Restore capability (if enabled) also prevents writes to the source LUN from being written to the clones that contain the replica while the reverse synchronization (restore) is happening. Protected Restore maintains the integrity of the replica, but allows, in effect, an instant-restore of the data on the replica.

Although the restored LUNs are available immediately, the array continues to restore from the replica in the background. If you attempt to run another replication or restore of a replica from the same application set while the array is completing the restore, Replication Manager waits for the first restore to complete before starting the next task.

Restrictions on VDEVs

Note the following restrictions on restoring from VDEVs.

- ◆ Restoring from a read-only mounted VDEV is not supported.
- ◆ When you are using VDEVs to store data, remember that there must be an unused VDEV "copy session" available to create these replicas. There is a theoretical maximum of 16 VDEV copy sessions defined for any particular standard device, however the available number of LUNs is reduced by the following factors:
 - BCVs associated with the source hypervolume use a copy session.
 - Other applications using SDDF sessions use up copy sessions.
 - One copy session is reserved to restore from the replica.
 - VDEVs associated with a source volume prevent restore from a BCV-based replica to the same source volume.

If all of the VDEV "copy sessions" are being used, you must first delete one of those replicas before you can create a new VDEV replica from that standard.

Symmetrix TimeFinder restrictions

Be aware of the following when restoring Symmetrix TimeFinder/Clone replicas:

- ◆ VDEVs associated with a source volume prevent restore from a Symmetrix clone replica to the same source volume.
- ◆ Replication Manager cannot perform a TimeFinder/Clone restore to a source device that has an application lock. The application lock must be removed.
- ◆ Because RAID 5 and RAID 6 BCVs associate with source devices, they prevent VDEV (TimeFinder/Snaps) replicas from being restored to the same source, and vice-versa. Note the following when restoring:
 - Before restoring a VDEV or TimeFinder/Clone replica to a source that also has a RAID 5 or RAID 6 BCV, delete Replication Manager RAID 5/RAID 6 BCV replicas for the same source, and manually disassociate the RAID 5/RAID 6 BCVs (using a **symmir** command).
 - Before restoring to RAID 5/RAID 6 BCVs, delete Replication Manager VDEV replicas for the same source.
- ◆ Expiring a TimeFinder/Clone or TimeFinder/Mirror replica, also expires any TimeFinder/Snap copy replicas associated with the expired clone or mirror replica, unless the associated TimeFinder/Snap copy replica is mounted.
- ◆ If a Symmetrix standard device has TimeFinder/Clone replicas and VDEV replicas, restoring a VDEV replica will cause any TimeFinder/Clone session for the standard to be terminated. This will result in a full copy the next time the TimeFinder/Clone target is updated by a replication.

Refer to the *EMC Replication Manager Administrator's Guide* for more information on configuring TimeFinder/Clones.

TimeFinder Clone VP Snap restrictions

Be aware of the following when restoring using the VP Snap job option:

- ◆ The VP Snap job's restore is similar to the Snap persistent restore and results in an additional session charged to the source device. The original CopyOnWrite session is preserved.

- ◆ The VP Snap restore sessions associated with the selected source device must be in the CopyOnWrite state and the other sessions must be in the Copied or Split state.
- ◆ Upon deleting a VP Snap restored replica, both the CopyOnWrite and the restore sessions are terminated.
- ◆ The VP Snap source device can have only one restored session at a time.
- ◆ Clone Split is not supported for a VP Snap restore session. Any newer restore requests on the same source wait for an invalid track to synchronize and terminate the older VP Snap restore session.
- ◆ If a Precopy or Online - Copy job runs with a source having a VP Snap restore, the restore session is terminated and the original CopyOnWrite session is preserved.
- ◆ Precopy or Online - Copy restore fails if the source has a VP Snap CopyOnWrite session.

Celerra restrictions

Be aware of the following when restoring Celerra replicas:

- ◆ Upon successful restore of a Celerra iSCSI SnapSure replica, Replication Manager automatically marks newer snap replicas within that application set as not restorable. In the event of a partial restore, the corresponding portions of the newer snap replicas are marked not restorable. In either case, the newer replicas cannot be mounted at all.
- ◆ This restriction applies to Celerra or VNXe iSCSI only, Celerra NFS does not have the same restriction.
- ◆ If you have SnapSure snaps and Celerra Replicator replications that share the same application set, and you restore a snap replication, any related Celerra Replicator replication sessions will be deleted. At this point no operations (mount, restore, or expire) can be done on the Celerra Replicator replicas until a new Celerra Replicator replica is created.

Shared storage issues

To use restore effectively, you must understand the restrictions restore places on you and the selections that you make when defining your application sets in Replication Manager.

When you restore, all data on each volume that contains any information defined as part of the application set is restored. A

volume is defined as any of the items listed in [Table 13 on page 231](#), depending on your environment.

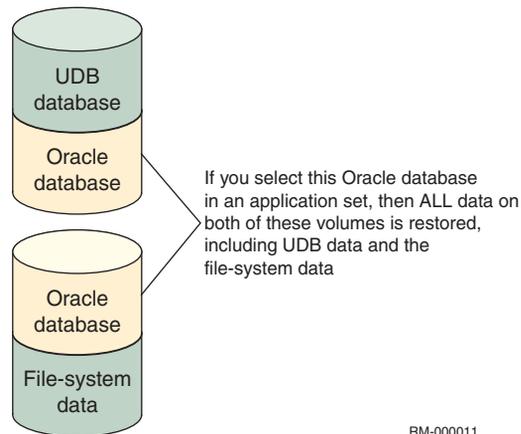
Table 13 Volume types by storage array (page 1 of 2)

Storage array	Volume manager	Restore granularity
Symmetrix (Mirrors)	No	BCVs.
Symmetrix (Mirrors)	Yes	All BCVs in the volume group (as defined by the volume manager).
Symmetrix (Clones)	No	Symmetrix Clones
Symmetrix (Clones)	Yes	All Clones in the volume group (as defined by the volume manager).
Symmetrix (Snaps)	No	Restores snapshots to the original source at device level.
Symmetrix (Snaps)	Yes	Restores snapshot to the original source at volume group level.
CLARiiON (clones)	No	Clone LUNs.
CLARiiON (clones)	Yes	All clone LUNs in the volume group (as defined by the volume manager).
CLARiiON (snapshots)	No	Snapshot sessions (defined in Snap cache).
CLARiiON (snapshots)	Yes	All snapshot sessions (as defined by the volume manager). If consistent split is selected then all LUNs in the replica are restored.
VNX (clones)	No	Clone LUNs.
VNX (clones)	Yes	All clone LUNs in the volume group (as defined by the volume manager).
VNX (snapshots)	No	Snapshot sessions (defined in Snap cache).

Table 13 Volume types by storage array (page 2 of 2)

Storage array	Volume manager	Restore granularity
VNX (snapshots)	Yes	All snapshot sessions (as defined by the volume manager). If consistent split is selected then all LUNs in the replica are restored.
Celerra SnapSure (local snapshots)	No	Celerra snapshots.
Celerra SnapSure (local snapshots)	Yes	All Celerra snapshots (as defined by the volume manager).
Celerra Replicator (remote snapshots)	No	Celerra LUNs.
Celerra Replicator (remote snapshots)	Yes	All Celerra LUNs (as defined by the volume manager).
VNXe SnapSure (local snapshots)	No	VNXe snapshots.
VNXe SnapSure (local snapshots)	Yes	All VNXe snapshots (as defined by the volume manager).
VNXe Replicator (remote snapshots)	No	VNXe LUNs.
VNXe Replicator (remote snapshots)	Yes	All VNXe LUNs (as defined by the volume manager).
RecoverPoint Appliance	Yes	Consistency Group
RecoverPoint Appliance	No	Full replica restore only.

In some circumstances, a Replication Manager restore overwrites some data that you did not intend to restore. These extra data elements are located on storage that shares volumes with the intended restored material. All data that shares these volume(s) will be restored and that restore can affect other applications or data accuracy. [Figure 120 on page 233](#) is an example of this kind of situation.



RM-000011

Figure 120 Shared storage issue

Reducing affected-entities issues

To reduce shared storage issues, consider and carefully plan your disk geometry based on the application sets that you plan to use to create replicas. For example, if you want to be able to easily restore at the tablespace level using Replication Manager, you might decide to set up your database so that each tablespace resides on its own device or volume group (when using Veritas Volume Manager) and that no other data resides on that device or volume group.

Restoring replicas that contain multiple applications

If you have two or more applications that contain related data and you want to be able to use Replication Manager to restore the data from both applications, you can create an application set that contains both applications.

When you create application sets with multiple applications, you can choose to restore all of the data from both applications at the same time. Refer to [“Federated compared with composite application sets” on page 260](#) for more information.

When you choose to restore replicas associated with these application sets, the procedure restores all the data that you elected to replicate from each of the applications in the application set.

Restoring to the R1 device in an SRDF environment

Restoring from the local copy to an R1 device requires additional manual steps (after the Replication Manager restore procedure) to manage the link state of the R1.

To change the link state of one device, change the state of all devices in the RDF group as follows:

1. Put all SRDF devices residing in the same RDF group into a `syndg` group.

To check which devices are present in the same RDF group, run the command `symrdf list`.

2. In the output of the `symrdf list` command, the G column under RDF Typ:G shows the RDF group number for each device.

From this column, determine the device IDs present in the same RDF group.

3. Use the `syndg create` command to put the SRDF devices in one device group.
4. Use the appropriate `syndg` command to set the link state (consistent to split or split to consistent) as necessary.

Restore restrictions in a MirrorView environment

You can take SnapView/Clone and SnapView/Snap replicas of the MirrorView/A or MirrorView/S secondary devices. These replicas cannot be restored to the secondary or the primary.

Troubleshooting restore in a Microsoft Cluster

In a Microsoft Cluster environment, a restore failure can result if file share resources are preventing disks from being unmounted due to open handles being held on the file system.

Use one of the following workarounds to prevent the restore failure:

Workaround 1

1. Open Cluster Administrator and take all the file shares related to the disks being restored offline.
2. Restore the replica.

Workaround 2 (only applicable for Windows 2003 clusters)

1. If restore fails after Workaround 1, put the disks into Extended Maintenance Mode manually using the cluster.exe tool (Microsoft Knowledge Base 903650).
2. Restore the replica.

Using application callout scripts

Application callout scripts allow you to add customized actions to Replication Manager at several points throughout the replication, mounting, and restore processes. Replication Manager calls these executable scripts based on the names of the scripts and their locations in the Replication Manager host.

The callout scripts:

- ◆ Must be located in the same directory as `irccd` on the host. If you installed the Replication Manager agent software to the default location, the directory for callout scripts is:
 - `/opt/emc/rm/client/bin/` (UNIX)
 - `C:\Program Files\EMC\rm\client\bin\` (Windows)
 - `C:\Program Files (x86)\EMC\rm\client\bin\`
- ◆ Must be an executable file or shell script (UNIX) or a `.bat` file (Windows).
- ◆ Must return zero as a successful result.
- ◆ Must be named using the naming conventions shown in the next section.

Application callout naming conventions

Non-federated

This section describes application callout naming conventions.

The convention for non-federated callout scripts (SharePoint excepted) is as follows:

```
IR_CALLOUT_<application_set_name>_<job_name>_<n>
```

where:

`<application_set_name>` is the name of the application set that contains the job that will run the script.

`<job_name>` is the name of the job that will run the script (within the application set defined above).

`<n>` is a number that defines when the script should run (refer to the list of choices in the next section).

Federated The callout naming convention used when setting up scripts for a federated application set is as follows:

```
IR_CALLOUT_<application_set_name> (hostname)_<job_name> (hostname)_<n>
```

where:

<application_set_name> is the name of the application set that contains the job that will run the script.

(hostname) is the name of the host where the script resides, or in the case of a mount host, the name of the host where the replica was created (not the name of the mount host). The parentheses are literal characters; they must be included as part of the name. A space must precede (hostname) in each case.

<job_name> is the name of the job that will run the script (within the application set defined above).

<n> is a number that defines when the script should run (refer to the list of choices in the next section).

Separate scripts must be specified for each host in the federated application set.

SharePoint A SharePoint job is comprised of SQL Server and/or SharePoint search indexes and a SharePoint VSS Writer job. The conventions for callout scripts for SharePoint reflect the structure of SharePoint jobs.

For the SQL Server/SharePoint search jobs the convention is:

```
IR_CALLOUT_<application_set_name> (hostname) sql_<SQL_job_name> (hostname)_<n>
```

where:

<application_set_name> is the name of the application set that contains the job that will run the script.

(hostname) is the name of the host where the script resides, or in the case of a mount host, the name of the host where the replica was created (not the name of the mount host). The parentheses are literal characters; they must be included as part of the name. A space must precede (hostname) in each case.

<SQL_job_name> is the name of the SQL Server/SharePoint search job that will run the script (within the application set defined above).

<n> is a number that defines when the script should run (refer to the list of choices in the next section).

An example is:

```
IR_CALLOUT_abcd (host087) sql_abcd_clone (host087)_20.bat
```

For the SharePoint VSS Writer job the convention is:

```
IR_CALLOUT_<application_set_name> (hostname) vss_<VSS_job_name> (hostname) vss_<n>
```

where

<application_set_name> is the name of the application set that contains the job that will run the script.

(hostname) is the name of the host where the script resides, or in the case of a mount host, the name of the host where the replica was created (not the name of the mount host). The parentheses are literal characters; they must be included as part of the name. A space must precede *(hostname)* in each case.

<VSS_job_name> is the name of the job that will run the script (within the application set defined above).

<n> is a number that defines when the script should run (refer to the list of choices in the next section).

An example is:

```
IR_CALLOUT_abcd (host091) vss_abcd_clone (host091) vss_20.bat
```

Specifying callout script processing time

Replication Manager starts the script at certain times throughout the process of replication, mounting, or restore. You can tell Replication Manager when to run the script based on the number that you put into the name.

For example, suppose you want to run your own script at the beginning of the prepare phase of a specific replication. If the application set was called *acct data* and the job was called *backup*, your script would be called:

```
"IR_CALLOUT_acct data_backup_20" (UNIX/Linux)
```

```
"IR_CALLOUT_acct data_backup_20.bat" (Windows)
```

Note: The previous script name is enclosed in quotes because there is a space in the name (*acct data*). Always use quotes when creating the file if the name includes spaces.

Script performance optimization

To optimize performance of job processing when user-written scripts (for example, pre- and post-replication, postmount, callout, backup scripts) are used, scripts should generate output only for significant events. Avoid commands that generate less important output, or redirect less important output to a separate log file. Output lines from scripts are processed by the Replication Manager logging facility, which prints lines to client and server log files, displays lines to the console progress window, and saves lines to the job history in the Replication Manager internal database. Because of this, scripts that generate a lot of output can slow down the job processing and consume extra amounts of Replication Manager Server CPU time.

Resolving script failures on Windows Server 2008 and Windows Server 2012

Callout script failure on Windows Server 2008 and Windows server 2012 may be related to User Account Control (UAC) preventing the creation of XML files. The workaround is to set the environment variable EMC_ERM_CALLOUT_DIR on the Replication Manager Agent host to point to a directory that is writeable by the SQL Server or Exchange account specified in the application set. For more information about the XML files, refer to [“Gathering information about alternate file locations” on page 241](#).

Application callout numbers for replication

[Table 14 on page 238](#) lists the entry points and entry point numbers for callout scripts called during replication.

Table 14

Callout script identifiers for replication

Callout script	The script is called
10	At the beginning of replication.
20	Before synchronizing the mirrors.
30	After Replication Manager has synchronized the mirrors.
40	After the application is quiesced.
50	After Replication Manager finishes the mirror split. (Not for federated or SharePoint.)
60	After the application is returned to normal processing.

Callout scripts 10 through 60 run on the production server.

Note: Callout scripts 10 through 60 do not have an XML file passed to them. Do not attempt to gather information from the XML file within those scripts. For more information about the XML files, refer to [“Gathering information about alternate file locations” on page 241.](#)

Application callout numbers for mount and restore

Table 15 on page 239 lists the entry points and entry point numbers for callout scripts called during mount and/or restore operations.

Table 15 Callout script identifiers for mount, failover, and restore

Callout script	The script is called
100	At the beginning of the process.
110	At the beginning of failover process (for Celerra iSCSI or VNXe iSCSI replica promotion).
200	Before checking that target devices are in the correct state.
300	Before the application recovery process starts. The 300 callout is valid only for mount operations in which some recovery occurs before file systems are made visible.
400	After checking the application state to verify application recovery is in progress.
500	After storage is recovered or mounted.
510	After the failover process is complete (for Celerra replica promotion).
550	After the network files have been copied and before the database is recovered. Use the 500 callout to make changes to the initialization file in Oracle before the application starts.
600	After application recovery is complete.

These scripts run on the production server in the case of a restore or on the mount host in the case of a mount operation.

Note: If you use the XML file that is passed to the callout scripts with a low-numbered callout script, the information passed in the XML file may be incomplete because all information may not be available during the early stages of mount/restore processing. For more information about the XML files, refer to [“Gathering information about alternate file locations” on page 241.](#)

Application callout numbers for unmount

Table 16 on page 240 lists the entry points and entry point numbers for callout scripts called during unmount operations. These scripts run on the mount host.

Table 16 Callout script identifiers for unmount

Callout script	The script is called
1100	At the beginning of the unmount operation.
1200	Before checking that the user environment can support the unmount.
1300	Before shutting down the user application.
1400	Before unmounting file system/deporting volume groups.
1500	At the end of the unmount operation.

Note: If you use the XML file that is passed to the callout scripts with a low-numbered callout script, the information passed in the XML file may be incomplete because all information may not be available during the early stages of unmount processing. For more information about the XML files, refer to [“Gathering information about alternate file locations” on page 241.](#)

Callout script permissions

Whenever Replication Manager runs a callout script, the script is executed using root privileges (UNIX) or administrator privileges (Windows).

Running PowerShell commands from callout scripts

To run PowerShell commands from callout scripts:

1. Specify the full path name to your PowerShell command file in the .bat file:

```
powershell -command C:\PshellCommands.ps1 <nul
```

2. Set the PowerShell execution policy so you can run your script. For example, the first line in the .bat file should look like the following for an unrestricted policy:

```
powershell -command set-executionpolicy unrestricted <nul
```

3. To ensure correct termination of your PowerShell session, add **<nul** to the end of the line that calls your PowerShell script.

Gathering information about alternate file locations

Replication Manager can mount data to an alternate location using one of two methods (alternate root path or substitution table). These methods provide different means to help Replication Manager determine the new location where data should reside on the mount or production host. Placing files in alternate locations causes special problems for callout scripts because the author of the script may not know in advance the new location of the application data.

To find the files mounted to alternate locations, Replication Manager creates an XML file that contains all the location data necessary to find the appropriate objects that have been mounted and perform operations on those objects. The XML information has the following structure:

```
<Record version='recordversion'>
  <AppRecord type='apptype'>
    /* Specific info. about applications goes here*/
  </AppRecord>
  <StorageRecord>
    /* Specific info. about storage goes here*/
  </StorageRecord>
</Record>
```

The following sections describe the specifics of the XML information.

Application attributes that are listed in [Table 17 on page 242](#) are all entered in standard XML format.

Application record portion of the XML

The application record of the XML defines the specifics that relate to the supported applications. [Table 17 on page 242](#) lists the possible attributes by application.

Table 17 Attributes by application type (page 1 of 2)

Application type	Possible XML attributes
Oracle	SID name='sid_name' alt='alternate_sid_name' Tablespace name='tablespace_name' Datafile (<i>within a tablespace</i>) path='data_file_path' alt='alternate_data_file_path' ArchiveLog name='archive_log_path' alt='alternate_arch_log_path' ControlFile name='controlfilepath' alt='alternate_control_file_path' InitFile name='init_file_path' alt='alternate_init_file_path' PasswordFile name='password_file_path' alt='alternate_password_file_path'
UDB	Instance name='instance_name' alt='alternate_instance_name' Database name='database_name' alt='alternate_database_name' Tablespace name='tablespace_name' Datafile (<i>within a tablespace</i>) path='data_file_path' alt='alternate_data_file_path' ControlFile name='controlfilepath' alt='alternate_control_file_path' ArchiveLog name='archive_log_path' alt='alternate_arch_log_path'

Table 17 Attributes by application type (page 2 of 2)

Application type	Possible XML attributes
Microsoft Exchange 2007	VSSWriter name="MicrosoftExchangeWriter (<instance_name>)" StorageGroup (<i>within VSSWriter</i>) name='storage_group_name' InformationStore (<i>within storage group</i>) name='information store name' Datafile (<i>within an Information Store</i>) path='data_file_path' alt='alternate_data_file_path'
Microsoft SQL Server	Database name='database_name' alt='alternate_database_name' Instance name='instance_name' alt='alternate_instance_name' Datafile (<i>within an instance</i>) path='data_file_path' alt='alternate_data_file_path'
SharePoint	Database name='database_name' alt='alternate_database_name' Instance name='instance_name' alt='alternate_instance_name' Datafile (<i>within an instance</i>) path='data_file_path' alt='alternate_data_file_path'
File System	FilesystemRecord name='filesystem_path' alt='alternate_filesystem_path' type='type_of_filesystem'

Storage record portion of XML

The XML file that contains the location data also contains important information about the storage devices and what information is stored on those devices.

The structure of the storage record is as follows:

```
<StorageRecord>
  <Device name='device_name'
alt='alternate_device_name'>
    <Filesystem name='fs_name' type='fs_type'>
    </Filesystem>
    <VolumeGroup name='vg_name' type='vg_type'>
    </VolumeGroup>
    <Gun source='GUN_source' target='GUN_target'>
    </Gun>
  </Device>
</StorageRecord>
<HostRecord name='host_name' OS='operating_system'
version='vers'></HostRecord>
```

The XML file is passed to each callout script as the first parameter when Replication Manager runs that script. If you need information in the XML file, you can access the XML by reading the filename from the script's first parameter. See the following example.

To access the contents of the XML file from a UNIX-based callout script, add a line to your script that is similar to the following:

```
cat $1 > /tmp/xml_content.xml
```

A similar line of code from a Windows-based script might look like the following:

```
copy %1 C:\Temp\xml_content.xml
```

Troubleshooting mount failures on the array

This section describes problems that may occur when mounting to a CLARiiON or Symmetrix array.

CLARiiON and Symmetrix mount failures

Problem a Replication Manager mounts or unmounts the wrong devices after host reboot.

Solution

Set the following environmental variable in order to update symapi db automatically on every reboot:

```
EMC_ERM_STARTUP_DISCOVER=1
```

CLARiiON only mount failures

If a mount of a CLARiiON replica fails with one or more of the following errors, follow the troubleshooting steps described below:

Problem Mount of a replica fails with an error message similar to the following:

```
000070: Attempt to mount file system G:\ has failed with
an error as follows: mount failure, see logs for more
details. /*e*/
```

Solution

Edit the list of mounted devices in the Windows Registry:

Note: If you are not familiar with Windows Registry Editor, you should use it with caution. If you make a mistake, your system may become unusable.

1. Click **Start > Run** and type **regedit**.
2. Using **regedit**, perform the following changes:
 - Navigate to the registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\MOUNTEDDEVICES
```
 - Look for drive entries that meet these criteria:
 - You are certain they are not mounted on the mount host.
 - They include \DosDevices as a prefix.

For example, if E:\ is not mounted delete the following entry:

```
\DosDevices\E:
```

Problem Mount of a CLARiiON iSCSI replica failed with an error similar to the following:

```
002020 ERROR: The host is not connected to the iqn:
iqn.1992-05.com.emc:apm000424036380000-9, any
operation for this iqn will fail because of that.
Please make sure it is connected.
002021 ERROR: Array
iqn.1992-05.com.emc:apm000424036380000-9 is not
accessible, please check its connectivity.
000014 ERROR: Storage Services operation PrepareMount
failed with an error as follows: Storage Services has
terminated with an error.
000014 ERROR: Storage Services operation processPrepare
failed with an error as follows: Storage Services has
terminated with an error.
000071 ERROR: The mount activity has failed.
```

Possible cause

The mount host is not logged into the target IQN.

Proposed solution

Use the Microsoft iSCSI Initiator on the mount host to log in to the target.

Problem CLARiiON replica mount fails with one or both of the following errors:

```
000358: Replication object clone LUN 123 from source LUN
45 on CLARiiON APM00030700175 was not able to be made
visible as a host device on host myhost123 after a bus
rescan.
000034: Function getHostDevice failed with an error as
follows: No non-PowerPath host device could be found
for "APM00030700175:0123". Please consult the Release
Notes for a list of possible causes and resolution
steps.
```

Possible causes and solutions

Possible causes and solutions for this problem include the following:

- ◆ Your mount host is not configured correctly.

Solution: Run the Config Checker to verify that the environment on the mount host conforms to requirements. Take corrective actions to ensure a clean run of the Config Checker. Retry the mount after corrections have been made.

- ◆ Incorrect zoning or other SAN issues.

Solution: Using Navisphere, verify that you can export LUNs to the mount host. First, add a LUN to the storage group of the mount host. Then, use the appropriate OS utilities to rescan for SCSI devices. If the mount host has never seen CLARiiON devices, rebooting the mount host may be necessary to see the LUNs the first time.

- ◆ You are using an unsupported host bus adapter (HBA).

Solution: The HBA in the host may not be supported by EMC on the CLARiiON. Check the EMC *Replication Manager Support Matrix*. You may need to upgrade to a more recent HBA. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

- ◆ You are using an unsupported OS driver, or driver firmware, for your host bus adapter.

Solution: Consult the HBA manufacturer's website to determine the latest EMC-approved firmware and OS driver and update if necessary.

- ◆ Your host bus adapter is configured incorrectly.

Solution: Each HBA has settings that may affect mounting of replications. If you are using an Emulex HBA, make sure the `exlcfg` utility was run with the `--emc` (dash dash) option.

- ◆ You are using an older version of EMC PowerPath®.

Solution: Some older versions of PowerPath may not work. Consult the EMC Replication Manager Support Matrix, and if necessary, make sure you have the most recent version of PowerPath installed.

- ◆ No snapcache available to mount "snap of clone" replica.

Solution: A "snap of clone" type mount was attempted, but the mount host could not create the "disposable" snap that it was going to use to mount because snapcache was not available.

Dedicate more free LUNs in the snapshot cache to support the associated replication.

- ◆ Not using the appropriate Storport drivers (with the appropriate Microsoft service pack) on Windows Server 2003 systems.

Solution: Use the Storport drivers that Microsoft recommends and be sure that the Microsoft Storport Hotfix has been installed on your Windows 2003 systems.

- ◆ Navisphere CLI and FLARE versions not compatible.

Solution: Verify that you are running compatible software versions.

Replication Manager can take advantage of Symmetrix and CLARiiON and VNX consistent-split technology. This chapter provides the following sections on how to use consistent-split technology:

- ◆ Defining consistent-split technology 250
- ◆ Creating consistent-split replicas 253
- ◆ Mounting consistent-split replicas 256
- ◆ Restoring consistent-split replicas 258

Defining consistent-split technology

Replication Manager can take advantage of consistent-split technology to create consistent copies of application data.

Symmetrix consistent split

The latest Symmetrix storage arrays with TimeFinder and consistent-split technology can ensure that the data stored in a replica is consistent even if the replica spans multiple devices stored on multiple Symmetrix arrays.

A replica created using consistent split is also known as a crash-consistent copy. After restarting the application, no part of the data will be transactionally out of sync, because any transactions that were in progress and not yet committed will not be part of the replica that gets mounted or restored. The replica will be consistent for the point in time when it was created. Consistency is assured not by the application, but by the hardware and TimeFinder or Open Replicator software.

CLARiiON and VNX consistent split

CLARiiON and VNX storage arrays with a supported FLARE release and the SnapView enabler software support consistent-split technology for both clones and snapshots. This feature preserves the point-in-time restartable copy of data in a set of source LUNs belonging to one or more applications.

During the creation of a consistent copy, SnapView delays any I/O requests to the set of source LUNs until the session has started on all LUNs (in the case of a snapshot) or during a consistent fracture (in the case of clones). This preserves the point-in-time restartable copy on the entire set of LUNs.

To configure CLARiiON consistent-split replications, consider the following:

- ◆ Production storage must reside on a CLARiiON array with FLARE Operating Environment version 2.19.
- ◆ All of the production storage for the Application Set must reside on the same CLARiiON array.
- ◆ The consistent-split replication cannot exceed the maximum clone or snap limit per CLARiiON array model. [Table 18 on page 251](#) provides the maximum figures for each array.

Table 18 Maximum snap or clone limits per consistent split replica

Array	Snapshot source LUN limit in consistent start		Clone source LUN limit in consistent fracture	
	FLARE 28 and prior	FLARE 29 and later	FLARE 28 and prior	FLARE 29 and later
AX4 Series	8	N/A	8	N/A ^a
CX3-20	8	N/A	8	N/A
CX3-20i	8	N/A	8	N/A
CX300 Series	8	N/A	8	N/A
CX400 Series	8	N/A	8	N/A
CX500 Series	8	N/A	8	N/A
CX3-40	16	N/A	16	N/A
CX3-40i	16	N/A	16	N/A
CX3-80	16	N/A	16	N/A
CX600 Series	16	N/A	16	N/A
CX700 Series	16	N/A	16	N/A
CX4-120	8	32	16	32
CX4-240	8	32	16	32
CX4-480	16	64	32	64
CX4-960	16	64	32	64

a. N/A denotes that CX3 and AX4 arrays do not support FLARE 29.

Restoring consistent split snapshot replicas

Replication Manager uses a single snapshot session to create a CLARiiON consistent-split snapshot, rather than one snapshot session per LUN. Since the replica is created as a single multi-LUN snapshot session, you cannot restore only a subset of the LUNs in a snapshot. Therefore, you can restore only the *entire* replica. This "all or nothing" restore limitation does not apply to consistent split clone replicas.

This chapter describes the considerations you should remember as you use consistent-split technology to create replicas, mount those replicas, and restore from those replicas.



CAUTION

Replicas created using consistent split without hot backup mode will be consistent as of the time they were created, in other words, they are in a crash-consistent state. That means that you can recover them to the point in time when the replica was created, but it is not possible to roll forward logs for applications that use log files.

If you have SAP with BRbackup compliant replicas, it is mandatory that you use online without hot backup mode.

Creating consistent-split replicas

Replicas created with consistent split have different requirements than replicas created without consistent split. Consistent-split options also differ depending on the application for which you are creating the replica. The wizard panel shown in Figure 121 on page 253 allows you to choose whether to use consistent split to create your replica.

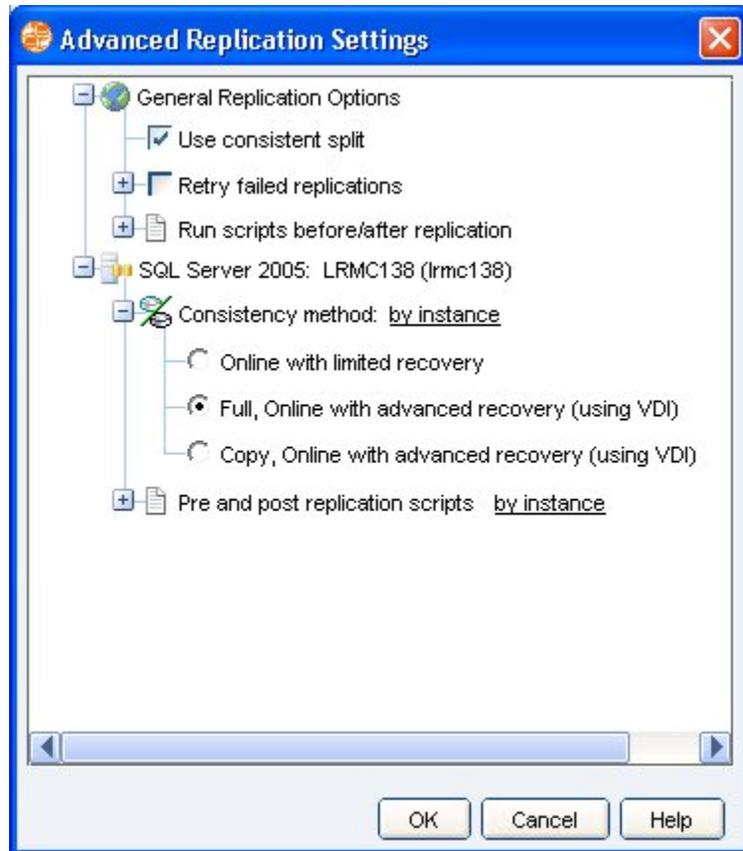


Figure 121 Advanced Replication Settings panel from job wizard

For configuration guidelines describing how to organize your application data on the storage array when using consistent split, refer to the appendix for the specific application found at the end of this guide.

Creating consistent-split replicas in Oracle

When you create an Oracle replica using consistent split, the replication options that you can choose from are expanded to include the following options:

- ◆ **Offline** — Replication Manager takes Oracle offline before using consistent split to create the replica.
- ◆ **Online with hot backup mode** — Replication Manager puts the database into hot backup mode before using consistent split to create the replica.
- ◆ **Online without hot backup mode** — Replication Manager creates a point in time replica with consistent split alone. This means that consistency is ensured by consistent split alone.

Note: This is the mandatory mode for SAP BRbackup compliant jobs. If you have SAP, you must use Online without hot backup mode.

A wizard panel allows you to choose the replication options for an Oracle consistent-split replica.

When you choose to create Oracle replicas using consistent split, the replicas automatically contain the Oracle datafiles, control files, and redo logs. These components are not optional because they are necessary for a crash restart. In addition, Replication Manager asks you to choose whether to include Initialization files as part of the replica.

When you choose Initialization files (either text-based p-files, or binary sp-files) enter a fully qualified path and filename of the file you want to replicate. Click **Browse** to find the file by navigating to it.

Note: Initialization files that you choose to add to the replica must reside on the storage array. You must choose initialization files if you want to mount the replica in Recover or Read-Only mode. If you choose initialization files, the entire file system where that file(s) reside will be replicated.

For detailed information on consistent split Oracle replication processing in all three modes, refer to “[Comparing Oracle replication choices](#)” on page 349.

Creating consistent-split replicas in SQL Server

When you create SQL Server replicas using consistent split, the replication options that you can choose from are expanded to include the following options:

- ◆ **Online with Advanced Recovery (Using VDI)** — Replication Manager uses Microsoft SQL Server Virtual Device Interface mode to quiesce the SQL data and also uses consistent split.
- ◆ **Online with Limited Recovery** — Replication Manager creates a point in time replica with consistent split alone; Microsoft SQL Server VDI mode is not used.

The wizard panel allows you to choose the replication options for a SQL Server consistent-split replica.

Note: If you choose to use consistent split and VDI mode to quiesce the SQL Server data in an application set that contains other applications, ensure that you understand how the applications are quiesced.

Both Online with Advanced Recovery (Using VDI) and Online with Limited Recovery using consistent split will freeze the input/output to the SQL Server application prior to creating the replica, leading to an inconsistency between SQL and other applications in the application set unless other provisions have been made to quiesce the data across the applications.

When you choose to create SQL Server replicas using consistent split, there are no optional components for you to specify. SQL Server replicas must include SQL Server datafiles, filegroups, and transaction logs to maintain consistency.

Note: Each of these components must reside on the storage array.

For detailed information on SQL Server replicas using consistent split, refer to [Appendix C, “SQL Server Procedures.”](#)

Mounting consistent-split replicas

Because a consistent-split replica can contain more components than a replica created without consistent split, the mount operation may require extra steps. The mount considerations related to mounting a consistent-split replica are described next.

Mounting Oracle replicas using consistent split

When you create an Oracle replica using consistent split, you must select parameter files that you want Replication Manager to copy to the Replication Manager Server. You can choose from either of the following types of initialization files:

- ◆ text p-file
- ◆ binary sp-file

If you are using binary sp-files, Replication Manager will discover the files and default the location so that you do not have to enter it manually.

Note: sp-files are binary files. If a database uses sp-files then the Oracle agent converts the file to a text file before sending it to the Replication Manager Server.

You can also choose whether to replicate the archive log directories. Replication of archive log directories is optional.

Mounting SQL Server replicas using consistent split

If a SQL Server replica was created using consistent split, you can later mount that replica in any of the following modes:

With or without VDI mode

- ◆ File System
- ◆ Attach Database

With VDI mode only

- ◆ Recovery
- ◆ No Recover
- ◆ Standby

If you choose Attach Database, then Replication Manager mounts the replicated file systems necessary and attaches the databases in those file systems. The resulting database is a point-in-time copy of the replicated database at the time the replica was created. You cannot apply logs to the database to roll it forward if the database was created using consistent split and you chose to create a replica without VDI.

Mounting file system replicas using consistent split

You can also mount UNIX and NTFS file system replicas that were created using consistent split. These replicas will be consistent as of the time they were created, however, applications within the file system will be in a crash-consistent state. That means that you can recover them to the point in time when the replica was created but it may not be possible to roll forward logs for applications that use log files. Ensure you understand these considerations before you create a replica using consistent split.

Restoring consistent-split replicas

The following sections describe considerations that apply when you restore an original production instance from a replica that was created using consistent-split technology. By definition, the restore process overwrites the production devices with the contents of the devices on which the replica resides.

Restoring Oracle from consistent-split replicas

When Replication Manager restores an Oracle instance from a replica that was created using consistent-split technology, *all components that you choose to restore will overwrite the existing components on the production server instead of sending these components to a temporary directory!* Restored components include redo logs, archive logs, control file, and tablespaces.



CAUTION

Restoring redo logs, archive logs, and control files that are part of the replica overwrites all existing files on the production host. Ensure that you understand the implications of restoring before you proceed.

In addition, you must be aware of the potential for *shared storage issues* when you are restoring an Oracle database. For example, if you do not choose to restore any of the optional components, but one or more of those optional components shares the same hypervolume with one or more of the required components, those optional components will be restored, overwriting the existing files on the production host.

Restoring SQL Server from consistent-split replicas

EMC does not recommend using a consistent-split replica to restore a SQL Server instance. Instead, you should create a replica using VDI for the purposes of restore.



CAUTION

If you do restore a replica created using consistent split and not VDI, you will not be able to roll the database forward beyond the time the replica was taken. Replicated databases can be rolled forward only if they were created using VDI.

Replication Manager can replicate, mount, and restore federated application sets. This chapter includes the following topics:

- ◆ Federated compared with composite application sets 260
- ◆ Creating federated application sets 262
- ◆ Creating federated jobs 265
- ◆ Running a federated replication 269
- ◆ Federated mount options 271
- ◆ Restoring from a federated replica 276

Federated compared with composite application sets

The terms *federated* and *composite* can both refer to a type of application set that contains data from more than one application source. So what is the difference between federated and composite application sets?

- ◆ The term *federated* refers to an application set that includes data sources that represent data *on multiple production hosts* that may even be located on more than one storage array. A federated application set requires a method to maintain consistency across all the components of the application set.

Note: In Symmetrix environments one host in the application set must have Fibre Channel connectivity to all arrays in order to facilitate the EMC Consistency Assist (ECA) functionality across all the components of the federated application set.

- ◆ The term *composite* refers to an application set that includes data sources from multiple applications *residing on a single production host*.

In both types of application set the array technologies must be the same across data sources.

Note: You can configure Replication Manager to replicate only those components of the federated environment that are located on Symmetrix storage arrays. Components and business processes located elsewhere (for example, as part of a federated managing application) cannot be part of the replica unless they are also located on Symmetrix storage.

Environments that support federated application sets

Federated application sets are supported for file system, Oracle, UDB, and SQL Server as summarized in [Table 19 on page 261](#).

Table 19 Summary of OS and application support for federated application sets

	Solaris	AIX	HP-UX	Windows
File system	Replicate, Mount, Restore	Replicate, Mount	Replicate, Mount	Replicate, Mount, Restore
Oracle	Replicate, Mount, Restore	Replicate, Mount	Replicate, Mount	
UDB	Replicate, Mount, Restore	Replicate, Mount		
SQL Server				Replicate, Mount, Restore

On Windows, data can reside standalone or in an MSCS cluster.

All Oracle, SQL Server, and UDB options that are supported in nonfederated environments (for example, RAC support with Oracle) are also supported in federated environments.

Federated replicas must reside on one or more Symmetrix storage array(s) with TimeFinder software and Enginuity Consistency Assist (consistent split) technologies.

Federated replicas can be created using any of the following replication technologies:

- ◆ TimeFinder/Mirrors
- ◆ TimeFinder/Clones
- ◆ TimeFinder/Snaps (VDEVs)

For more information on each of these replication technologies, refer to [“Available replication technologies” on page 58](#).

Federated application sets are not supported in a VMware VMFS configuration.

Creating federated application sets

This section describes how to create an application set that includes federated data. The following considerations apply to federated application sets:

- ◆ Number of hosts you select is limited only by server computing resources and agent activity.
- ◆ User must have valid credentials for all applications that have been chosen as part of the federated application set.
- ◆ A federated application set is defined when you select multiple Oracle, UDB, Microsoft SQL Server and/or file system data sources from more than one host.
- ◆ Microsoft SQL Server instances can only be combined with other Microsoft SQL Server instances and NTFS file systems. Microsoft SQL Server instances cannot be combined with Oracle instances or non-Windows applications or file systems in a single federated application set.
- ◆ Federated application sets support clusters in the same way that nonfederated application sets do.
- ◆ If your environment includes RAID 5 or RAID 6 BCVs and other BCVs other than RAID 5/RAID 6, you should create replicas using TimeFinder/Clones instead of TimeFinder/Mirrors.
- ◆ At least one host in the federated application set must have visibility to all Symmetrix arrays that are included in the federated application set.

Note: The required visibility to storage can be accomplished by exporting one gatekeeper device from each Symmetrix to at least one of the production hosts involved in the federated application set. Replication Manager automatically detects which host has this visibility and uses that host to perform the split operation.

- ◆ Federated application sets cannot contain both basic disks and dynamic disks. All the disks must be of the same type.

To create a federated application set:

1. In the tree panel, right-click **Application Sets** and select **New Application Set** from the context menu.
2. Click **Next** on the **Welcome** screen for the Application Set Wizard. The **Application Set Name and Objects** screen appears; refer to [Figure 122 on page 263](#) for an example of the screen. This screen is used to select application objects in a federated data source.

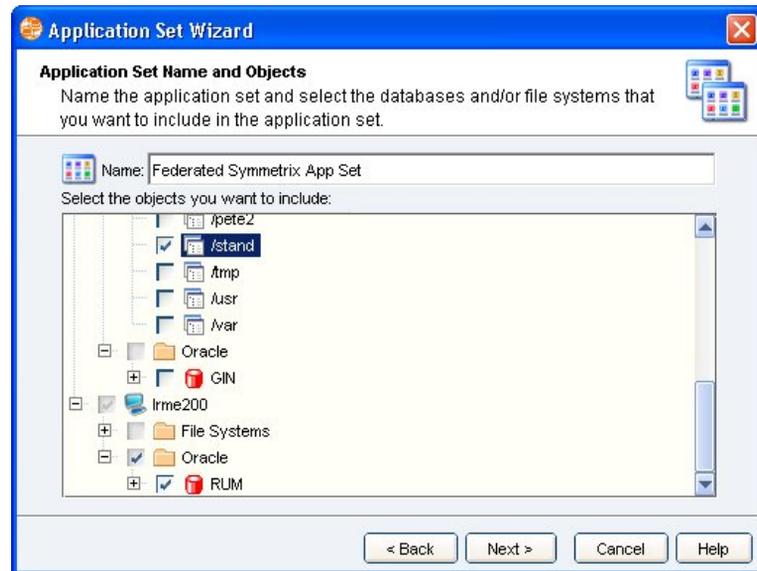


Figure 122 Selecting application set objects (federated application set)

3. To create a federated application set, select one or more supported applications from more than one production host. Complete the remaining screens in the Application Set Wizard as you would for any other application set.

Note: For specific information about other Application Set Wizard panels, click **Help**.

4. The final screen of the Application Set Wizard shows a summary of application set data in a tree view with separate information for each host that is part of the federated application set. Refer to [Figure 123 on page 264](#) for a sample federated application set summary screen.



Figure 123 Federated application set summary screen

Restrictions on composite or federated application sets

There are certain restrictions associated with changing composite or federated application sets you have created in Replication Manager. Those restrictions are as follows:

- ◆ Once you have created a federated application set you cannot modify it in a way that would reduce the number of hosts down to just one (making it nonfederated).
- ◆ You cannot modify a nonfederated application set to add one or more hosts (making it federated).

In order to make changes like those described above, you must create a new application set and optionally delete the old application set.

After adding components to a composite or federated job, you must update mount settings for *all* existing jobs via the Job Properties panel.

The next section describes the specifics of creating a job when your application set includes federated data.

Creating federated jobs

When you create a job for a federated application set, it will be necessary to configure the following settings separately for each part of the federated data:

- ◆ Configure separate replication options for each application that is part of a federated application set.
- ◆ Configure separate mount options for each host that is part of the federated application set.
- ◆ Optionally configure separate pre- and post-replication scripts for each host that is part of the federated application set.

To create a job for a federated application set:

1. If you did not access the Job Wizard automatically from the Application Set Wizard, right-click **Jobs** in the tree panel and select **New Job** from the context menu.
2. Select a federated application set from the **Select the application set you want to create jobs for** drop-down list and click **Next** to continue.
3. Complete the **Job Name and Settings** screen of the wizard. Note that you can only choose Symmetrix replication technologies that support federated application sets.

Note: All source and target volumes should be on the same type of storage. Do not mix Mirror (BCVs) with RAID 5 or RAID 6 protected storage or devices protected with emulation mode.

4. Click **Advanced** to set the Advanced Replication Settings shown in Figure 124 on page 266.

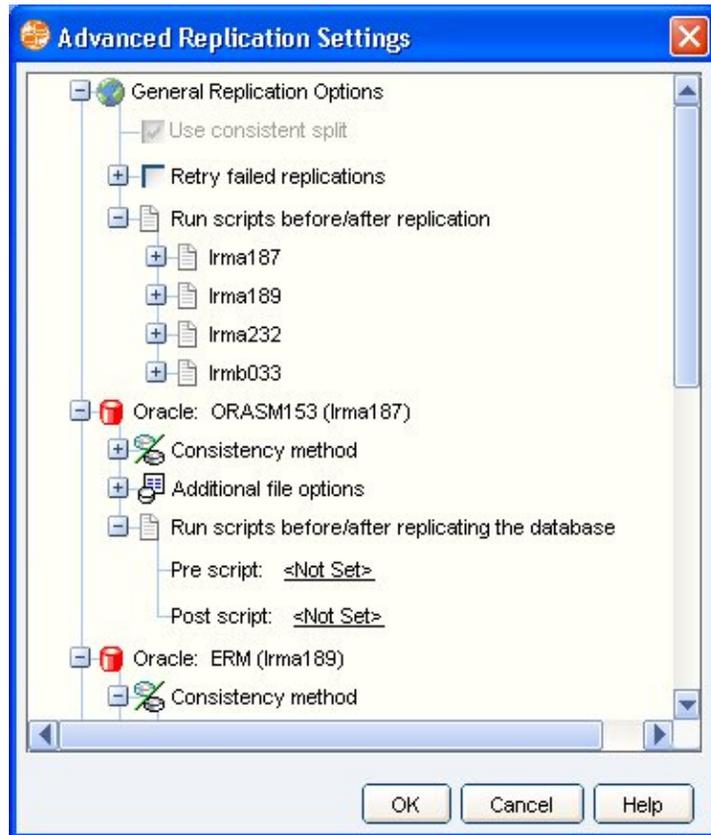


Figure 124 Advanced Replication Settings panel

Table 20 on page 267 describes the advanced settings that apply specifically to federated jobs. For information on other fields, refer to EMC Replication Manager online help.

Table 20 Advanced replication settings that apply to federated jobs

Advanced replication setting	Information
Use consistent split	Use consistent split is greyed out and selected for all jobs related to a federated application set. That is because federated jobs require consistent split to be selected.
Run scripts before/after replication	<p>For federated jobs, you can specify pre- and post-replication scripts for each production host that will run before and after the replication as a whole. These are not related to the replication of any specific application components:</p> <ul style="list-style-type: none"> • The general pre-replication scripts specified here run before any pre-replication scripts that are defined for each application. • The general post-replication scripts specified here run after all post-replication scripts that are defined for each application.
Replication Manager offers one of the following two options for each application	
Run scripts before/after replicating the database (specified under the node for each application)	In addition, you can specify scripts to run before and after you replicate each database (if you chose to specify scripts by database).
Pre- and post-replication scripts instance (specified under the node for each application)	Alternatively, you can specify scripts to run before and after you replicate each instance (if you chose to specify scripts by instance).

5. Click **Next** to continue. Complete the **Replication Storage** screen of the wizard. If you select a pool, ensure that it has Symmetrix storage that can be used to replicate a federated application set.

Note: If you decide not to mount during the job, subsequent mounts should use pools to help Replication Manager select the appropriate storage.

The next section describes setting the Mount Options and completing the wizard.

Running a federated replication

Processing for a federated job occurs in a slightly different way from processing a nonfederated job. Consider the federated environment shown in [Figure 125 on page 269](#).

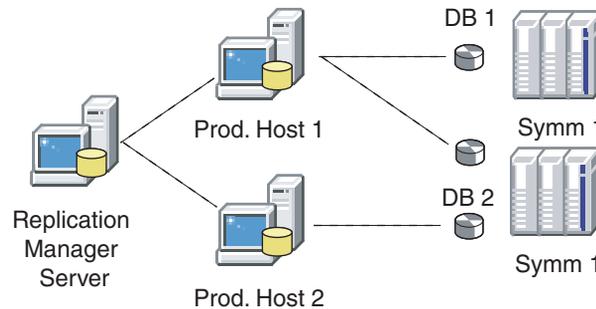


Figure 125 Federated environment

In this environment, the applications are running on two different production hosts and each host is controlling the activities of the application running on that host. Notice that Prod Host 1 has components on both Symmetrix arrays and host 2 has components on only one array.

In this environment, Replication Manager would perform the following steps with a synchronization point at each step. In other words, Replication Manager would run each of the following steps separately on production host 1 and production host 2, however Replication Manager pauses processing after each step until all hosts have reached the end of that step before any host can proceed to the next step:

1. Get application detail for the applications on each host. (Log in to each database and get information about the tablespaces and datafiles that reside on each host.)
2. Get storage detail. (Collect information about the devices where the application objects found in step one are stored.)
3. Identify storage devices required for replication. (Identify suitable target storage to hold the replica and start establishing the mirrors.)
4. Prepare applications on each host for the split.

Note: Step 5 runs from the master process only, not from each individual process on the hosts.

5. Determine the host with visibility to all storage arrays and issue the split command from that host.
6. Return applications to their normal state.
7. Collect catalog data and files required for mounting the applications stored on each host.
8. Run other optional processes such as mount and/or backup operations.

Federated mount options

As part of creating a job, users define the mount options for the federated application set. They can choose whether or not to mount data located on each production host separately.

To set the mount options:

1. Access the Job Wizard as defined in the previous section, or expand the **Application Sets** folder and the specific application set that generated the replica. Right-click the replica and select **Mount** from the context menu.
2. Complete the **Mount Options** screen of the wizard. Federated data spans more than one production host. Therefore, you must complete a separate set of mount options for each production host. Refer to [Figure 126 on page 271](#).

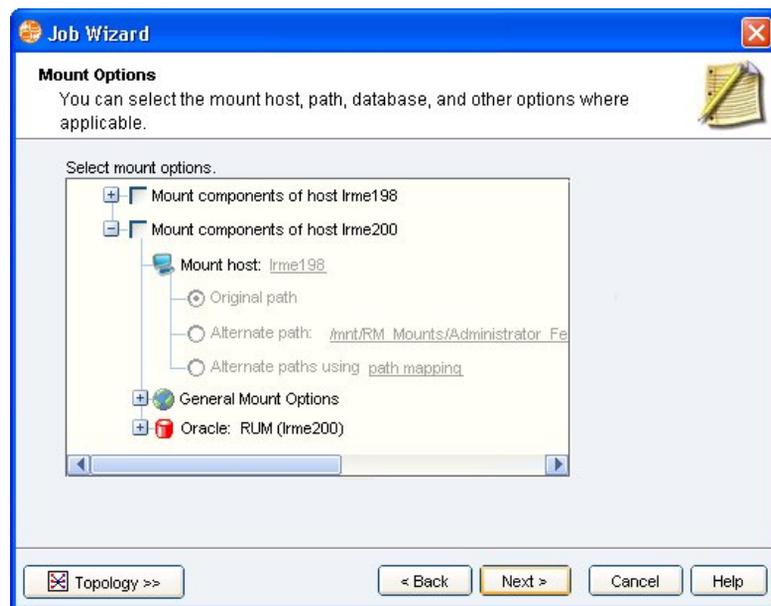


Figure 126 Federated mount options screen

For each production host, you can complete the fields described in [Table 21 on page 272](#) to configure the mount options for a federated job.

Table 21 Mount options configured for each production host separately

Mount option	Information
<p>Mount components of host: <host> If you choose this checkbox you must specify the remaining mount options for this host. Mount options are described below.</p>	<p>Select this checkbox to signify that you want to mount the components of the named production host onto a mount host. You can choose whether or not to mount data that resides on each production host separately.</p> <p>Note: EMC recommends that each federated job specify a mount as part of the job.</p>
<p>Mount host</p>	<p>This list allows you to choose the name of the host to which you want to mount this part of the federated replica.</p> <p>You can choose to mount each production host to a different mount hosts or all of the data to a single mount host.</p> <p>Note: The host you select must have the same operating system as the production host named above.</p>
<p>Original path</p>	<p>Choose this option to mount to the same path used on the production host.</p> <p>Note: You can mount two or more production hosts to the same mount host, but Replication Manager will not warn you if the mount locations conflict. If conflicts occur, the job fails at runtime.</p>
<p>Alternate path</p>	<p>Choose this option to mount to an alternate path that you specify.</p>
<p>Alternate path using path mapping</p>	<p>Choose this option to mount to alternate paths defined using path mapping.</p>

- Also complete the **General Mount Options** for each application in the federated application set. Refer to [Figure 127 on page 273](#).

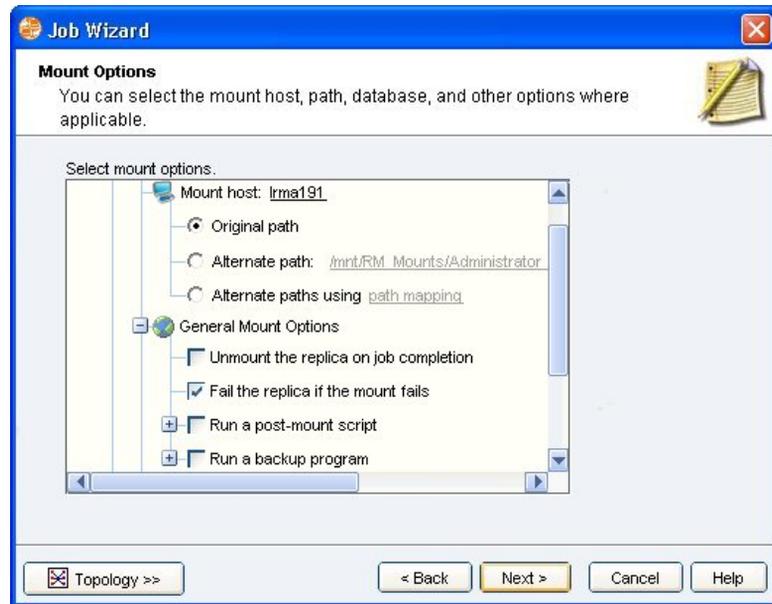


Figure 127 General Mount Options for each application

Most of the general mount options do not change for federated application sets. The following options have some modified behavior:

- If you select **Unmount the replica on job completion**, Replication Manager unmounts the part of the replica that is mounted on this mount host. Replication Manager will not unmount parts of the replica mounted on other mount hosts, unless you choose this checkbox for all hosts.

Note: If you have SAP BRbackup compliant jobs, this option is forced (unchecked).

- If you select the checkbox **Fail the replica if the mount fails**, then if any one mount fails, the entire replication fails and anything that had been mounted will be unmounted.

Note: If you have SAP BRbackup compliant jobs, this option is forced (checked).

If you clear the checkbox **Fail the replica if the mount fails**, then these are the behaviors:

- If all mounts succeed, then the replication status will say Mounted.
 - If some mounts succeed, but some others fail, the status will be Mounted, but there will also be a warning flag attached to the replica.
 - If all mounts fail, the status will say Mount Failed, and a warning flag will be attached to the replica.
 - If you choose to include post mount scripts and/or backup scripts, remember that these scripts must reside on the selected target mount or backup host (respectively).
 - If you select the checkbox **Fail the replica if the post mount script fails**, then if the post mount script fails, the entire replication fails and anything that had been mounted will be unmounted.
 - If you clear the checkbox **Fail the replica if the post mount script fails**, then these are the behaviors:
 - If all scripts return success, then the replication status will say Mounted.
 - If some scripts fail, the status will be Mounted, but there will also be a warning flag attached to the replica.
4. Once you complete the mount options screen, click **Next** to continue. The **Schedule the Job** screen appears.
 5. Complete the **Schedule the Job** screen of the wizard. There are no changes to scheduling that relate to federated jobs specifically. Click **Next** to continue.

Note: When Replication Manager is creating a federated replica that is consistent across hosts, each host proceeds based on its own resources and processing speed and then faster hosts wait for slower hosts to maintain consistency across the entire application set. Completion of federated jobs depends upon the processing speed of the slowest host in the application set.

6. The wizard completes with a summary that describes what options you have selected. Verify the options and click **Finish** to save your job settings.

Note: If you are creating federated jobs, Replication Manager does not allow copy jobs of a federated job.

Partial mounts of federated databases

It is possible to mount data from selected hosts when mounting a federated application set, leaving data from other hosts unmounted. This might occur because you have selected to mount only certain hosts or it may occur because mount operations failed on some hosts and not on other hosts during a mount operation.

In a partial mount situation, Replication Manager allows you to choose to mount all the unmounted portions of the replica, or unmount all mounted portions from the context menu of a replica:

- ◆ If you choose mount in a partially mounted replica, Replication Manager attempts to mount the parts of the replica that were not already mounted.
- ◆ If you choose unmount in a partially mounted replica, Replication Manager attempts to unmount the parts of the replica that are currently mounted.

Restoring from a federated replica

Beginning in Replication Manager Version 5.3, you can restore from a federated replica of data. [Table 19 on page 261](#) summarizes the support per operating system and application for restoring from a federated replica.

Full and partial restores are supported. Restore options are the same as for nonfederated replicas.

All Replication Manager components (Console, Agent, and Server) must be at a minimum of Version 5.3 for federated restore support.

With proper planning and management, Replication Manager users can maintain a balance between the amount of storage available for replicas and how often a replica can be created.

This chapter includes the following sections that describe how to manage replicas effectively:

- ◆ Viewing replicas 278
- ◆ Replicating data on remote storage arrays..... 283
- ◆ Replicating data on Celerra or VNX network file system..... 292
- ◆ Availability of RecoverPoint replicas 298
- ◆ Controlling replica expiration and deletion 300
- ◆ Rotating replicas..... 303
- ◆ Setting retention periods..... 305
- ◆ Understanding Replication Manager scripting 308
- ◆ Acknowledging a failed replica 310

Viewing replicas

To view replica information in the content panel:

1. Expand **Application Sets** in the tree.
2. Select the particular application set that contains the replica you want to view. Information on each replica is displayed in the content panel.

To view replica properties, including replica tablespaces or file systems in a tree diagram:

1. Expand **Application Sets** in the tree.
2. Expand the particular application set that contains the replica you want to view.
3. Right-click the replica whose structure you want to view.
4. Select **Properties**.

A separate window displays replica information in four tabs: General, Objects, Storage, and History Logs. Each tab is shown and described next.

General tab

The General tab shown in [Figure 128 on page 279](#) contains information such as the date when the replica was created; the current replica state, host and application set with which the replica is associated; the mount information; and a checkbox to allow users to enable expiration of the replica.



Figure 128 View replica General tab

Objects tab

The Objects tab shown in [Figure 129 on page 280](#) describes the structure of the data stored in the replica, including the nodes that are either portions of a database or file system, depending on the application associated with the replica. For example, if the application set specified an entire SQL Server database, the tab displays the database as the top level, and then displays file groups and datafiles.

To view what elements are part of the replica on the Objects tab:

1. Click the plus sign (+) to expand a node.
2. Use the scroll bars as needed.

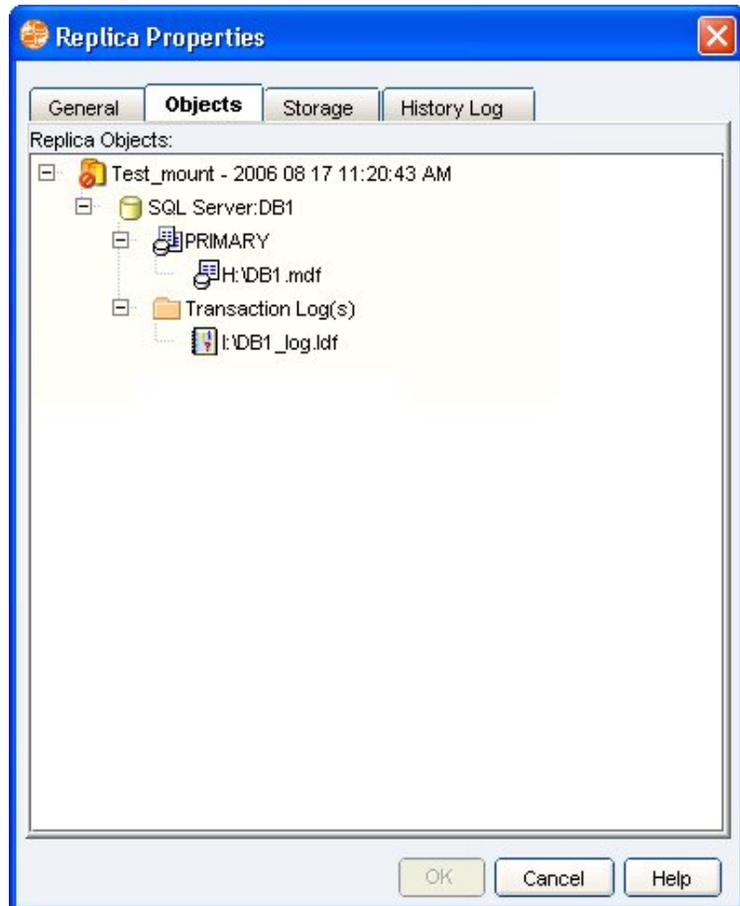


Figure 129 View replica Object tab

Storage tab

The Storage tab shown in [Figure 130 on page 281](#) shows the following information:

- ◆ The type of storage that holds the replica
- ◆ The devices that are available
- ◆ The device type for each device
- ◆ The size of each device
- ◆ The storage array where each device is located

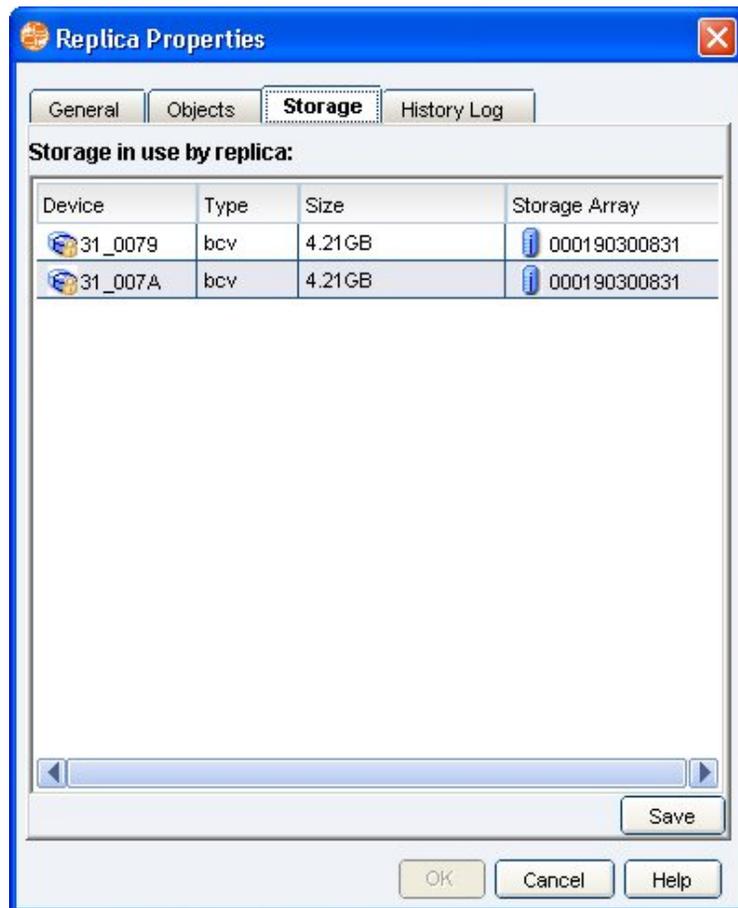


Figure 130 View replica Storage tab

History Log tab

The History Log tab shown in [Figure 131 on page 282](#) shows the details of each step Replication Manager took to create the replica. The details shown here are also listed in the Replication Manager log files. You can display the log in Table View (shown here) or click the **Text View** button to see it in Text View.

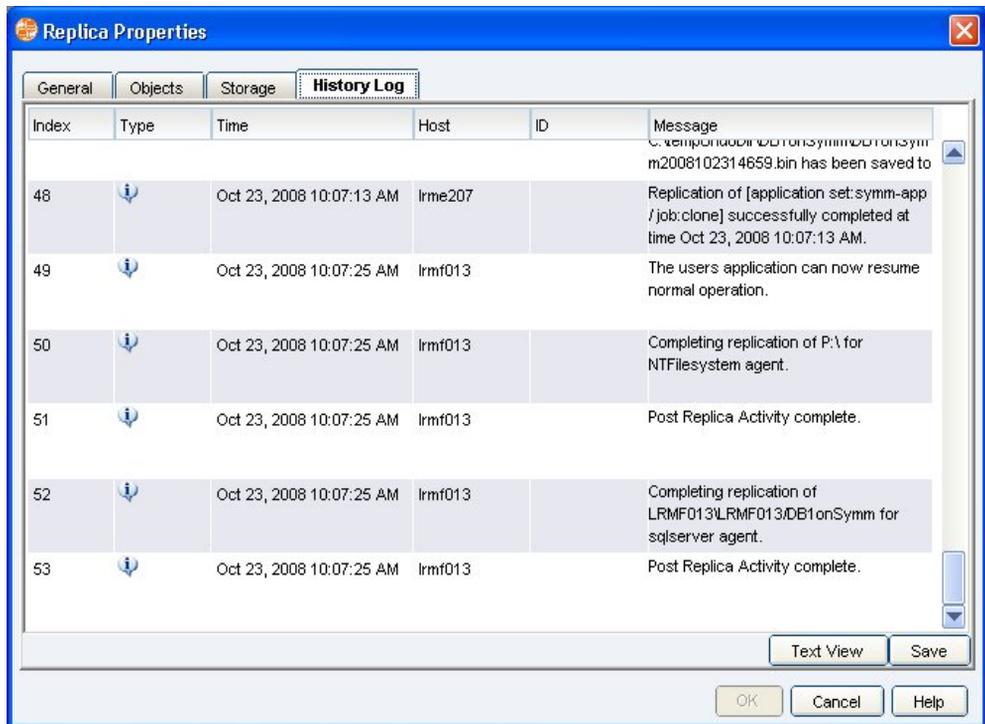


Figure 131 View replica History Log tab

Replicating data on remote storage arrays

Replication Manager supports SRDF and EMC SAN Copy functionality for creating a replica of data on a separate storage array. Remote replica functionality is supported on arrays equipped with remote replication software sold separately.

Note: For information about how to configure your servers and storage arrays to support EMC SAN Copy, refer to the *EMC Replication Manager Administrator's Guide*.

Table 22 on page 283 describes what array combinations are supported.

Table 22 Remote replication array combinations

Technology used	Source array type	Target array type
SRDF	Symmetrix	Symmetrix
Full SAN Copy	Symmetrix CLARiiON VNX	CLARiiON VNX
Incremental SAN Copy	CLARiiON VNX	CLARiiON VNX
MirrorView/S	CLARiiON VNX	CLARiiON VNX
MirrorView/A	CLARiiON VNX	CLARiiON VNX
Replicator	Celerra VNXe	Celerra VNXe

Creating an SRDF replica

When you create an SRDF replica, Replication Manager can use a Symmetrix (R1/Standard) as the source and the target replica will be created on a BCV of the R2. When you run the job to create a remote replica using SRDF/S, Replication Manager performs the following steps automatically:

1. Creates an SRDF R2 copy on the remote Symmetrix system.
2. Creates a BCV copy of the R2 on the remote Symmetrix system as shown in [Figure 132 on page 284](#).

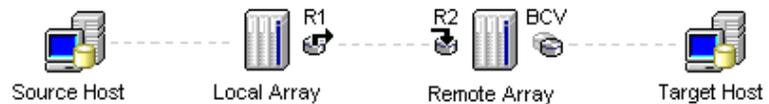


Figure 132 Creating a BCV on a remote Symmetrix

Note: The storage pool you use to create a remote BCV replica should include enough BCVs from the remote Symmetrix to create the replica. You do not have to include the R2 devices in the storage pool.

Creating a SAN Copy replica

You can create the following types of SAN Copy replicas:

- ◆ Full SAN Copy replicas of Symmetrix standards or CLARiiON or VNX LUNs onto a CLARiiON or VNX clone on a remote storage array (out-of-frame)
- ◆ Full SAN Copy replicas of CLARiiON or VNX LUNs onto a CLARiiON or VNX clone within the same storage array (in-frame)
- ◆ Incremental SAN Copy replicas of CLARiiON or VNX LUNs to a remote storage array (out-of-frame)
- ◆ Incremental SAN Copy replicas of CLARiiON or VNX LUNs onto a CLARiiON or VNX clone within the same storage array (in-frame)

Each of these types of SAN Copy replicas are described in detail in the following sections.

Full SAN Copy replica on a remote array

When you create a Full SAN Copy replica or a SAN Copy replica of an existing replica, Replication Manager can use a Symmetrix standard or CLARiiON or VNX LUN as the source. Then, Replication Manager performs the following steps automatically:

1. Creates a copy of the source disk (either a TimeFinder/Snap (VDEV) on Symmetrix or a SnapView snapshot on CLARiiON or VNX). Replication Manager creates a copy of the source disk to reduce the amount of time that the application must be quiesced. After the snapshot has been created, the application can come back online.
2. Creates a Full SAN Copy replica on a CLARiiON or VNX clone located on a remote CLARiiON or VNX storage array using the new VDEV or snapshot as the source. Refer to [Figure 133 on page 285](#).

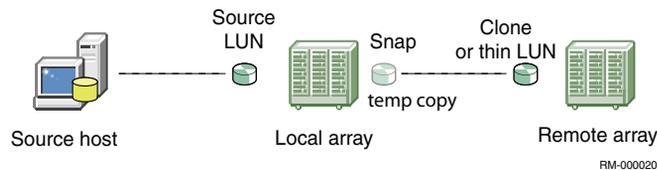


Figure 133 Creating a Full SAN Copy replica on a remote CLARiiON (out-of-frame)

Note: Only one kind of storage must be added to your storage pool to support a Full SAN Copy replica. Although Replication Manager will use VDEVs (for a Symmetrix source) or Snap cache (for a CLARiiON source) on the source storage array, you do not have to add these to the pool in order to create a valid pool for the replication. You must include CLARiiON LUNs on the target storage array to hold the replica. All Symmetrix devices used in a SAN Copy replication operation (intermediate devices or standards) must be exported to the remote CLARiiON. See the *EMC Replication Manager Administrator's Guide* for full details.

Full SAN Copy replica within local array (in-frame)

When you create a Full SAN Copy replica or a SAN Copy replica of an existing replica, Replication Manager uses CLARiiON or VNX LUNs as the source. Then, Replication Manager performs the following steps automatically:

1. Creates a copy of the source disk using a SnapView snapshot on CLARiiON or VNX. Replication Manager creates a copy of the source disk to reduce the amount of time that the application must be quiesced. After the snapshot has been created, the application can come back online.

2. Creates a Full SAN Copy replica on a CLARiiON or VNX clone located within the same CLARiiON or VNX storage array using the new snapshot as the source. Refer to [Figure 134 on page 286](#).



Figure 134 Creating a Full SAN Copy replica within a CLARiiON or VNX (in-frame)

When source LUNs are on multiple CLARiiON arrays and users select in-frame target arrays, the entire replication is localized in that array. Refer to [Figure 135 on page 286](#).

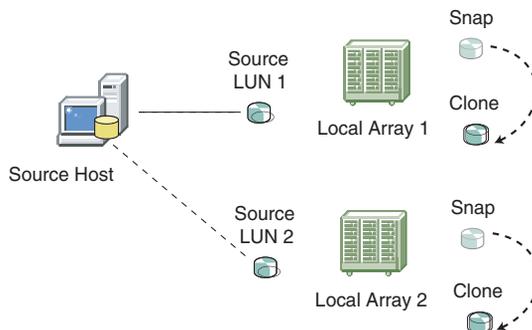


Figure 135 Creating a Full SAN Copy replica on local CLARiiON or VNX target arrays (in-frame)

Incremental SAN Copy replica

When you create an Incremental SAN Copy replica, the source data must be located on a source LUN in a CLARiiON or VNX storage array. Replication Manager performs the following steps automatically:

1. Checks to see if a previous Full SAN Copy replica has been created by the job; if not, a Full SAN Copy replica will be created. Subsequent job runs will create an Incremental SAN Copy. If the job has previously created a replica, then the changes will be updated on the same remote clone used before.

2. Creates an Incremental SAN Copy session for the source devices and the selected target devices. The target devices are cataloged as the replica. Refer to [Figure 136 on page 287](#).

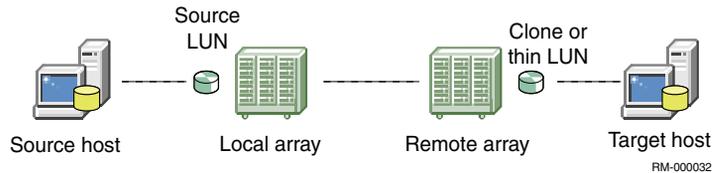


Figure 136 Creating an Incremental SAN Copy replica on a remote CLARiiON or VNX (out-of-frame)

Note: To support an Incremental SAN Copy replica, you must include CLARiiON or VNX LUNs on the target storage array to hold the incremental SAN Copy data.

You can create a SAN Copy replica by following these steps:

1. Create a new job or modify an existing job.
2. Select one of the SAN Copy storage technologies on the second wizard panel:
 - Incremental SAN Copy
 - Full SAN Copy
3. On the **Replication Storage** panel, set up a storage pool that contains the appropriate types of storage to complete the replica process. For specific information about storage needs, click **Required Storage**. Remember that SAN Copy replicas usually need two different types of storage to complete the replication.

Copy a replica using Full SAN Copy

When you create a copy of a replica using Full SAN Copy, Replication Manager can use a replica of a Symmetrix standard, RBCV, snapshot or CLARiiON or VNX LUN as the source. Then, Replication Manager performs the following steps automatically:

1. Creates a copy of the source disk (either a TimeFinder/Snap (VDEV) on Symmetrix or a SnapView snapshot on CLARiiON or VNX). Replication Manager creates a copy of the source disk to reduce the amount of time that the application must be quiesced. After the snapshot has been created, the application can come back online.

2. Creates a Full SAN Copy replica on a CLARiiON or VNX clone located on a remote CLARiiON or VNX storage array using the new VDEV or snapshot as the source. Refer to [Figure 137 on page 288](#).



Figure 137 Creating a Full SAN Copy of an existing replica

Note: Two different kinds of storage must be added to your storage pool to support a Full SAN Copy replica. You must include VDEVs (for a Symmetrix source) or SnapView snapshots (for a CLARiiON or VNX source) on the source storage array. In addition, you must include CLARiiON or VNX LUNs on the target storage array to hold the replica. All Symmetrix devices used in a SAN Copy replication operation (intermediate devices or standards) must be exported to the remote CLARiiON or VNX. See the *EMC Replication Manager Administrator's Guide* for full details.

Mounts and restores of SAN Copy replicas

Replication Manager can mount and restore Full SAN Copy replicas just like any other replica that you create using the traditional replication jobs.

Incremental SAN Copy replicas can be mounted only as read only, or as a snapshot of the replica. Choose **Create and mount a snap of the replica** to cause changes you make to the mounted replica to be discarded on unmount. If you choose this option, Replication Manager creates a snapshot of the target clone to mount and deletes that snapshot when the replica is unmounted.

Note: Incremental SAN Copy restore is a Full SAN Copy restore from the remote snapshot to the source LUN.

Modifying CLARiiON or VNX SAN Copy options

Use the **Advanced** tab under **CLARiiON Properties** or **VNX Properties** to specify how Replication Manager should use clone LUNs and SAN Copy options:

1. Select **Do not remove clone LUNs from clone group** if you do not want Replication Manager to use clones from other clone groups to create a replica. Replication Manager will fail a job if all available clone LUNs are already in existing clone groups for different source LUNs.

Do not select **Do not remove clone LUNs from clone group** (default) if you want Replication Manager to use clone LUNs regardless of whether those LUNs are part of an existing clone group. This is the default Replication Manager behavior.

2. Choose the appropriate SAN Copy session options:
 - **Session Throttle** — Select the value that controls the I/O rate for a SAN Copy session. You can set the session throttle to a value between 1 and 10 where 1 is the lowest rate and 10 is the highest. The default throttle value is 6.
 - **Link Utilization** — Select the available link bandwidth for the copy session. You can use the entire bandwidth of your network connection for incremental SAN Copy sessions, or you can specify only a portion of it. The minimum allowable value is .016 MB/s and the maximum value is 2048 MB/s. The default value is 1.5 MB/s (equivalent to a T1 line).
3. Select the **Override Clone Sync Rate** option to adjust the clone synchronization rate for each CLARiiON array in your replication environment. The default synchronization rate is **Medium**. You may need to lower this setting if you are using a lower-end CLARiiON model. Also keep in mind that when the sync rate for the target LUN is set to **Low**, it may slow the replication process.

Note: Setting the synchronization rate for an array from one Replication Manager Server does not affect the sync rate of other Replication Manager Servers using the same array.

4. Click **OK** to save the configuration.

Remote SRDF/S replication

Replication Manager can also create remote replicas using TimeFinder/Mirror, TimeFinder/Clones, and TimeFinder/Snaps, across an SRDF link. Creating these remote replicas requires two Symmetrix storage arrays connected via an SRDF/S link.



CAUTION

Replicas created using Remote SRDF/S cannot be restored using Replication Manager. These replicas can be mounted to an alternate host if the devices that constitute the replica are visible to the mount host to which you plan to mount the replica. The replicas can also be used as the source for SAN Copy replicas to a CLARiiON array.

Replication of MirrorView/A or /S secondary

Replication Manager supports creation of replicas of the MirrorView/A or MirrorView/S secondary using SnapView/Snap and SnapView/Clone. It also supports creation of Full SAN Copy and Incremental SAN Copy replicas of the SnapView snaps and SnapView clones of the secondary (created from a copy job).

The following facts are true when replicating MirrorView:

- ◆ Synchronization of MirrorView/A sessions must be in manual update mode because synchronization is controlled by Replication Manager.
- ◆ The user must manually establish the MirrorView session.
- ◆ Replication Manager does not manage the MirrorView link.
- ◆ Replication Manager is not involved with creation of the secondary.
- ◆ The MirrorView/A consistency group, if used, must consist of LUNs that are all part of a given application set. Additionally, the consistency group cannot contain any additional LUNs that are not part of the application set.

Replicating data on Celerra or VNX network file system

Replication Manager creates replicas of network file systems (NFS) or Oracle databases that reside on NFS or NFS with the use of the Oracle 11g Direct NFS (dNFS) client. For information about how to discover an NFS using Replication Manager, refer to the *EMC Replication Manager Administrator's Guide*.

Celerra NFS storage failover configurations

This section explains how to create application sets and jobs to replicate data stored on Celerra NFS storage. Celerra NFS environments offer a feature called Celerra storage failover. Replication Manager includes special configuration options that allow replication of data before and after Celerra storage failovers, which cause another Celerra in the NFS environment to become the primary storage. A typical Celerra storage configuration is shown in [Figure 138 on page 292](#).

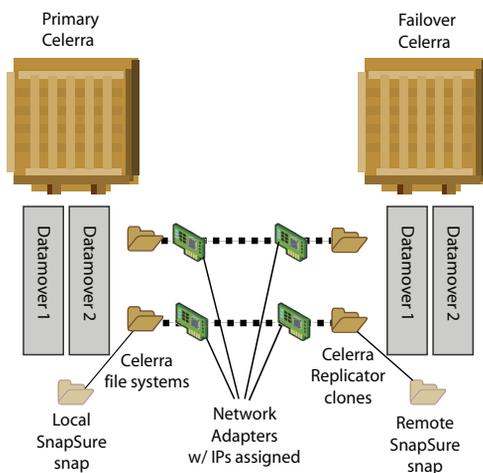


Figure 138 Celerra storage failover configuration

In this configuration, there are two Celerras each containing two datamovers and file systems reside on the datamovers. Information is transferred via network adapters between the Celerras. The NFS on the primary Celerra is a read write file system and is considered active. The NFS on the failover Celerra is read-only and considered passive.

Replication Manager for Celerra or VNX NFS allows customers to choose the target Celerra or VNX where the replica will be created and Replication Manager controls which replication technologies are used to create a replica on that Celerra or VNX based on the location of the data when the job runs.

Note: Replication Manager only supports application sets with a one-to-one relationship between two Celerras or VNXs. You cannot support configurations that include more than one target Celerra or VNX.

Each network adaptor is assigned an IP address so that it can be referenced across a network. This configuration allows for storage failover to the related NFS through normal Celerra procedures.

Configuring application sets for Celerra NFS

Replication Manager can support Celerra NFS environments like the one illustrated above. Before creating an application set, it is important to complete all the necessary steps to prepare the NFS on the Celerra, export the NFS to the production host, and (in the case of remote replication and failover between two Celerras) establish a replication session between the production and replica storage. For more information on these steps, refer to the *EMC Replication Manager Administrator's Guide* and the Celerra documentation.

On the Application Set Name and Objects wizard panel, there are two types of file systems available for Celerra environments (local file systems and network file systems). Refer to [Figure 139 on page 294](#) for an illustration of the wizard panel.

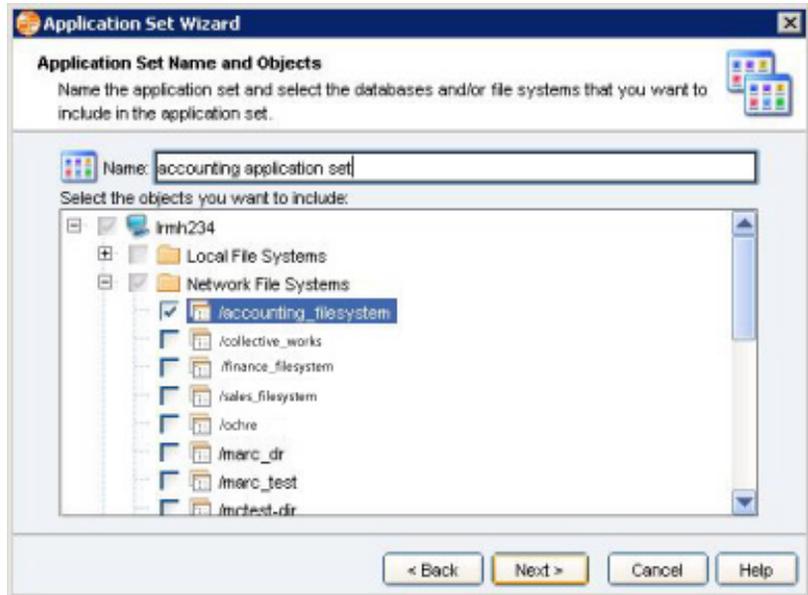


Figure 139 Celerra Application Set Name and Objects wizard panel

Note: Do not attempt to mix local file systems and network file systems in the same application set. This causes the replication to fail. Also, when choosing a network file system, make sure you are choosing an NFS that resides on supported storage. The list of network file systems includes all network file systems exported to the host.

Configuring jobs for Celerra network file systems

When creating a job for an application set that includes Celerra network file systems, the job wizard allows you to choose the **Snap Destination (Celerra)** and uses that information to determine if Replication Manager should perform a local or remote replication. Also, Replication Manager offers an advanced setting that allows you to restrict the job so that it only runs if the source and destination reside on the same Celerra, referred to as an in-frame replication.

Configuring jobs to replicate Celerra network file systems requires some planning. The following scenarios describe two possible configurations for NFS replicas.

Configuring NFS jobs to run in-frame only

Jobs can be configured to run only when the source file system and the target file system are in-frame, in other words only when both reside on the same Celerra or VNX.

To configure NFS jobs exclusively for “in-frame” replication, choose the **Advanced** replication options and select the checkbox **Only run job if in-frame replication**. Refer to [Figure 140 on page 295](#) for a view of that screen.

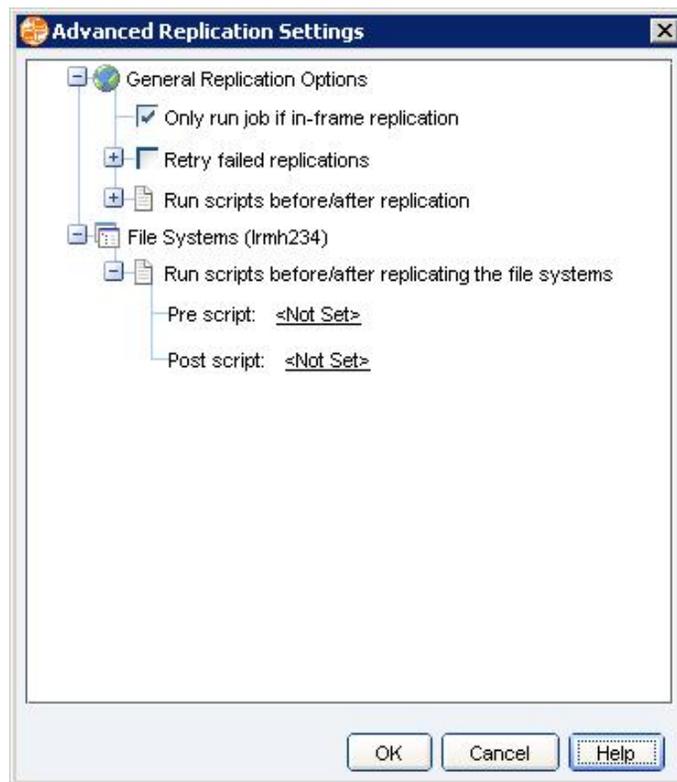


Figure 140 Advanced Replication Settings (Job Wizard)

Typically, two jobs are configured, one on the Primary Celerra and another on the Failover Celerra. This ensures that a job always runs on the active Celerra or VNX, creating SnapSure snaps of data on whichever Celerra or VNX is currently active. Refer to [Figure 141 on page 296](#) for a graphical representation.

Notice that Job 1 on the left is creating “in-frame” replicas on the Primary Celerra. Job 2 is not running because the source is on a different Celerra from the target for Job 2. Also note that if a failover occurs, Job 2 would create replicas and Job 1 would not.

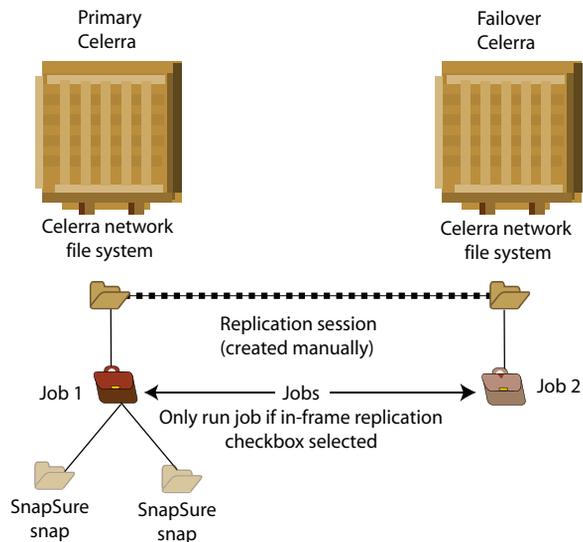


Figure 141 Jobs configured with “Only run if in-frame replication” checkbox selected

Configuring NFS jobs to run whether in-frame or not

Alternatively, jobs can be configured to run regardless of the location of the source. To configure a job to run whether the source and target are in-frame or not, choose the **Advanced** replication options and clear the checkbox **Only run job if in-frame replication**. Refer to [Figure 140 on page 295](#) for a view of the Advanced Replication Options panel.

Notice that the job depicted in [Figure 142 on page 297](#) runs even though the target is configured on the Failover Celerra, which is currently passive. If a failover occurs in this scenario, the job continues to run, but instead of using Celerra Replicator to create a replica across the WAN, Replication Manager now uses Celerra SnapSure to create a replica of the active file system located on the new primary Celerra, which was the failover Celerra prior to the failover.

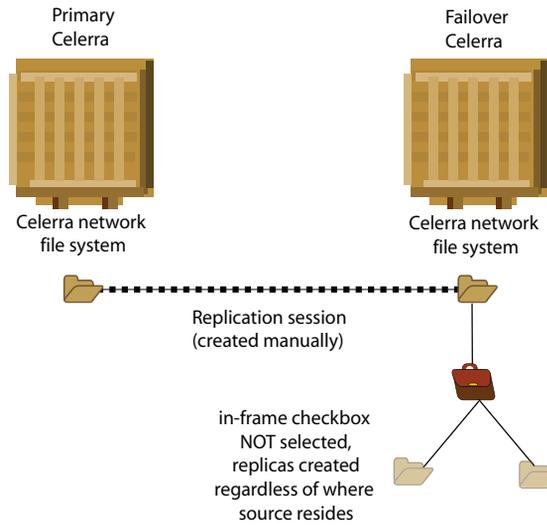


Figure 142 Jobs configured with “Only run if in-frame replication” checkbox cleared

Availability of RecoverPoint replicas

When data transfer on a RecoverPoint consistency group has started and an application set and job has been created for the consistency group, an application-consistent (specific-point-in-time) replica can be created.

Furthermore, when at least one application-consistent replica is created for the application set, a crash-consistent (any-point-in-time) replica can be created for any time since the start of the RecoverPoint protection window for the consistency group; this includes times before the creation of the application-consistent replica. [Figure 143 on page 299](#) illustrates this.

Any subsequent crash-consistent replicas created for the application set ensure the continued ability to mount RecoverPoint bookmark replicas within this period (the original application-consistent replica need not be present).

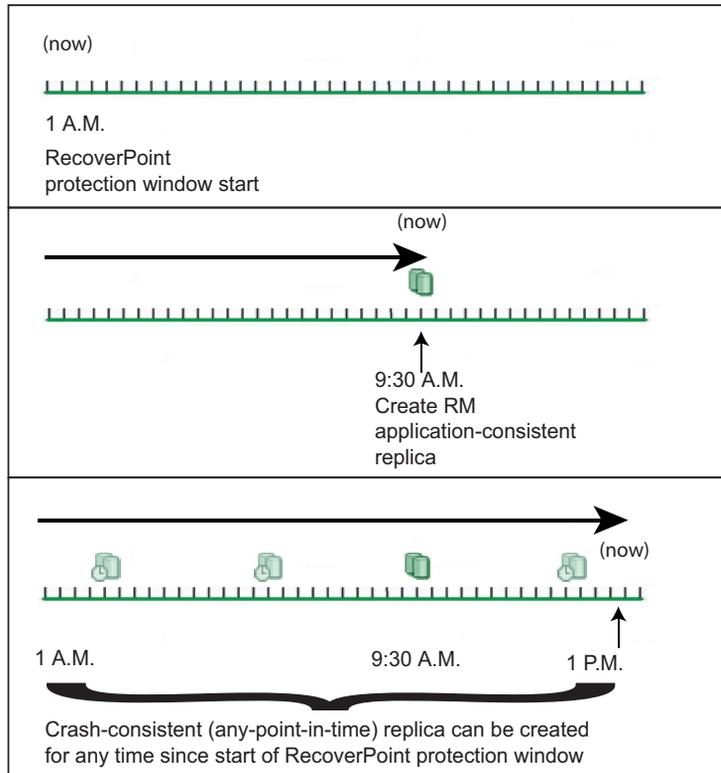


Figure 143 Availability of RecoverPoint replicas

Controlling replica expiration and deletion

Any role except operator can delete a replica at any time to reclaim the disk resources that are being used by that replica.

Deleting a replica

To delete a replica:

1. Expand **Application Sets** in the tree.
2. Select the particular application set that contains the replica you want to delete.
3. Right-click the replica you want to delete and select **Delete** from the context menu.
4. Click **OK** to confirm delete.



CAUTION

If you delete a replica, you will lose the data on that replica. Also, if a TimeFinder/Snap of a TimeFinder/Clone or TimeFinder/Mirror replica exists, deleting the source Clone or Mirror replica expires and deletes all the associated Snaps.

Enable or disable replica expiration

To enable or disable replica expiration:

1. Expand **Application Sets** in the tree.
2. Select the particular application set that contains the replica you want to modify.
3. Right-click the replica you want to modify and select **Properties**. The **General** tab allows you to postpone the delete operation of the replica as necessary.
4. Click **OK** to modify the date when the replica will be deleted as you have specified.

Enabling and disabling expiration is not available for RecoverPoint replicas.

Disabling replica expiration

To suspend automatic expiration (prevent a replica from being automatically deleted) for all current and subsequent replicas created using an application set:

1. Expand **Application Sets** in the tree.
2. Right-click the application set for which you want to suspend automatic expirations.

Note: Alternatively, you can right-click a specific replica in the content panel to disable expirations for that specific replica only.

3. Select **Disable All Expirations** from the shortcut menu.
4. Click **Finish** to complete the operation.

Note: Disabling Replica Expirations also prevents expiration of subsequent replicas associated with that application set (until you resume suspended expirations).

A replica for which Expirations have been disabled cannot be deleted even if its expiration time is reached or a replica rotation determines that it should be deleted.

Note: When a replica expiration is disabled, it can prevent replica rotations from occurring as planned if that replica is part of a rotation. This situation could also cause newer replicas in the rotation to be deleted “prematurely.”

Enabling replica expirations

To resume replica expirations (allow replicas to be automatically deleted) for one or more replicas created by an application set:

1. Expand **Application Sets** in the tree.
2. Right-click the application set for which you want to resume automatic expiration.

Note: Alternatively, right-click a specific replica in the content panel to enable expirations for that specific replica only.

3. Select **Enable All Expirations** from the shortcut menu.
4. Click **Finish** to complete the operation.

Automatic expiration resumes according to the retention periods or rotations set in the related job.

Note: You can modify the expiration time of a time-based replica under the replica's properties on the **General** tab.

A replica that has expirations enabled may be deleted at any time based on the rotation period or replica interval that you have determined. However, if you want to ensure that the replica is deleted immediately, you should explicitly delete it. For more information, refer to [“Deleting a replica” on page 300](#).

Modifying replica rotation

To modify a replica rotation to delete replicas based on a maximum number of replicas:

1. Click **Jobs** in the tree. Individual jobs appear in the content panel.
2. Right-click the job associated with the replica you want to modify and select **Properties**.
3. Click the **Replication** tab.
4. Select the **Limit replica count to** option and enter the desired number of replicas for that rotation.
5. Click **Finish** to complete the operation.

Note: For specific information about each wizard panel, click **Help**.

Rotating replicas

Replication Manager can create replicas using either timed expirations or rotations. Timed expirations are designed to keep the replica for a specified length of time, after which the system automatically deletes the replica. A replica rotation is not based on a time frame, rather it is based on reaching a maximum number of replicas. A replica rotation combined with a schedule can ensure that a certain number of replicas remain in existence without exceeding a pre-specified number of replicas. Rotational sets can be an effective means of controlling disk usage.

This section describes what a replica rotation is and how to create one. When you define a job, you can choose to use timed expirations or replica rotations. Refer to the wizard panel shown in [Figure 144](#) on [page 303](#).

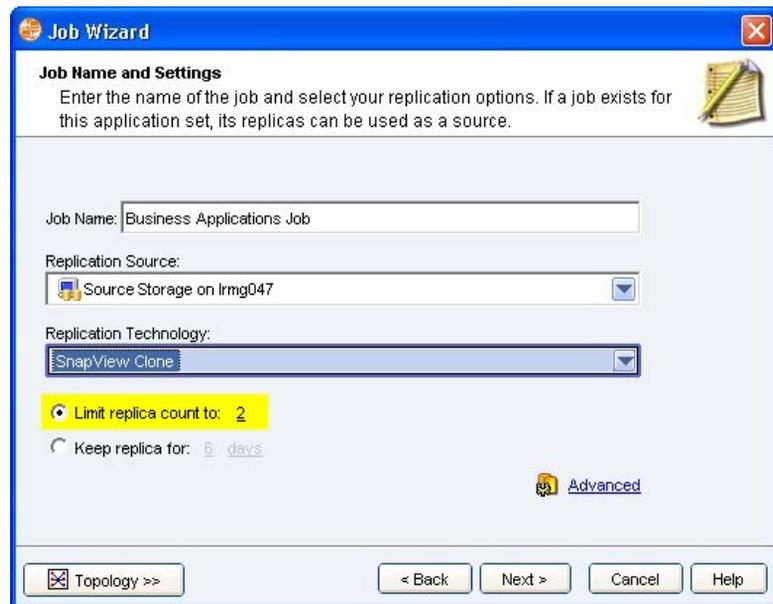


Figure 144 Creating a job for a replica rotation

In [Figure 144](#) on [page 303](#), the user has specified a maximum of two replicas. That means that when a schedule runs to create a replica from the job, Replication Manager checks to see how many replicas of the application set, that were created using the same job, currently exist. If the number is currently two or more, then Replication

Manager deletes the oldest replica that is not currently protected from expiration. If the replica to be deleted is mounted and is part of the rotation, it will be automatically unmounted then deleted. If the replicas are not part of the rotation, the job fails.

Note: In this example, between the time that the oldest replica is deleted and the time that the new replica has been created, there are a maximum of two replicas in existence.

If Replication Manager is unable to delete enough replicas to bring the replica count down to two replicas, then the job fails. Therefore, you can use replica rotations to limit the amount of storage that Replication Manager uses to create replicas of a single application set.

Note: If the time based retention policy is set such that target devices are not freed up frequently, the target devices can be exhausted and unavailable for fresh replication. To avoid this, be sure you have enough target devices when using a time based rotation policy. Another way to safeguard against this, is to set a storage pool, dedicate it to a job and, based on job schedule/devices in the application set, adequate devices are added to it.

For information on how to prevent a replica from being deleted when conditions dictate that it should be deleted, refer to [“Disabling replica expiration” on page 301](#) or [“Enabling replica expirations” on page 301](#).

Rotation is not supported for RecoverPoint replicas.

Setting retention periods

ERM Administrators and Database Administrators should plan appropriate replication schedules and retention periods to balance:

- ◆ Needed replica capabilities (for example, decision support, or data protection)
- ◆ Duration of the replica
- ◆ Available replica storage space

Any role except Operator can set retention periods for jobs so that replicas automatically delete on a reasonable schedule. Each schedule entry can be set to run at various intervals.

For example, suppose your company wants to replicate a certain critical database every hour and retain at least the two most recent replicas at all times. You could set up a schedule entry to replicate every hour with a retention period of three hours. With such a schedule in place, there are always two replicas in existence. If you set the duration of each replica to only two hours, there would be a window during which the replication was taking place when only one replica was available.

Alternatively, you could set the rotation to three and set up a schedule to run a replication every hour.

Setting retention periods is not supported for RecoverPoint replicas.

Setting retention periods for replicas

To set retention periods for all replicas created by a certain job:

Note: Setting the retention period as described here does not change the retention period for existing replicas, only for subsequent replicas created after performing the following steps.

1. Click **Jobs** in the tree. Individual jobs are displayed in the content panel.
2. Right-click the job associated with the replica you want to modify and select **Properties**.
3. Click the **Replication** tab, and select **Keep replica for:** Refer to [Figure 145 on page 306](#).

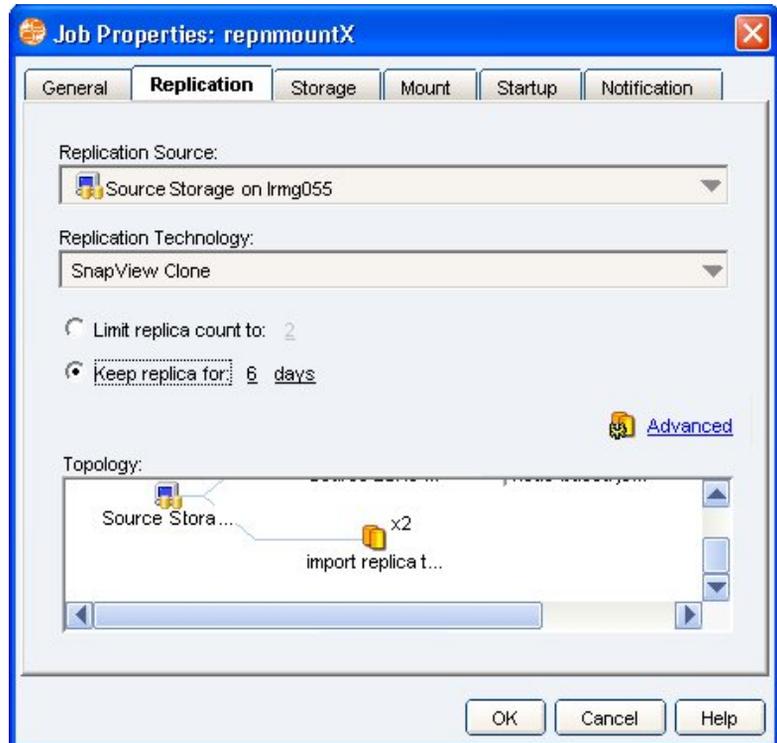


Figure 145 Setting replica retention period

4. Modify the retention period of the job.
5. Click the **Startup** tab and adjust the scheduled time so that the replication occurs at the desired time.
6. Click **OK** to complete the operation.

For the most part, you can choose a replica retention time and the expiration occurs as scheduled. If Replication Manager creates a replica with a retention period of *n* months, the expiration occurs on the same day of each month.

If your job uses retention periods to define when the replica expires then when the job runs, no other replica can be mounted to the target mountpoint specified in that job. If an existing replica is mounted there, the job will fail and the old replica will remain mounted. This differs from jobs that use rotations. If a rotational job encounters

another replica mounted to the target mountpoint, it automatically unmounts that replica and expires it before creating and mounting the new job.

Establishing and reestablishing mirrors

Mirror limitations on Symmetrix

Mirror synchronization functionality differs depending on the storage array(s) you are using. The following sections describe mirror capabilities for each storage array.

TimeFinder Multi-BCV functionality allows you to incrementally establish up to 16 mirrors simultaneously with a single standard device. However, you can create more than 16 replicas of the same standard. You are limited only by the available storage in the Symmetrix array. If a mirror is already associated with a standard device, a short reestablishing process is required to regain synchronization.

TimeFinder/Clone functionality is controlled via copy sessions, which pair the source and target devices. Copy sessions on a Symmetrix array are limited to 16 sessions per source device.

Mirror limitations on CLARiiON

CLARiiON storage arrays can have up to eight consecutive replicas of CLARiiON-based LUNs at any one time. Any attempt to create more than eight replicas on CLARiiON storage results in a failed replication attempt.

Understanding Replication Manager scripting

Replication Manager offers several scripting options to allow you to control what happens before, during, and after the replication process. This section describes:

- ◆ What scripting options are available to you
- ◆ Why you should choose a particular scripting option
- ◆ Where you can get more information about these scripting options

Scripting options

Replication Manager includes the following scripting options:

- ◆ **Pre-replication scripts** — Runs on the production host before Replication Manager creates the replica. This script is used to prepare the application(s) on the production host for replication processing.

You can specify parameters for the pre-replication script. For more information on the syntax and usage of these parameters, see online help from the Advanced Replication Settings dialog.

- ◆ **post-replication scripts** — Runs on the production host after Replication Manager has finished creating the replica. This script is typically used to perform tasks that return the production applications to normal processing.

You can specify parameters for the post-replication script. For more information on the syntax and usage of these parameters, see online help from the Advanced Replication Settings dialog.

- ◆ **Postmount scripts** — Runs on the mount host after Replication Manager has successfully mounted the replica to the mount host. This type of script is typically used to perform some special processing on the mount host related to the reason for mounting the replica.
- ◆ **Backup scripts** — Similar to mount scripts, however, a backup script allows you to specify a separate backup host on which the script can run. For example, if you have a dedicated backup host, separate from the mount host and you want to run a third-party backup process on that host, you would use a backup script for that.

- ◆ **Callout scripts** — Runs at specific times during the replication, mount or restore processing. These scripts give you more precise control over when your scripts run. Special callout scripts can be defined for federated jobs as well as non-federated jobs.

You can use these scripting options during replication, mount, and recovery processing to customize your results. Scripts can monitor progress, modify the environment, or carry out additional processing.

Additional references

Pre- and post-replication scripts

For further reading on the scripting options listed above, consider these suggestions.

For more information on pre- and post-replication scripts by application, see the following sections:

- ◆ For Oracle scripts, refer to [“Using pre- and post-replication Oracle scripts” on page 389](#).
- ◆ For UDB scripts, refer to [“Using pre- and post-replication UDB scripts” on page 424](#).
- ◆ For SQL Server scripts, refer to [“Using pre- and post-replication SQL Server scripts” on page 483](#).
- ◆ For Exchange scripts, refer to [“Using pre- and post-replication Exchange scripts” on page 555](#).

Mount and backup scripts

For more information on using mount and backup scripts, refer to [“Using mount and backup scripts” on page 186](#).

Callout scripts

For more information about how to use Replication Manager callout scripts, refer to [“Using application callout scripts” on page 235](#).

Acknowledging a failed replica

In the Replication Manager Console's content panel, a failed replica is marked with a flag. To acknowledge the failed replica and remove the flag, right-click the replica and select Acknowledge from the context menu.

If you want to display the flag on a replica again, right-click the replica and select Unacknowledge from the context menu.

To acknowledge all failed replicas in an application set, right-click the application set and select Acknowledge Application Set from the context menu.

This appendix includes the following sections that cover these subjects:

- ◆ Configuring Oracle for Replication Manager 312
- ◆ Configuring SAP for Replication Manager 328
- ◆ Configuring ASM RAC mount to ASM RAC..... 334
- ◆ Oracle configuration detailed concepts 337
- ◆ Understanding Oracle application sets and jobs..... 349
- ◆ Mounting Oracle replicas..... 361
- ◆ Restoring Oracle replicas 379
- ◆ Using pre- and post-replication Oracle scripts 389
- ◆ Using the root user to perform Oracle operations 395
- ◆ Storage Foundation for Real Application Clusters 397
- ◆ Oracle troubleshooting..... 402

Configuring Oracle for Replication Manager

Replication Manager interoperates with the following Oracle environments:

- ◆ Standalone Oracle databases
- ◆ Oracle using Real Application Cluster (RAC)
- ◆ Oracle using Automatic Storage Management (ASM)

Replication Manager protects these environments by creating and managing application sets that contain one or more Oracle databases or specific Oracle tablespaces. These application sets are used to create replicas that can be mounted or later restored to the original production environment.

Replication Manager does not truncate archive logs or manage the Flash Recovery Area free space. To protect these parts of your Oracle environment, configure the appropriate retention policy with backup tools such as RMAN to avoid filling up the archive log directory location. Consult Oracle documentation for additional details.

Before using Replication Manager, verify that Oracle is configured properly for your environment. The next few sections document Oracle configuration requirements.

Replication Manager supports Oracle in a wide variety of platforms as shown in [Table 23 on page 314](#) which outlines the supported platforms briefly. For complete support information, refer to the *EMC Replication Manager Support Matrix*. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

Oracle configuration notes

When considering the Replication Manager Oracle support table in the *EMC E-Lab Interoperability Navigator Replication Manager Support Matrix*, review these configuration notes:

UNIX

On UNIX, Replication Manager:

- ◆ Supports Oracle stand-alone databases that reside on file systems, raw devices and ASM.
- ◆ Supports Oracle RAC databases on raw devices, ASM devices and Veritas cluster file systems (Solaris-only).

- ◆ Supports Oracle ASM on raw devices only. Use of custom character special files created with tools such as mknod or mksf is not supported on AIX and HP platforms for RM versions greater than 5.4.1. For RM versions lower than 5.4, only external redundancy level disk mirroring is supported.
- ◆ 5.2.x, Replication Manager does not support a separate ORACLE_HOME for the ASM and database instances for non-Linux platforms.

LINUX On Linux Replication Manager:

- ◆ Supports Oracle stand-alone databases that reside on file systems, raw devices, ASM, VNX File and Celerra-based network file systems (NFS).
- ◆ Supports Oracle RAC databases on raw devices, ASM devices VNX File and Celerra-based network file systems (NFS). Mixing of RAW and ASMLib is not supported so CRS disks must also be on ASMLib..
- ◆ Supports Oracle ASM on raw devices, ASMLib devices. ASMLib devices are recommended. Use of block devices or character special files other than raw interface devices are not supported. Mixing of RAW and ASMLib is not supported. For RM versions lower than 5.4.0, only external redundancy level (ASM disk mirroring) is supported.. When using ASMLib devices, the disks must be presented to the ASM instance using the "ORCL:" diskstrings, as the use of the "/dev/oracleasm" disk path is not supported.
- ◆ Supports mounting of ASM replicas from production RAC to a target RAC if the database exists on ASMLib volumes. Mixing of RAW and ASMLib is not supported so CRS disks must also be on ASMLib.
- ◆ For Oracle version 11gR2 and onwards, RM support Oracle ASM on ASMLib devices only. ASM on raw devices is not supported.
- ◆ For RHEL 6. 0, Oracle Linux 6.0 and SUSE SLES 11 onwards, RM does not support databases residing on raw devices.
- ◆ Replication Manager does not support the use of Oracle Clustered File System (OCFS).

Windows On Windows Replication Manager:

- ◆ Supports Oracle stand-alone databases that reside on file systems only.

- ◆ Does not support ASM and RAC configurations.

Table 23 Replication Manager platform support for Oracle

Oracle configuration and Replication Manager features	Windows	AIX	HP-UX	Solaris	Linux	Notes
Standalone database (general)	Yes	Yes	Yes	Yes	Yes	
Standalone database with ASM	No	Yes	Yes	Yes	Yes	
Standalone database on network file systems	No	Yes	Yes	Yes	Yes	
RAC database	No	Yes	Yes	Yes	Yes	
RAC database with ASM	No	Yes	Yes	Yes	Yes	
RAC database on network file systems	No	Yes	Yes	Yes	Yes	
Catalog with RMAN	No	No	Yes	Yes	Yes*	
Mount to a RAC	No	No	No	No	Yes*	*Only for ASM with ASMlib
SAP BRbackup integration	No	No	Yes	No	Yes	*Not with ASM storage
ASM extended features (ASMCFS and DVM)	No	No	No	No	No	

General Oracle configuration for Replication Manager

The following checklist can help you configure Oracle to work with Replication Manager.

These steps apply to all environments:

- ❑ If the production Oracle server is running Multi-Threaded Server (MTS), verify that the `tnsnames.ora` file contains at least one entry with a dedicated connection. This dedicated connection should be used when defining an Oracle application set. Consult Oracle documentation on MTS for instructions.
- ❑ Replication Manager communicates with Oracle through the TCP/IP listener. The listener on the production Oracle server must be configured according to Oracle standards.

- ❑ If you plan to create online replicas, put the database in Archive Log mode. To check whether the database is in archive log mode, use the following SQLPlus command:

```
archive log list;
```

- ❑ If you plan to create replicas without using consistent split, verify that controlfiles and redo logs reside on devices separate from the datafiles devices.
- ❑ If you plan to create replicas using consistent-split technology, controlfiles and redo logs have to be located on a supported storage array.
- ❑ If you are creating replicas of the flashback recovery area and/or archive logs, those components must be located on a supported storage array.
- ❑ Verify that the `tnsalias` string is valid for use with Replication Manager and that the listener is configured properly. To do this, use the `tnsping` command as follows:

```
tnsping myalias
```

If the system responds to `tnsping`, it is working properly; if not, refer to [“Configuring the Oracle TCP/IP listener” on page 338](#) for more information about how to configure the listener.

- ❑ In case of Oracle 11gR2, the ASM instance user should have the SYSASM and SYSDBA privileges. Also, the ASM instance user ‘sys’, is not a requirement; any user with SYSASM privilege is supported for authenticating to the ASM instance.
- ❑ Set the parameter REMOTE_LOGIN_PASSWORDFILE to EXCLUSIVE and create a password file in the standard location:

```
$ORACLE_HOME/dbs/orapw<sid>
```

Note: In Windows environments, the password file is created when the NT service for the instance is created.

Configuring Oracle for Windows

The following checklist items can help you configure Oracle to work with Replication Manager in a Windows environment.

These steps apply to Windows environments:

- ❑ If you are integrating Replication Manager with Oracle in a Windows environment, and you are experiencing difficulties connecting to Oracle, EMC recommends that you try connecting without SQL*Net Authentication. To do that edit the SQLNET.ora file and add the following lines:

```
SQLNET.AUTHENTICATION_SERVICES = (none)
```

```
SQLNET.AUTHENTICATION = (none)
```

On Windows hosts, Replication Manager requires that you install the Oracle binaries on the mount host even if no Oracle operations are performed during the mount. In other words, Oracle binaries are required on the Windows mount host even when you choose the **Do not perform database operations** or the **Generate database recovery scripts** options. Oracle binaries are not required for the **Do not perform database operations** on UNIX mount hosts.

Configuring Oracle for UNIX/Linux

The following checklist items can help you configure Oracle to work with Replication Manager in a UNIX or Linux environments.

These steps apply to UNIX or Linux environments:

- ❑ It is recommended that the Oracle SID be listed in the oratab file. If the SID is in oratab, Replication Manager detects it and presents it as a selectable object when you create an application set. Additionally, its corresponding oracle home will be defaulted automatically. This recommendation applies specifically to standalone databases.

You can find the `oratab` file in `/var/opt/oracle` (Solaris) or `/etc` (HP-UX, Linux, AIX).

For a SID that is not listed in `oratab`, nothing is defaulted, and the SID is not presented as a selectable object; you need to use the Add Instance option when creating an application set.

See [“Adding an undiscovered Oracle instance \(UNIX only\)” on page 339](#) for more information.

- ❑ Be aware that if you choose to allow Replication Manager to perform Oracle operations as the root user (not recommended), extra configuration steps are required. Refer to [“Using the root user to perform Oracle operations” on page 395](#) for more information.

- ❑ With Red Hat Linux 5.6, when creating a job, the following error is received:

```
026380 ERROR: The file system
/secondarypath/oradata/a/b/systema.dbf is not located
on a supported storage technology, please ensure the
application set is properly configured to include only
storage that is supported by Replication Manager.
```

This error occurs because /secondarypath is a symbolic link to /primarypath and the /etc/fstab does not contain an entry for /primarypath.

Workaround: Please ensure there is an entry in the /etc/fstab file for the /primarypath. You may need to issue the mount -a command to cause all filesystems listed in the /etc/fstab directory to be remounted and recreate the application set

Configuring Oracle ASM for Replication Manager

Replication Manager supports ASM disk groups but does not support ASM Dynamic Volume Manager (ADVM) and ASM Clustered File System (ACFS). Replication Manager cannot be deployed in environments where these features are leveraged (even if the disk groups contained in the current application set do not use them).

The following checklist items can help you configure Oracle ASM to work with Replication Manager.

These steps apply to Oracle ASM environments only:

Note: For more information on Oracle ASM configurations with Replication Manager, refer to [“Using Oracle ASM with Replication Manager” on page 344](#).

- ❑ The ASM instance user ‘sys’ is the only user account currently supported for authenticating to the ASM instance for Oracle releases less than Oracle 11gR2.

For Oracle 11gR2, any ASM instance user is supported for authentication, provided the user has SYSASM and SYSDBA privileges.

- ❑ Datafiles, online redo logs, controlfiles, and sp-files can be located on ASM disk groups composed of raw volumes or ASMLib volumes (Linux).

Note: The practice of building Oracle databases on Linux raw volumes (/dev/raw/rawX) has been deprecated in newer releases of Oracle. Replication Manager supports this legacy feature but recommends transitioning to newer methods in Linux environments.

Using the ASMLib driver in the ASM environment implies that the ASM disks provided to the ASM instance follows default naming conventions defined in ASMLib documentation. For instance, if an ASMLib volume is given the label "DISK1", Replication Manager recognizes and maps DISK1 as long as it is presented to ASM as ORCL:DISK1. Direct path references to the underlying /dev/oracleasm/disks/DISK1 path are not supported by Replication Manager.

Note: ASMLib volumes and EMC PowerPath can be used together. For example ORCL:DISK1 can point to /dev/emcpowerd1. Again ORCL:DISK1 is what should be presented to the ASM instance when creating/managing ASM diskgroups, in order for Replication Manager to translate the paths to the physical disks correctly.

- ❑ Replication Manager automatically renames ASMLib volumes as necessary to prevent them from conflicting with production volumes, and then restores the original names upon restore.
- ❑ If an existing ASM database is recreated using a create controlfile command and the existing data or log files are specified using short-form Numeric filenames, Replication Manager will not be able to replicate that ASM database. Numeric filenames in the v\$datafile view are unsupported. Only fully-qualified ASM filenames are supported.
- ❑ On Solaris, HP-UX, and Linux platforms, Replication Manager supports the use of two separate ORACLE_HOME directories, one for ASM and one for the Oracle databases. The ASM and database installations can be owned by different operating system users belonging to different groups. The environment must list the ASM instance explicitly in the /etc/oratab file.
- ❑ The kfed binary must be built and available on the \$ORACLE_HOME/bin directory on any host where the ASM disk group rename functionality will be used.

To create this binary, type the following commands on a command line:

```
cd $ORACLE_HOME/rdbms/lib
gmake -f ins_rdbms.mk ikfed
```

Note: kfed binary need not be built for Oracle 11gR2.

For Oracle 11gR2 on Linux, Replication Manager uses *renamedg*, a tool provided by Oracle for renaming a diskgroup. This tool requires patch 9316059 to be installed as a prerequisite for diskgroup rename functionality. Refer to the readme file provided with 9316059 for details on how to apply the patch.

- ❑ If you intend to replicate the Oracle archive log directory, configure it to reside on its own ASM disk group.
- ❑ Verify that the ASM disk groups that are replicated use external redundancy. High and normal redundancy levels are not supported by Replication Manager. Refer to [“Oracle ASM redundancy level restrictions” on page 347](#) for more information.
- ❑ On Linux and Solaris platforms, character special files created in arbitrary locations by the **mknod** or **mksf** command as raw interfaces to Oracle ASM disks are not supported by Replication Manager.

Note: However, this functionality is supported for AIX and HP-UX.

The following example, `c0t2d0s4`, is a supported character special file:

```
# ls -lL /dev/rdisk
crw-r----- 1 root sys 32, 20 Feb 24 09:14 c0t2d0s4
```

In the following, `disk1` is an example of a character special file that is not supported:

```
# cd /mydisks
# mnod disk1 c 32 20
# ls -l /mydisks
crw-r--r-- 1 root other 32, 20 May 7 09:50 disk1
```

Configuring Oracle RAC for Replication Manager

The following checklist items help you configure Replication Manager to replicate an Oracle database that is part of a Real Application Cluster (RAC) environment.

Users that have installed Replication Manager with Oracle ASM and RAC should consider the following:

- ❑ The default ASM instance name is automatically entered in the relevant configuration panels of the Replication Manager Console where a connection to the ASM instance is required. For example, during application set creation, the production ASM instance to connect to is defaulted for the RAC node where the application set is being defined. For node 1, this would typically be +ASM1.

For 11gR2, a connection to the mount host ASM instance is also required, when mounting a replica. The same defaulting mechanism is used there also.

- ❑ ASMLib volumes are required if you plan to mount an ASM RAC replica to a separate ASM RAC cluster. Refer to [“Configuring ASM RAC mount to ASM RAC” on page 334](#) for more specifics on configuring RAC to RAC environments.
- ❑ Replication Manager automatically renames ASMLib volumes as necessary to prevent them from conflicting with production volumes, and then restores the original names upon restore.
- ❑ For 11gR2 RAC configurations, Replication Manager supports Single Client Access Name (SCAN) feature for use with Oracle RAC, however the SCAN IP cannot be used to register a host with Replication Manager. Instead, choose the hostname, public IP or virtual IP to register this kind of host with the Replication Manager server. Additionally, complete the steps described in [“Configuring the Oracle TCP/IP listener” on page 338](#).

Configuring Replication Manager for RAC awareness

Replication Manager can operate in 2 modes when run on a RAC configuration

1. **RAC agnostic** in which Replication Manager only runs on a single node within the RAC. If the node is down, the operation fails.
2. **RAC aware** in which:
 - During replication, if the node used for application set creation is not accessible, Replication Manager runs the replication on another node in the RAC.

- During a restore operation, if the node on which the replica was created is not accessible, Replication Manager runs the restore on another node in the RAC.

Note: If you do not want to configure RAC awareness, you can register Replication Manager client using its physical IP. Replication Manager console does not provide a GUI option to set RAC awareness, it should be configured through the IP using which the Replication Manager client has been registered.

The following are the requirements for configuring RAC awareness with Replication Manager:

- ◆ Replication Manager client should be installed on all nodes of the RAC.
- ◆ Replication Manager client should be registered using
 - VCS Failover IP or with the IP resource configured in Oracle service group for the database or Oracle RAC VIP managed by VCS in case of SFRAC.
 - Oracle RAC VIP/ Failover/Cluster IP in case of other RACs
- ◆ The database user, i.e., SYSDBA privileged user credentials should be the same for all DB instances running in the RAC. All Oracle password files across all production RAC nodes should be in sync.
- ◆ RAC-aware feature will work only if the node failover occurs prior to running the job, that is, prior to start of replication or start of restore. If failover happens in the midst of running replication or restore, the operation fails.
- ◆ In Advanced Replication settings, select **SP file** in the **Copy parameter file to RM Server** option. The SPfile can be a global SPfile on shared location and not referencing the local SID (RAC instance) or the SPfile of a RAC instance running on the node selected for job creation. Replication Manager does not support use of init file in **Copy parameter file to RM Server** option.
- ◆ Replication Manager does not support OMF for database Creation of non-ASM databases.

Configuring Replication Manager to work with failover standalone database in a cluster

In addition to supporting RAC databases, Replication Manager can work with standalone databases that are configured to failover from node to node in the cluster. During replication, if the node that was used for application set creation is not accessible, Replication Manager runs the replication on another node in the cluster.

Similarly, during a restore, if the node on which the replica was created is not accessible, Replication Manager runs the restore on another node in the cluster. During appset creation, the oracle agent verifies if the selected instance is currently running on the node used for appset creation, else, an appropriate error is displayed.

The following are the requirements for using standalone database that fails over from node to node in the cluster with Replication Manager:

- ◆ Replication Manager client should be installed on all the nodes of the cluster.
- ◆ Replication Manager client should be registered using:
 - VCS failover IP or IP resource configured in Oracle service group for the database in case of a VCS cluster.
 - Failover/ Cluster VIP in case of other clusters.
- ◆ If failover happens in the midst of running replication or restore, the operation fails. Node failover should occur prior to running the job, that is, prior to start of replication or start of restore.
- ◆ The oratab file should have an entry for all possible SIDs that can run on the given node (passive and active instances).
- ◆ The tnsnames.ora files on all nodes should contain entries of all standalone instances.
- ◆ The following files should be accessible to all nodes of the cluster where the database can run:
 - Database init/spfile
 - Password file

Note: Create the files on a shared filesystem or manually copy the files to their corresponding location on all nodes (\$ORACLE_HOME/dbs).

Configuring Oracle 11g R2 RAC One Node

- ◆ The dump directory location specified in the init file must exist on all nodes in the cluster.

Replication Manager can be configured to use the Oracle 11g R2 RAC One Node feature on Linux. The requirements for using the feature are listed here:

- ❑ Replication Manager client should be installed on all the nodes of the RAC.
- ❑ Register the host using RAC VIP of the node that currently owns the RAC One Node database. SCAN IP cannot be used for registering the host.
- ❑ In case of a failover, the RAC VIP and the database should come up on the same RAC node. In case, the RAC VIP and the instance come up on different nodes of the cluster, the user needs to migrate the RAC VIP to the appropriate node using the **ifconfig** command. For example:
 - ❑ If Omotion is used to migrate the instance, on the node previously owning the RAC VIP, run the command **ifconfig eth0:<n> <RAC VIP> netmask <Netmask> down**
 - ❑ On the RAC node where the database has come up after a failover, run the command **ifconfig eth0:<n> <RAC VIP> netmask <Netmask> up**
 - ❑ Check the node presently owning the RAC VIP resource in the output of `crs_stat` and confirm that it is pointing to the correct node.
- ❑ If there is a failover during a replication or restore process, the operation fails. Node failover should occur prior to running the job, that is, prior to start of replication or start of restore.
- ❑ The `tnsnames.ora` file on all the nodes that are candidates for failover must contain an entry for all the possible RAC one node database instances.
For example, let us consider a two node RAC configuration with node names `node1` and `node2`, and the RAC One Node database name as `racone`. Let us assume that the database is initially running on `node1` as `racone_1`. In case a failover happens via Omotion to `node2`, the instance would be renamed to `racone_2`. However, in case `node1` goes down, the instance name remains as

racone_1 when it comes up on node2. Consequently, the tnsnames.ora file on both the nodes should have entries for both the database instances. Sample entries are given below

```
RACONE_1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = SCAN_IP) (PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = racone)
      (INSTANCE_NAME = racone_1)
    )
  )
```

```
RACONE_2 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = SCAN_IP) (PORT =
1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = racone)
      (INSTANCE_NAME = racone_2)
    )
  )
```

- ❑ In Advanced Replication settings, select **SP file** in the **Copy parameter file to RM Server** option. Replication Manager does not support use of init file in **Copy parameter file to RM Server** option.
- ❑ **Mount as Real Application Cluster** is not supported for RAC One Node database replicas.

Configuring Oracle Celerra NFS DR environment

Replication Manager can create replicas of both the production and disaster recovery sites of a Celerra NFS environment configured using Celerra Replicator to transfer data between the two sites. In order to work in this environment, there are some very specific replication requirements, outlined here:

- ❑ Celerra Replicator sessions between the two sites must be configured manually in advance. Refer to “Configuring Celerra network file system targets” section of the *EMC Replication Manager Administrator’s Guide* for more information on the setup steps.

- ❑ The Oracle instance on the disaster recovery site must mimic the source Oracle instance on the production site with regard to the following characteristics:
 - Same Oracle instance name in the application set
 - Same ORACLE_HOME
 - Same OS user and group user
 - Same init file (with some exceptions)

One difference that is allowed is that the Oracle instance on the disaster recovery site can be a standalone instance, even if the production site includes a RAC implementation.

- ❑ The hostname used in the Replication Manager application set for a Celerra NFS DR site should resolve to a virtual IP so that the same hostname (or virtual IP) can be used for both production and disaster recovery application sets. In this way, the application set can access the correct host regardless of the failover status of the DR environment.
- ❑ The path structure of the production and disaster recovery sites must match when setting up this kind of disaster recovery environment for Oracle.

One possible Oracle Celerra DR configuration is illustrated in [Figure 146 on page 326](#).

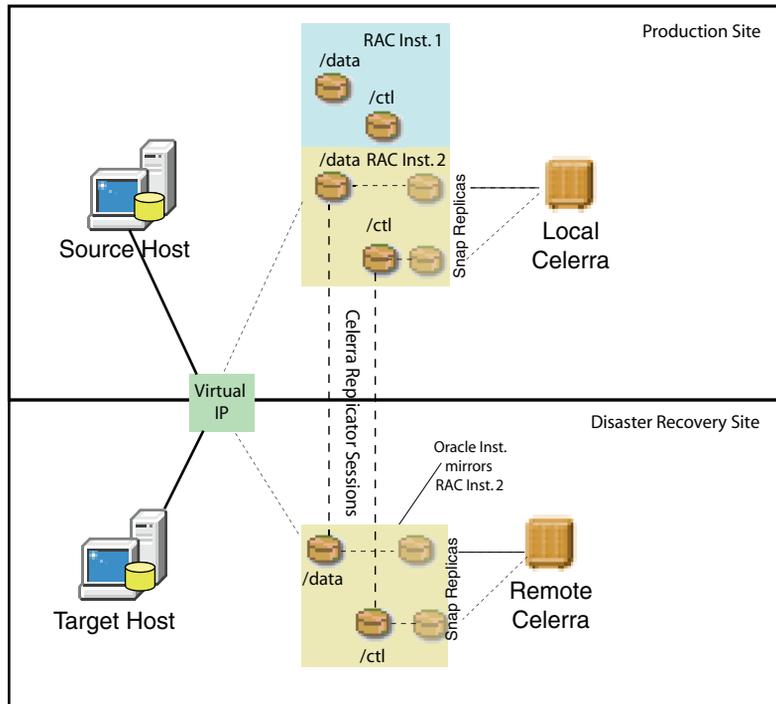


Figure 146 Oracle Celerra DR configuration with network file systems (RAC on production site; standalone on the disaster recovery site)

Note: Note that the Oracle instance on the disaster recovery site is not a RAC instance. You can choose to implement the disaster recovery site as a stand-alone instance that has similar characteristics to one node of the production RAC site if you so desire.

- ❑ Configure jobs to create replicas in one of the following ways:
 - Configure two jobs, one with the target Celerra on the production site and one with the target Celerra on the remote site. Select the checkbox under the Advanced link within General Replication Options entitled **Only run job if in-frame replication**. See [Figure 84 on page 130](#) for an illustration of this panel. Selecting this checkbox configures the job to run only when source and destination are on the same Celerra.

One job runs when both source and target are on the production Celerra, the other job runs when both source and target are on the disaster recovery Celerra.

- Configure only one job and clear the checkbox entitled **Only run job if in-frame replication**. This job creates local snaps prior to a failover and can run remote snaps after a failover if the original environment is still in existence.

Configuring SAP for Replication Manager

SAP (Systems, Applications and Products in Data Processing) provides enterprise software applications and support to businesses globally and is built on relational databases such as Oracle and DB2. Replication Manager provides an integrated solution for those with SAP-Oracle environments to backup the Oracle database.

Note: SAP with Replication Manager is currently only supported on Linux and HP-UX platforms.

BRbackup and SPLITINT

BRbackup is the SAP tool for backing up the SAP environment, including all Oracle database components (data files, control files, archive log files) as well as non-Oracle files and directories.

SPLITINT is an interface provided by SAP. In this context, Replication Manager provides a SPLITINT implementation to integrate with BRbackup. BRbackup will handle the database operations such as placing the database in hot backup mode and delegate the disk mirror splitting operations to SPLITINT (the Replication Manager job). After the disk mirror splitting operation is complete, be sure to take the database out of hot backup mode.

Overview

In order to facilitate the integration with BRbackup, these steps need to be followed:

1. You must set up a shared directory between the production host (with read access) and mount/backup host (with read write access). Both shared directories should have the same absolute path on both hosts. NFS is the recommended way to achieve this (although it is not required).
2. You must create a Replication Manager application set for the Oracle database that the SAP instance runs on and a job with the BRbackup compliant option.

- On the Replication Manager mount host, create a symbolic link from the splitint.sh script (located under <RM_install_directory>/client/bin) into the SAP BR*Tools directory, typically /sapmnt/<SID>/exe as "splitint".

For example:

```
l/bin/ln -s <RM_install_directory>/client/bin/splitint.sh /sapmnt/<SID>/exe/splitint
```

- Make sure the splitint executable has permissions set so that the user running brbackup can execute the command.
- You must copy the split_options parameter from the Replication Manager replication options panel to your BRbackup init file and your init file must specify a backup_type of "online_mirror".
- The BRbackup command will be run from the mount host and initiate the job created in [Step 2](#).

Configuration procedure

When you create a replica that includes SAP data, Replication Manager offers BRbackup compliant related options. To configure these options:

- From the Job Wizard, Job Name and Settings screen, click **Advanced**. The **Advanced Replication Settings** screen appears as shown in [Figure 147 on page 329](#).

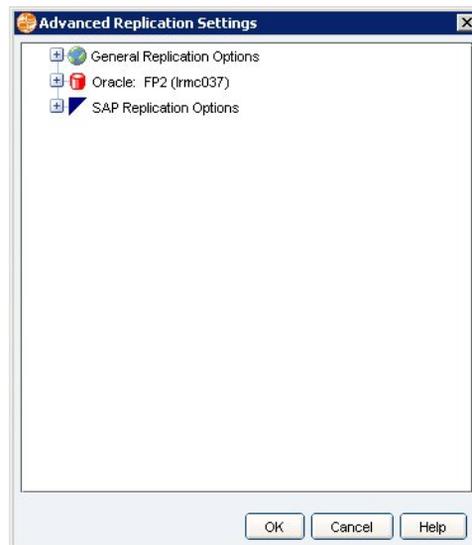


Figure 147 Advanced Replication Settings showing SAP Options

2. From the **Advanced Replication Settings** screen, expand the **SAP Replication Options** tree to view the specific SAP options:
 - **SAP BRbackup Compliant Job** — You must select this to enable BRbackup compliant job. Until you do so, the remaining SAP options are disabled and greyed out.
 - **BRbackup shared directory** — Allows you to specify the shared directory for BRbackup synchronization files between the production and mount host.
 - **Copy the SPLIT_OPTIONS parameter to your BRbackup init file** — Displays the text for you to copy and paste in your BRbackup init file. This parameter is what will link the BRbackup process to the Replication Manager job.
3. Select **SAP BRbackup Compliant Job** to enable BRbackup compliant replicas.

An informational message displays explaining that SAP BRbackup must be invoked to run BRbackup jobs, and that you will be required to choose a mount host (not the production host) to serve as your backup host. In addition, a red asterisk appears next to **BRbackup shared directory** indicating that this is a required field. Click **OK** to dismiss this Information dialog.

4. Click **<Not Set>** next to **BRbackup shared directory**. A dropdown menu and Browse button become active.
5. Enter the name of your shared directory or use the **Browse** button to select it. Once you enter this information, the Copy split_options parameter information becomes active, an informational text area displays the script commands and parameters needed for your init file, as shown in [Figure 148 on page 331](#). You must copy this text and paste it into your init file.

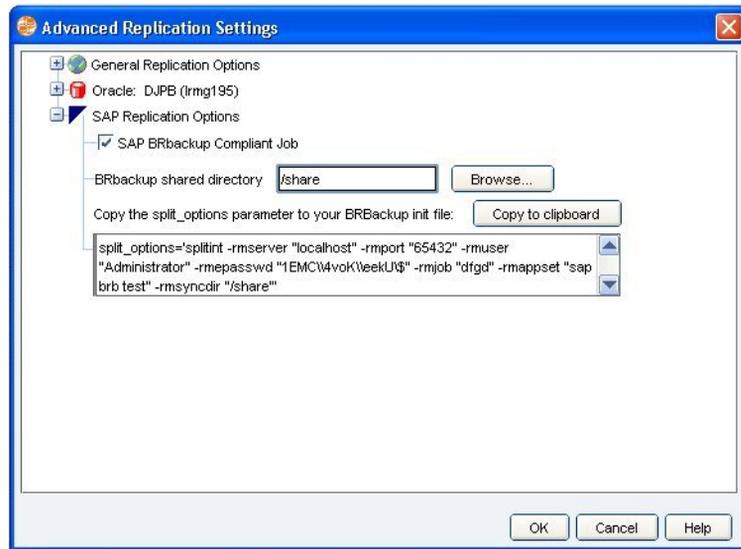


Figure 148 SAP Replication Options (expanded and selected)

6. After the text, “**Copy the split_options parameter to your BRbackup init file**”, click **Copy to clipboard**.

The Oracle consistency option is set to online without hotbackup mode since BRbackup handles the hotbackup mode commands, and hence Replication Manager's job avoids this to prevent conflicts.

Note: Alternatively, you could highlight the text, right-click and copy or use **Ctrl-C** to highlight and copy the text of the script.

Note: If the Replication Manager user password changed, the split_options need to be re-taken (copied to the clipboard) and pasted to the brbackup init file to reflect the password changes.

7. Paste the text stored in the clipboard into your init file.
8. Click **OK** in the Advanced Replication Settings dialog to return to the Job wizard.

As you progress through the Job wizard, note that certain replication, scheduling, and mount options are forced to certain values when using SAP BRbackup. These forced values are depicted in [Figure 149](#) on page 332.

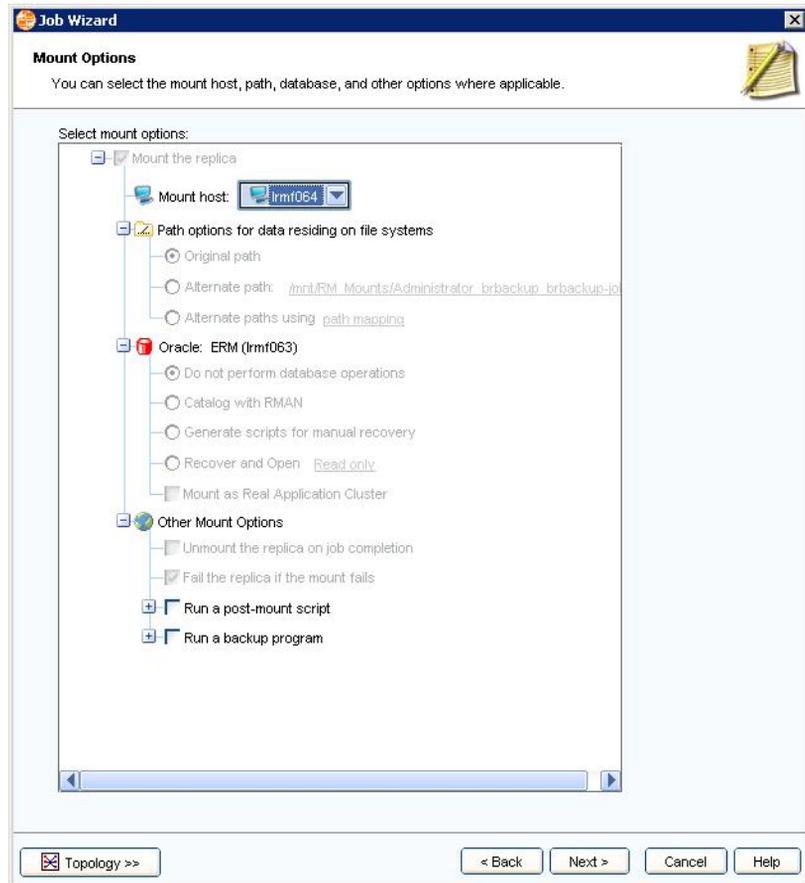


Figure 149 SAP Brbackup forced mount options

On the Mount Options panel, the following mount options are enforced:

- a. The act of mounting itself is enforced. The BRbackup compliant job should ensure that the resulting replica is mounted to the BRbackup backup host where it can run the backup.

- b. The path where the replica gets mounted needs to be the same as that of the production host because the BRbackup utility will look for the files to backup in the same path as the production database.
- c. The recovery option of the database is set to *Do not perform db operations*. This is because the replica being backed up by BRbackup needs to be in hotbackup mode and recovering the database would remove this and invalidate the copy being backed up.
- d. The option *Unmount the replica upon mount completion* is disabled because the job needs to leave the replica mounted in order for BRbackup to backup the replica.
- e. The option *Fail the replica if mount fails* is checked because the overall requirement of the BRbackup compliant job is to provide a mounted replica to the backup host. Therefore a mount failure should invalidate the replica.

Configuring ASM RAC mount to ASM RAC

Replication Manager allows you to mount an Oracle 10g, 11g, or 11gR2 ASM RAC replica to a standalone server or to another existing ASM RAC cluster. RAC to RAC mount is supported in the following environments only:

- ◆ ASM RAC on Red Hat Linux version 4.x or 5.x
- ◆ ASM RAC on Oracle Enterprise Linux version 4.x or 5.x
- ◆ ASM RAC on SUSE Linux version 10.x or 11
- ◆ ASM RAC (11gR2 onwards) on AIX version 6.1 or later

RAC to RAC mount is not supported for cross platform mounts. The platform where the production RAC instance resides must match exactly the platform of the mount environment.

This section outlines the configuration prerequisites when configuring the target ASM RAC cluster. The target cluster must be configured in advance in order to support mounts to a RAC cluster:

- ❑ The configuration of the cluster to which the replica is mounted must match the production environment, typically with the same number of nodes on the production and mount clusters.
- ❑ initfiles on the production system must contain global parameters for all instances.
- ❑ SSH passwordless authentication must be enabled between all nodes of the target ASM RAC cluster for the owner of the Oracle binaries, the owner of ASM instance and the root user.
- ❑ Oracle Clusterware should be setup in advance on all nodes of the alternate ASM RAC cluster. This includes configuration of the Oracle Cluster Registry (OCR) and voting disks.
- ❑ Any node of the target ASM RAC cluster should be designated as a Replication Manager mount host. When you choose the mount option **Mount as Real Application Cluster**, Replication Manager will mount to this host and propagate the mounted instance to all the nodes of the target ASM RAC cluster.
- ❑ **Applicable for Linux platforms:** ASMLib is required on all nodes of the target ASM RAC cluster.
- ❑ **Applicable for AIX platforms:**

- Prior to mount, all the replica disks must have the required permissions for ASM candidate disks on all the nodes of the target RAC. Validate this by logging into the ASM instance on each of the target RAC nodes and execute the command `select name from v$asm_disk`. The replica disks should be listed in the output.
- After every unmount, remove permissions from the unmounted replica disk so that no two disks of different replicas belonging to same diskgroup are ASM candidate disks.
- For restore operations, target LUNs should not be visible to the production RAC.
- The reserve policy of each replica disk should be set to `no_reserve`. See [“Setting the reserve policy on AIX shared disks”](#) for more details.

A replica previously mounted as standalone to one node of a RAC cluster prevents an attempt to mount another replica as Real Application Cluster to any other node of the same RAC cluster. For versions of Oracle before 11gR2, if a RAC replica is mounted as standalone to a node in the target RAC cluster, you should not mount another replica as RAC to another node of the same RAC cluster.

For versions of Oracle before 11gR2, if a RAC replica is mounted as standalone on a node N1 in the target RAC cluster, you should use node N1 for mounting any other replicas as standalone. Do not mount a RAC replica as standalone to a different node in the RAC cluster if another replica is already mounted to a node in the RAC.

It is possible to first perform **Mount as Real Application Cluster** to one node of a RAC cluster followed by a standalone mount to any of the other nodes of the that same RAC cluster.

Creating an application set for an Oracle 11gR2 RAC database

To create an application set for an Oracle 11g R2 RAC database:

1. Login as a grid user.
2. From the `$ORACLE_HOME/bin` directory, stop the listener using the following command:

```
lsnrctl stop
```

3. Login as the database user.

- Navigate to the \$TNS_ADMIN location and edit the tnsnames.ora file and add an entry like the following example:

```
RACDB1 =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST =
SCAN.rmcluster.local) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = racdb)
      (INSTANCE_NAME = racdb1)
    )
  )
```

Set the parameter INSTANCE_NAME to the Oracle RAC instance running on the node. This ensures a connection to the correct instance on the registered node during application set creation (racdb1 in the example above).

- Login as the grid user, go to \$ORACLE_HOME/bin directory and start the listener using the following command:

```
lsnrctl start
```

- Create an application set through the Replication Manager Console.

Setting the reserve policy on AIX shared disks

To enable simultaneous access to a AIX disk device from multiple nodes, set the appropriate Object Data Manager (ODM) attribute to the correct value.

- ◆ Disk type: ESS, EMC, HDS, CLARiiON, or MPIO-capable disks
- ◆ Attribute: reserve_policy
- ◆ Value: no_reserve

To determine the current value of the attribute, run the following command on all cluster nodes for each disk device that you want to use: `/usr/sbin/lsattr -E -l hdiskn`

To set the correct value for the attribute, run the following command on all cluster nodes for each disk device that you want to use:

```
/usr/sbin/chdev -l hdiskn -a reserve_policy=no_reserve
```

where *n* in *hdiskn* is the number of the disk device in your environment.

Oracle configuration detailed concepts

The following sections provide more detailed information about certain Oracle setup and configuration concepts to help administrators understand how to setup Replication Manager in an Oracle environment.

A note about Oracle user privileges

Replication Manager is designed to work with a range of users that can have different Oracle privileges. Each privilege level grants users more access to Replication Manager functionality.

[Table 24 on page 337](#) outlines what a user with certain Oracle user privileges can and cannot do in Replication Manager.

Table 24 Oracle user privileges and Replication Manager

Oracle user privilege level	Allowed Replication Manager tasks
USER without DBA role	Cannot use Replication Manager to create Oracle replicas.
USER with DBA role	<p>Can perform the following tasks:</p> <ul style="list-style-type: none"> • Create application sets • Create replicas (provided the replication method does not take the database offline) • Mount using the Generate scripts for manual recovery option <p>Cannot perform these tasks:</p> <ul style="list-style-type: none"> • Create offline replicas • Mount using Read Only mode • Mount using Read/Write mode • Restore a replica • Use the Copy BCT File option
USER with SYSDBA role	Full access to all Replication Manager functionality.

Note: For Oracle 11gR2, the ASM instance user must have SYSASM privilege to perform any operations in Replication Manager.

Collecting important Oracle information

Before you create Oracle replicas, you need to collect some important information, such as the location of the log files. This section describes how to collect the information you will need:

1. Access SQLPlus:
 - In Windows, navigate to **Start > Programs > Oracle > OraHome > Application Development > SQLPlus**.
 - In UNIX or Linux, log in as an Oracle user and enter:


```
$ORACLE_HOME/bin/sqlplus
```
2. Use the following commands in [Table 25 on page 338](#).

Table 25 Accessing Oracle information

Information needed	Example command
Location of datafiles	<code>SELECT name FROM v\$datafile;</code>
Location of redo logs	<code>SELECT * FROM v\$logfile;</code>
Location of control files	<code>SELECT * FROM v\$controlfile;</code>
Location of archive logs	<code>archive log list;</code>
List of users with SYSDBA privileges (users with SYSDBA privileges are required to execute offline backups)	<code>SELECT * FROM v\$pwfile_users;</code>
Database mode	<code>SELECT * FROM v\$database;</code>
List of users with SYSASM privileges (users with SYSASM privileges are required to perform Replication Manager operations on Oracle 11gR2 ASM databases)	<code>SELECT * FROM v\$pwfile_users;</code>

Configuring the Oracle TCP/IP listener

Without the Oracle TCP/IP Listener configured properly and running, Replication Manager cannot create an application set or replica of an Oracle database.

To configure and start the listener:

1. Set up the following two files in the directory `$ORACLE_HOME/network/admin`:
 - `tnsnames.ora`
 - `listener.ora`
2. After the files are present, start the listener with the following command:
`lsnrctl start`

For more information about the Oracle TCP/IP Listener, refer to the Oracle documentation set.

Adding an undiscovered Oracle instance (UNIX only)

If Replication Manager is operating in an environment with Oracle Enterprise Manager (OEM) and RAC set up in the Oracle environment, it may not be possible to discover SIDs that are not listed in the primary `oratab` file. If this situation occurs, it is possible to add an instance (SID) that was not discovered during application set creation.

To add an Oracle instance that was not discovered in UNIX environments:

1. Right-click **Application Sets** on the tree panel and click **New Application Set**.
2. Click **Next** on the welcome screen and expand a host from which the instance you would like to use is accessible. Also expand the **Oracle** folder under that host.



Figure 150 Adding an instance during application set creation

3. Select the **+Add Instance** checkbox shown in [Figure 150 on page 340](#). An Application Credentials panel like the one shown in [Figure 151 on page 341](#) appears to allow entry of the specifics of how to connect to the Oracle Instance.

Set Application Login (Oracle: DUAL)

Enter the Oracle database username needed to connect to the database, and the operating system username you used during the installation of Oracle.

Database

Username:

Password:

Connect String:

OS Username:

TNS_ADMIN:

ORACLE_HOME:

Figure 151 Application credentials panel

Set Application Login (Oracle: DUAL)

Enter the Oracle database username needed to connect to the database, and the operating system username you used during the installation of Oracle.

Database

Username:

Password:

Connect String:

OS Username:

TNS_ADMIN:

ORACLE_HOME:

ASM

Username:

Password:

Instance Name (SID):

ORACLE_HOME:

Figure 152 Application credentials panel (cont.)

Note: In case of RAC environments, select the global database name as the instance, instead of the specific instance running on the RAC node for which you are creating application set.

Deploying Replication Manager in an MSCS Cluster

Replication Manager can be configured to support Oracle deployed in a MSCS Cluster with or without Oracle Failsafe. The configuration prerequisites associated with this configuration:

1. Ensure that `listener.ora` and `tnsnames.ora` are configured to use the virtual IP address to refer to the host where the Oracle database resides.
2. Install the Replication Manager Oracle Agent on all physical nodes of the MSCS Cluster.
3. Register the Oracle Agent only once with the Replication Manager Server using the virtual IP address that always accesses the active node of the cluster.

With this configuration, a single application set and job can be used to perform online replications of this environment, regardless of which node is currently active.

In addition to these requirements, there are other considerations when creating offline replications or restoring in an MSCS Cluster environment. These considerations are outlined in [“Performing offline replications in an MSCS Cluster with Oracle FailSafe” on page 356](#) and [“Restoring Oracle in an MSCS Cluster” on page 382](#).

Choosing an Oracle archive log location to replicate

Replication Manager uses the following process to choose which archive log location to replicate:

- ◆ If a `log_archive_dest` value is specified as `LOCATION=USE_DB_RECOVERY_FILE_DEST`, then Replication Manager uses that location.
- ◆ If `USE_DB_RECOVERY_FILE_DEST` is not used, Replication Manager chooses the first defined `log_archive_dest` location. The location must be local to the host.

In the example below, `dest_1` will be the archive log location used by Replication Manager, because it is the first defined location and `USE_DB_RECOVERY_FILE_DEST` is not used:

```
log_archive_dest_1 = 'LOCATION=/arch/ '
log_archive_dest_2 = 'LOCATION=/arch2 MANDATORY'
```

In the following example, `dest_2` will be used, because `USE_DB_RECOVERY_FILE_DEST` is specified:

```
log_archive_dest_1 = 'LOCATION=/arch/ '
log_archive_dest_2 = 'LOCATION=USE_DB_RECOVERY_FILE_DEST'
```

Using Oracle ASM with Replication Manager

Oracle ASM environments must be configured in specific ways in order for Replication Manager to interact with the Oracle ASM environment successfully. This section provides the details of those special configuration needs.

Oracle ASM disk partition restrictions

For Replication Manager to work properly, a partition containing the entire LUN must be used for the disks that are presented to ASM. Additionally, the Oracle ASM disk cannot be in a partition that includes the partition table because the partition table would be overwritten.

For example, the Linux partition `/dev/sdb1` would be appropriate where:

```
fdisk /dev/sdb
Command (m for help): par
Disk /dev/sdb: 4294 MB, 4294967296 bytes
255 heads, 63 sectors/track, 522 cylinders
Units = cylinders of 16065 * 512 = 8225280 bytes

Device Boot    Start    End    Blocks    Id    System
/dev/sdb1      1        522    4192933+  83    Linux
```

As another example, on Solaris, typically slice 0 points to the partition table. Slice 1 is the swap partition. Slice 2 points to the whole disk including the partition table. In this case slice 6 should be used.

Oracle ASM database layout restrictions

Replication Manager requires a specific database layout to support replication, mount, and restore of data using Oracle ASM. All datafiles, online redo logs, and controlfiles must reside in ASM disk groups. Replication Manager always replicates data at disk group granularity.

Oracle ASM rebalancing restrictions

When Oracle stores datafiles in an Oracle ASM environment, ASM is constantly rebalancing data to try to maintain a balanced load across the existing disks. This does not present a problem for Replication Manager because Replication Manager communicates with Oracle before creating the replica to verify that load-balancing changes have been suspended during the time the replica is created.

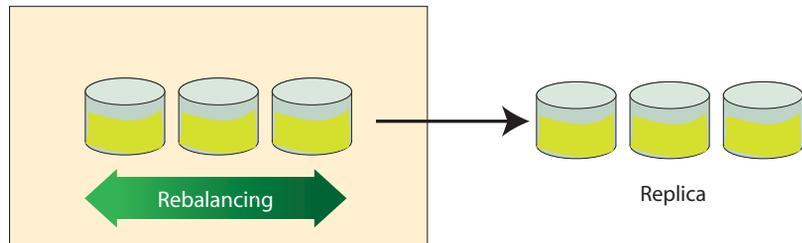


WARNING

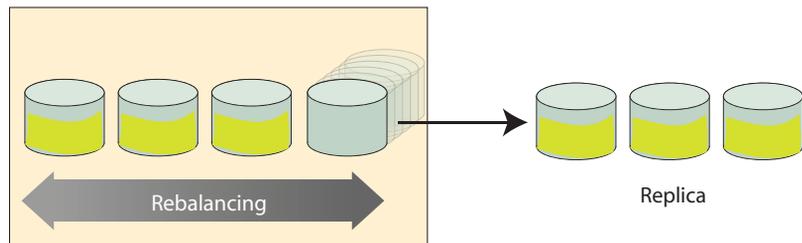
If an administrator changes the number of physical disks in an Oracle ASM disk group after replicas of that disk group have been created, that change could cause problems when Replication Manager restores those replicas. Replication Manager will issue a warning in this case.

Before changing the number of disks in an Oracle ASM disk group, EMC recommends that you delete all existing replicas of the data on that disk group and re-create those replicas after the change. Refer to [Figure 153 on page 346](#) for a graphical representation of potential issues with a restore in this situation.

ASM disk group (has 3 disks) rebalancing to distribute data evenly



After replication, add a disk to the ASM disk group



Rebalancing re-enabled, restore of the replica is not recommended

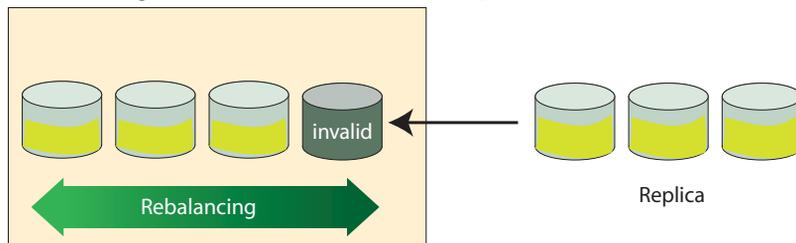


Figure 153 Affect of changing the number of disks in an Oracle ASM disk group

The case of adding a disk to the disk group is one scenario, another scenario that can cause issues is if a disk is dropped from the disk group between the time the replica is taken and the time it is restored. EMC does not recommend restoring a replica with a different number of disks than the current disk group configuration. In the case shown above the 4th disk would become invalid and would have to be explicitly excluded from ASM manually.

For Oracle running on UNIX platforms and Linux, Replication Manager requires that the ASM instance on the host be running when a restore takes place. Additionally, EMC recommends that all disk groups to be restored be in a mounted state. This allows Replication Manager to run additional checks before the restore to verify that the layout of the ASM disk groups are the same as during the replication.

If ASM is running, but the involved disk groups are not all mounted, the restore will proceed, but the layout checks will not happen correctly, and warnings will be displayed.

At the end of the restore, Replication Manager attempts to remount the involved disk groups. Replication Manager does not attempt to restart the database or perform any type of database recovery, once the restore has finished.

Oracle ASM redundancy level restrictions

Oracle ASM has three different redundancy levels to choose from:

- ◆ Normal redundancy — Institutes a single software mirror for each disk in the disk group.
- ◆ High redundancy — Institutes two software mirrors for each disk in the disk group.
- ◆ External redundancy — Depends upon the hardware layer to provide any redundancy, whether through mirroring or RAID or some other means.

Replication Manager supports all levels of ASM redundancy. The following requirements and restrictions apply for normal and high redundancy:

- ◆ All disks of a given diskgroup will be replicated. For normal and high redundancy, this means that mirrored copies in all failure groups will be added to the replica. The allocated storage for the replica has to account for all copies of each disk.
- ◆ All disks of all failure groups must be visible to the production host.
- ◆ Diskgroups can span multiple storage arrays, however all the disks must be visible to the production host and, if the replica needs to be mounted, the targets must be visible to the mount host.
- ◆ All disks of all failure groups must be online at the time of replication.

The example below describes a valid diskgroup configuration:

Host 1 has a database built on a DISKGROUP DG1 with two failure groups of one disk each (normal redundancy). Failure group 1 is located on Symmetrix array A and failure group 2 is located on Symmetrix array B. Host 1 sees both Symmetrix arrays. A replica is created. The target LUNs for this replica are located respectively on Symmetrix A and B. Both Symmetrix arrays are visible to Host 2 (the mount host), and the target LUNs are zoned to host 2.

The examples below describe an invalid configuration:

- ◆ Host 1 has a database built on a DISKGROUP DG1 with two failure groups of one disk each (normal redundancy). Failure group 1 is located on Symmetrix array A and failure group 2 is located on Symmetrix array B. Host 1 sees Symmetrix A, but Symmetrix B is remote. The replication cannot proceed.

- ◆ Host 1 has a database built on a DISKGROUP DG1 with two failure groups of one disk each (normal redundancy). Failure group 1 is located on Symmetrix array A and failure group 2 is located on Symmetrix array B. Host 1 sees both Symmetrix arrays. A replica is created. The target LUNs for this replica are located respectively on Symmetrix A and B. Symmetrix array A is visible to host 2 but Symmetrix array B is not. The replication can occur but the mount will fail because the redundancy level of the diskgroup will not be satisfied, due to the second array not being visible.

Oracle ASM mount considerations

When you mount Oracle ASM disk groups back to the production host, or you mount several replicas of the same disk group to the same mount host, you must instruct Replication Manager to rename the disk groups to prevent naming collisions. To rename ASM disk groups as part of a mount operation, select the **Rename ASM Disk Groups** checkbox in the Oracle mount options panel and choose a prefix in the **Rename using the prefix** field.

Understanding Oracle application sets and jobs

You can control what an Oracle application set can replicate by selecting a database or a set of tablespaces to include in your replica when completing the Application Set Wizard. You cannot specify replications for selected datafiles; the finest level of granularity that you can specify is the tablespace.

Note: The software automatically selects the corresponding Oracle system tablespaces (SYSTEM, USERS, OEM_REPOSITORY, INDX, and UNDO) when present. It is not possible to deselect these system tablespaces because they are not displayed in the tablespace selection list.

Before each replication, Replication Manager discovers the location of the data to replicate. It identifies the pathnames for all the datafiles in the requested tablespaces.

Note: For partial restores, no other data besides the specific application you are trying to restore should be located in the same volume group. Refer to [“Restore limitations for data in file systems” on page 599](#) for more information.

Oracle instances that reside on Symmetrix systems can be included as part of a federated application set. Some restrictions apply and special configuration steps may be required. [Chapter 7, “Configuring Federated Data,”](#) provides full details on how to create Oracle federated application sets.

Comparing Oracle replication choices

When you configure the job to produce the replica, Replication Manager offers multiple methods for creating a replica of an Oracle database. The methods are:

- ◆ Without consistent split:
 - Online using hot backup mode (SYSDBA credentials not required)
 - Online without hot backup mode (SYSDBA credentials not required. Not available for a RecoverPoint job.)
 - Offline by shutting down the database (requires SYSDBA credentials to complete)

- ◆ With consistent split:
 - Online using hot backup mode (SYSDBA credentials not required)
 - Online without hot backup mode (SYSDBA credentials not required. Not available for a RecoverPoint job.)
 - Offline by shutting down the database (requires SYSDBA credentials to complete)

The following sections describe these options.

Online using hot backup mode (no consistent split)

Oracle online replication with hot backup mode allows you to create a replica of the data while the database continues to serve data. The database server tracks changes to the database in an online redo log, and can later apply those changes to the database, after the point-in-time replica has been created.

You can provide pre- and post-replication scripts to perform any specific tasks before and after the replica is created; however, if you do choose to specify pre- and post-replication scripts, the pre-replication script must invoke the Oracle Online Backup mode and the post-replication script must release the database from Online Backup mode. If you do not provide pre- and post-replication scripts, Replication Manager performs these tasks for you.

Configuring a replica with this option does not require SYSDBA credentials because this method does not have to shut down the database.

The Oracle Agent also performs the following tasks:

- ◆ Catalogs the replica by copying all relevant information to the Replication Manager Server database.
- ◆ Saves the archive redo logs covering the time since the start of the replication (after the first new log was started).

Note: You can back up (as part of a separate file system backup or third-party backup routine) any additional archive redo logs occurring after one replication (and up to the next) to enable recovery and roll-forward work to extend past the time the replica was taken. You can back up the archived redo logs from the production server and, with the most recent replica mounted to an alternate host, you can back up the Oracle datafiles at the same time.

Offline by shutting down the database (no consistent split)

For a step-by-step explanation of this process, refer to [Table 26 on page 352](#).

Oracle offline replication creates a replica of the data by first shutting down the database, and then creating the replica of the database in that offline state. To select this option, the user must have SYSDBA credentials.

You can provide pre- and post-replication scripts to perform any specific tasks before and after the replica is created; however, if you do choose to specify pre- and post-replication scripts, the pre-replication script must include the command to shut down the instance. If you provide a pre-replication script for an offline replication, Replication Manager checks that the database instance is shut down before splitting the mirrors. If it is not, Replication Manager fails the replication. If a post-replication script is specified for the job, the script must include the command to start up the database instance. If you do not provide pre- and post-replication scripts, Replication Manager performs these tasks for you.

This option has been included to allow easier integration of Replication Manager with third-party products (such as BR Backup) that must activate hot backup mode independent of Replication Manager. Automatic recovery is not guaranteed to work for replicas taken with the option because Replication Manager is not managing the archive logs. Replication Manager issues a warning during mounts and restores.

The Oracle Agent catalogs the replica by copying all relevant information to the replication server database.

For a step-by-step explanation of this process, refer to [Table 26 on page 352](#).

Comparison of replication options without consistent split

Table 26 on page 352 illustrates Oracle replication processing in both online and offline modes (without consistent split).

Table 26 Online and offline replications (without consistent split) (page 1 of 3)

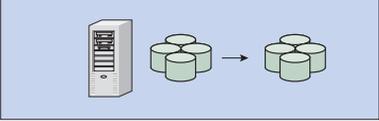
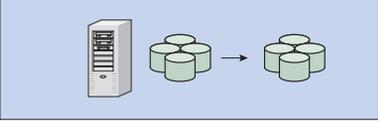
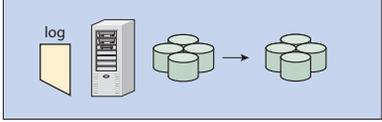
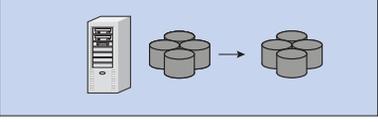
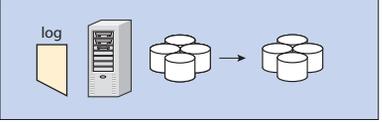
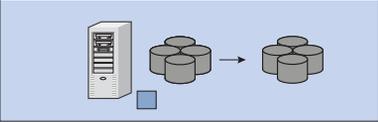
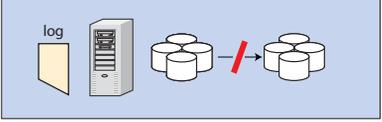
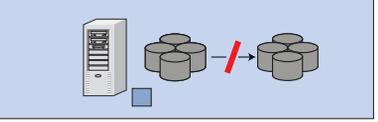
Online replication	Offline replication
<p>1. Establishes mirrors (if necessary).</p> 	<p>Establishes mirrors (if necessary).</p> 
<p>2. Signals Oracle to start a new Oracle redo log (performs a log switch).</p> 	<p>Shuts down the database.</p>  <p>(If you specify a pre-replication script, it must shut down the database.)</p>
<p>3. Puts the selected tablespaces into Online Backup mode and stops Oracle ASM rebalancing activity (if applicable).</p>  <p>(If you specify a pre-replication script, that script must put the database into Online Backup mode.)</p>	<p>Creates a backup control file (to use later in step 6) and stops Oracle ASM rebalancing activity (if applicable).</p> 
<p>4. Splits the mirrors/creates the replica.</p> 	<p>Splits the mirrors/creates the replica.</p> 

Table 26 Online and offline replications (without consistent split) (page 2 of 3)

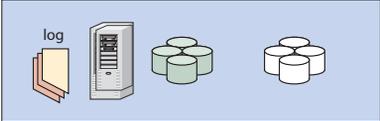
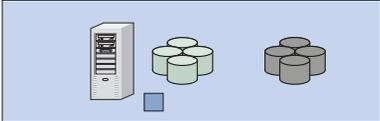
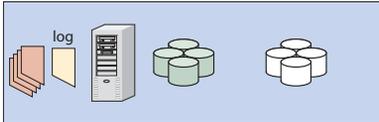
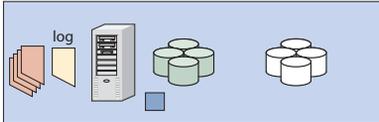
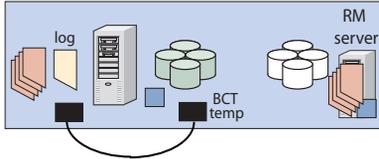
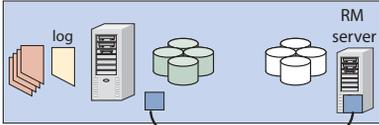
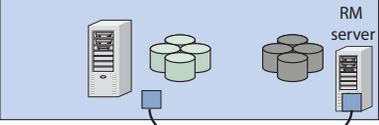
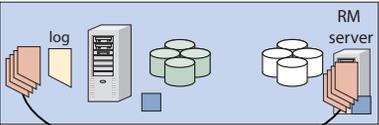
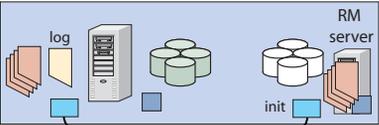
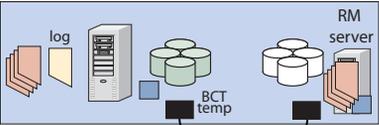
Online replication	Offline replication
<p>5. Takes tablespaces out of Online Backup mode.</p>  <p>(If you specify a post-replication script, it must take the database out of Online Backup mode) and restarts Oracle ASM rebalancing (if applicable).</p>	<p>Restarts the production database.</p>  <p>(If you specify a post-replication script, it must restart the database, if you want it restarted.) and resumes Oracle ASM rebalancing (if applicable).</p>
<p>6. Starts a new Oracle redo log and archives the redo log(s) created during the mirror split (performs a log switch).</p> 	
<p>7. Creates a backup control file in a temporary directory on the production host.</p> 	
<p>8. If the Copy BCT file option was selected, the BCT file is updated (incremented to the next backup slot) then copied to the temp area.</p> 	

Table 26 Online and offline replications (without consistent split) (page 3 of 3)

Online replication	Offline replication
<p>9. Transfers the b/u control file and catalog information to the Replication Manager Server.</p> 	<p>Transfers the backup control file and catalog information to the Replication Manager Server.</p> 
<p>10. Copies archived redo logs.</p> 	
<p>11. Copies init file.</p> 	
<p>12. Copies BCT file if applicable.</p> 	

Integrating replication with Oracle Block Change Tracking

Replication Manager supports the Oracle Block Change Tracking (BCT) feature in environments where RMAN is supported. Replication Manager supports this feature during replication by offering an option, **Copy BCT file**, in the Job Wizard. The BCT file will be available when you mount a replica using the **Catalog with**

RMAN mount option only. Other mounts or restores of the Oracle replica will not utilize the BCT file.

Replication Manager requires that BCT be enabled on the production database before this option is used. You can determine if BCT is enabled using the following SQLPlus query:

```
SELECT STATUS FROM V$BLOCK_CHANGE_TRACKING;
```

If you choose that option during job configuration, Replication Manager:

1. Invokes the Oracle stored procedure to advance the BCT file to the next available incremental backup slot.
2. Copies the BCT file so that it is part of the replica.

If **Copy BCT file** is selected but block change tracking is not enabled on the production database, the replication fails.

Controlling where temporary files are stored

To set an environment variable to control the storage location of temporary files created during an Oracle replication:

1. In the Replication Manager Console, right-click the Oracle host and select **Properties**.
2. On the **Advanced** tab, select **New**.
3. In the New Environment Variable dialog box, type **ERM_TEMP_BASE** for the variable name and specify a path for the base directory where logs are to be stored temporarily. Replication Manager adds another directory that represents the Oracle Database name to which the logs are related

Note: Verify you have enough space to hold all the archive logs that are generated during replication.

4. Click **OK** in the New Environment Variable dialog box.
5. To commit the environment variable, click **OK** in the **Advanced** tab.

If you do not set the environment variable, the temporary files on UNIX clients are stored in the default temp directory. On Windows clients, Replication Manager determines the location of temporary files by searching for them first in the path specified by the TMP environment variable, then by the path specified by TEMP.

The *Replication Manager Administrator's Guide* and the online help describe the use of the Advanced tab in more detail.

Online using hot backup mode (with consistent split)

You can choose to create a replica using both Oracle's hot backup mode and the consistent split capabilities of EMC's storage arrays. This option performs all the processing described in "[Online using hot backup mode \(no consistent split\)](#)" on page 350 but also uses the consistent split capabilities of the storage array.

Configuring a replica with this option does not require SYSDBA credentials because this method does not have to shut down the database.

For a step-by-step explanation of this process, refer to [Table 27 on page 358](#).

Online without hot backup mode (with consistent split)

This option keeps the Oracle database online and creates a consistent replica of the database using only the storage arrays consistent split capabilities.

As with the previous option, configuring a replica this way does not require SYSDBA credentials because this method does not have to shut down the database.

For a step-by-step explanation of this process, refer to [Table 27 on page 358](#).

Offline by shutting down the database (with consistent split)

You can choose to create a consistent replica by both shutting down the Oracle database and using the consistent split capabilities of EMC's storage arrays. This option performs all the processing described in "[Offline by shutting down the database \(no consistent split\)](#)" on page 351 but also use the consistent split capabilities of the storage array.

Configuring a replica this way requires that the user have SYSDBA credentials because the method shuts down the database.

Performing offline replications in an MSCS Cluster with Oracle FailSafe

When Replication Manager performs a offline backup, it attempts to take the Oracle instance offline prior to performing the replication. If that instance is part of a cluster resource group, MSCS may try to restart the instance right after the shutdown, causing Replication Manager's replication to fail.

To successfully perform offline replications in an MSCS cluster environment with Oracle Failsafe:

1. Create a pre script that uses Oracle Failsafe's `fscmd` command to offline the database resource. Use the `offline=immediate` switch to shutdown the Oracle instance cleanly.
2. Create a post script that uses Oracle Failsafe's `fscmd` command to online the database resource after the replication is complete.

For more information on Oracle Failsafe and the syntax of the `fscmd` command, refer to your Oracle documentation.

Comparison of replication options with consistent split

Replication Manager performs different steps depending on which replication option you choose. [Table 27 on page 358](#) and [Figure 154](#)

on page 360 illustrate consistent-split Oracle replication processing in all three modes.

Table 27 Online and offline replications (with consistent split) illustrated (page 1 of 2)

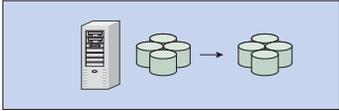
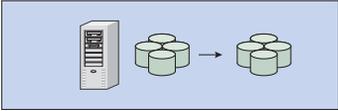
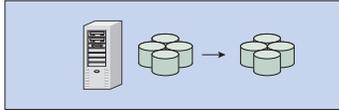
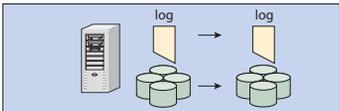
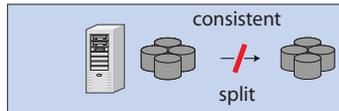
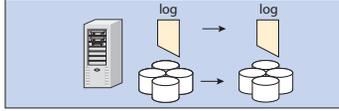
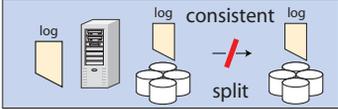
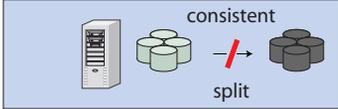
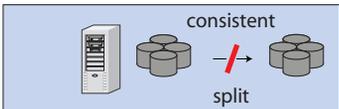
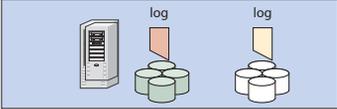
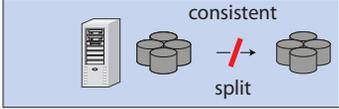
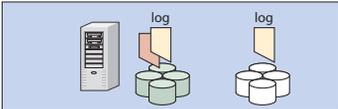
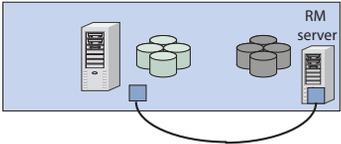
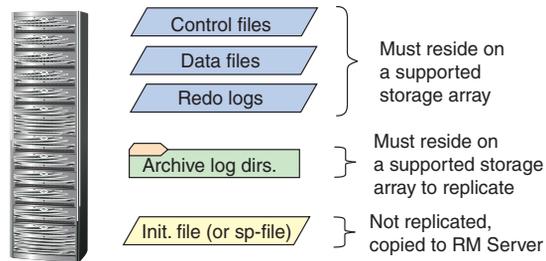
Online consistent-split replication (with hot backup mode)	Online consistent-split replication (without hot backup mode)	Offline consistent-split replication (by shutting down the database)
<p>1. Establishes mirrors (if necessary).</p> 	<p>Establishes mirrors (if necessary).</p> 	<p>Establishes mirrors (if necessary).</p> 
<p>2. Signals Oracle to start a new Oracle redo log (performs a log switch).</p>  <p>Logs are on supported array storage.</p>		<p>Shuts down the database.</p>  <p>(If you specify a pre-replication script, it must shut down the database.)</p>
<p>3. Puts the selected tablespaces into Online Backup mode.</p>  <p>(If you specify a pre-replication script, it must put the database into Online Backup mode.) If you select the entire database for replication, RM puts the database (10g and above) into hot backup mode.</p>		

Table 27 Online and offline replications (with consistent split) illustrated (page 2 of 2)

Online consistent-split replication (with hot backup mode)	Online consistent-split replication (without hot backup mode)	Offline consistent-split replication (by shutting down the database)
<p>4. Splits the mirrors/creates the replica using consistent-split technology.</p> 	<p>Splits the mirrors/creates the replica using consistent-split technology.</p> 	<p>Splits the mirrors/creates the replica using consistent-split technology.</p> 
<p>5. Takes tablespaces out of Online Backup mode.</p>  <p>(If you specify a post-replication script, it must take the database out of Online Backup mode.)</p> <p>If you select the entire database for replication, RMAN releases the entire database (10g and above) from hot backup mode.</p>	<p>This creates a crash-consistent copy of the database.</p>	<p>Restarts the production database.</p>  <p>(If you specify a post-replication script, it must restart the database, if you want it restarted.)</p>
<p>6. Starts a new Oracle redo log and archives the redo log(s) created during the mirror split (performs a log switch).</p> 		<p>Transfers the catalog information to the Replication Manager server.</p> 



RM-000015

Figure 154 Oracle consistent-split prerequisites

Upgraded Oracle jobs

Beginning in version 5.2.2, Replication Manager no longer catalogs the password file, and instead creates a new password file on the mount host. In order to create this file, a SYS password must be provided. For pre-5.2.2 jobs that specified a mount after replication, Replication Manager automatically assigns a SYS password that matches the password of the dba user in the application set. These jobs will continue to run without user action. However, if you modify such a job, you will be required to enter a new value for the SYS password.

Mounting Oracle replicas

The Oracle Agent enables Replication Manager to initiate and control mounts of Oracle databases (including control files, online redo logs, and optionally the Flash Recovery Area and archive logs).

The Oracle Agent can mount a replica to an alternate host, as long as that host has an installed Oracle database server. EMC recommends that you use the same version of Oracle on both the production and alternate hosts.

Replication Manager can perform the following types of mounts:

- ◆ On an alternate mount host to the same location as the production host
- ◆ On an alternate mount host to a new location (determined by adding an alternate root to the pathname)
- ◆ On alternate mount host to a new location (determined by path mapping)

Note: Alternate root or path mapping does not apply to ASM disk groups. ASM disk groups can be renamed using a prefix only. This is described in the section entitled [“Oracle ASM mount considerations” on page 348](#).

- ◆ To the original production host in a new location (determined by adding an alternate root to the pathname)
- ◆ To the original production host in a new location (determined by path mapping)
- ◆ Using an alternate SID name to prevent SID name collisions when mounting multiple replicas of the same database to the same mount host
- ◆ Using an alternate database name (EMC does not recommend a restore after such a mount)

Note: Specific information about how alternate paths and path mapping work can be found at [“Mounting using alternate path” on page 174](#) or [“Mounting using path mapping” on page 177](#).

Mount setup and prerequisites

In order to perform a successful mount, certain prerequisites and setup steps need to be followed. This section outlines those prerequisites.

The Oracle agent can mount a replica to an alternate host as long as the following conditions are met:

- ◆ The mount host has an Oracle database server installed (“software only” install will suffice).

Note: Mounts using the **do not perform database operations** option in UNIX environments do not require that Oracle binaries be installed on the mount host.

- ◆ The mount host has Replication Manager Agent software installed.
- ◆ The mount host and production host are registered with the Replication Manager software.
- ◆ The mount host and the production host have identical operating system, volume manager, file system, and application versions.
- ◆ Mounts to a Linux ASM RAC cluster have all cluster nodes pre-configured as described in [“Configuring ASM RAC mount to ASM RAC” on page 334](#).
- ◆ For Oracle 10g (Release 1) in Windows environments, mounts can only succeed when mounting to the host that is NOT acting as a domain controller.
- ◆ For Oracle 10.2 on Linux 32-bit hosts, available mount options are limited to Do not perform database operations and Generate scripts for manual recovery.
- ◆ For Oracle 11gR2, the ASM instance on the mount host must be started before a mount operation is performed. Each ASM instance on the mount host must have a valid and accessible spfile. This is a mandatory requirement.

User account recommendations

For Oracle versions prior to Oracle 11gR2 in the case of an Oracle ASM setup, you must use a SYS user account. However, for Oracle 11gR2, any ASM instance user (SYS or any other user) that has SYSASM and SYSDBA privileges should be used.

Customizing the initialization parameters used for mount

For more information on how to accommodate Oracle binaries in a file system, refer to “Mounting replicas containing Oracle binaries” on page 376.

Use the Replication Manager Console to enter customized initialization parameters for mounting Oracle data. When the mount operation runs, these parameters will be appended to the copy of the production initialization file used with the mounted database. In previous versions of Replication Manager, customized parameters were entered by editing a special init file. Any existing customizations in this file will continue to be recognized but use of the Console is encouraged.

When the mount operation runs, this init file will be combined with the existing init file to create a new version of the init file used with the mounted database.

Setting parameters with the Console

To customize initialization parameters:

1. In the Replication Manager Console, click **Customizing Initialization Parameters** under Mount options (in the Job wizard, Job properties, or Mount wizard). The **Customize Initialization Parameters** window appears.
2. The following operations are available:
 - To add a new parameter, click **New** and enter a parameter name and value.
 - To disable a parameter but not remove it from the list, clear the checkbox.
 - To remove a parameter, select it and click **Delete**.

Notes on setting custom initialization parameters

Note the following when setting initialization parameters using the Replication Manager Console:

- ◆ The values that you set here override any parameters that are set in an init file.
- ◆ Only the parameters that were added using this table are displayed in the list.
- ◆ Use care when specifying parameter names and values. An invalid entry can have an undesirable result (for example, it can prevent the Oracle instance from starting).

- ◆ No special permissions are required of the user account running Replication Manager Console.
- ◆ The following parameters cannot be set in the Replication Manager Console: `control_files`, `db_recovery_file_dest`, `log_archive_dest*`, `db_name`, and `instance_name`.
- ◆ In general, if the file or directory specified by a parameter value does not exist, Replication Manager does not create it. The exceptions are values for `background_dump_dest`, `core_dump_dest`, and `user_dump_dest`. In these cases, Replication Manager creates the directory on mount and deletes it on unmount.

In 11gR1 and later, the dump directories parameters still exist but are deprecated in favor of `diagnostic_dest`. In that case, if you set any of the `*_dump_dest` parameters to a value, it will be ignored in favor of an automatically generated directory under the location pointed by `diagnostic_dest`.

Prerequisites to support Recovery Manager (RMAN) integration

There are specific prerequisites that must be met in order for Replication Manager to integrate with RMAN. These prerequisites are described here:

- ◆ RMAN catalog database must exist and be accessible on the same network as the mount host.
- ◆ The `tnsnames.ora` file on the mount host must contain a `tnsalias` that points to the RMAN catalog database where Replication Manager should catalog the replica.
- ◆ The catalog and catalog owner must be created prior to mounting a replica to be cataloged.
- ◆ Production database must be registered in the RMAN catalog before mounting the replica.
- ◆ The Oracle version running the RMAN catalog database must be equal to or greater than the highest Oracle version of all production databases registered to that catalog.
- ◆ Production databases that integrate with RMAN must be at least Oracle 10g.

Replication Manager can be configured to skip the cataloging of Oracle data files for all mounts on a specified mount host. This can be accomplished by setting the following environment variable to 1:

```
ERM_ORACLE_RMAN_NO_CATALOG_DATAFILES
```

Set the variable for the mount host with the Replication Manager Console, under Host properties, Advanced tab. If you do not set the environment variable, all datafiles and logs will be cataloged. In some cases it is desirable to prevent the cataloging of datafiles, since that is not required for backups and skipping that cataloging can improve backup performance in environments with many datafiles.

Oracle mounts to an alternate location

When you mount an Oracle database to an alternate location on an alternate mount host, you select the alternate mount host and the new path where the data should be mounted on that mount host. Replication Manager mounts the data to the alternate mount host in the new location. Alternate path mounts invalidate any flashback recovery area that is part of the replica.

This method allows you to mount several copies of Oracle data from different servers onto the same mount host. However, an Oracle restriction prevents you from running two databases with the same name on the same production host at the same time, even if they are using different SIDs.

Replication Manager includes an Oracle Database Rename feature to modify the databases after such a mount. If the database names are different and there is a collision of SIDs, you can also instruct Replication Manager to rename the SID that it is using for one or both of the databases.

Note: If you use the Oracle database rename feature during a mount of a certain replica, subsequent mounts of the same replica will also require you to use the database rename function. Additionally, EMC does not recommend that you use a replica that has undergone a database rename as a restore candidate.

Recovery considerations for renamed databases or SIDs

If you need to recover an Oracle replica that you mounted and you chose to rename the database for that replica, the following considerations apply:

- ◆ When using the **Generate scripts for manual recovery** mount option, there are some manual procedures that must be performed before you recover the database. For more information about using the **Generate scripts for manual recovery**, see *“Mounting and generating scripts for manual recovery”* on page 371.

- ◆ If you rename the database, remember that flashback recovery is disabled.
- ◆ The **Check for existing SID on mount host** option should remain enabled in most cases to protect from inadvertently overwriting critical files that already exist on the mount host. In some cases, skipping that SID check is desirable, such as when one mount fails, resulting in files being left behind that cause subsequent mounts to fail the SID check. In this case there is no value in the incomplete instance and therefore no reason to check for an existing SID and cause the mount to fail. However care should be taken when clearing that checkbox.

Oracle mounts to the production host

When you choose to mount an Oracle replica to the production host, you must select an alternate path where that replica data can reside on the production host. Then the system will not overwrite the production data already mounted on that host. Replication Manager also changes the appropriate application data.

The ability to mount data back onto the same production server can reduce the overall number of servers needed to review the data. No extra mount server is necessary. Data can be mounted to an alternate location by changing the root path or by using path mapping. However, because the database names will be identical if you mount the same data back to the production host, you will not be able to restart the database on the production host (unless you shut down the production database or use the database rename feature as part of the mount operation).

Oracle mount database recovery options

Oracle mount options provide a wide range of choices regarding how much processing Replication Manager performs on the mount host, after surfacing Oracle. A sample of the Oracle Mount Options panel is shown in [Figure 155 on page 367](#).

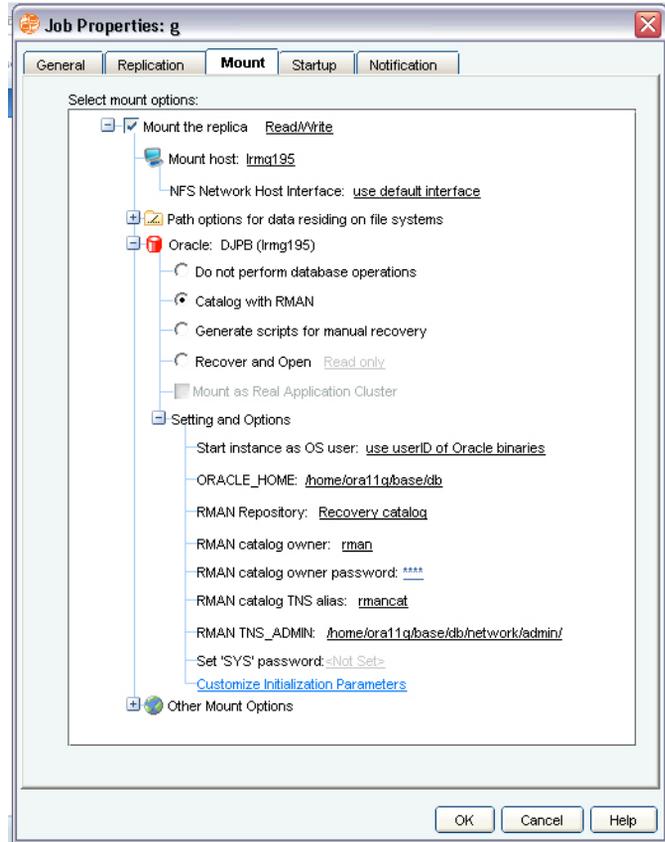


Figure 155 Oracle Mount Options panel

This section describes each of these options in detail:

- ◆ Mounting without database operations
- ◆ Mounting for use with Recovery Manager (RMAN)
- ◆ Mounting and generating scripts for manual recovery
- ◆ Mounting and recovering read only
- ◆ Mounting and recovering read/write
- ◆ Mounting a replica of a Veritas Quick I/O environment
- ◆ Mounting to a Real Application Cluster (RAC)
- ◆ Mounting replicas with Oracle ASM data
- ◆ Mounting replicas containing Oracle binaries

The details of each of these is outlined in the following sections:

Mounting without database operations

This operation, (formerly known as an Oracle file system mount) mounts the file systems that are part of the replica on the mount host along with their metadata collected at the time of replication. Refer to [“Controlling where temporary files are stored” on page 355](#) for more information. Subsequently, the Database Administrator can start the instance and recover the database manually by using the control file to bring the database online.



CAUTION

Mounting without database operations can be performed as many times as you like as long as the database has not been recovered manually. Because Replication Manager cannot predict what the user will do with data recovered in file system mode, Replication Manager does not save the control files during an unmount.

Mounting for use with Recovery Manager

Oracle installations with storage protected by Replication Manager can integrate with Oracle's Recovery Manager (RMAN) utility. Replication Manager can catalog a database replica contents with Oracle Recovery Manager (RMAN). This catalog can then be used to perform any RMAN operations outside of Replication Manager.

RMAN integration is supported on Linux, on Solaris, and HP-UX. Note that on Solaris and HP-UX, if ASM is used, the replica must be mounted on an alternate host in order to enable cataloging with RMAN; a replica of an ASM database mounted on a production host cannot be cataloged with RMAN (unless the ASM version is 11gR2 in which case production mount is allowed).

Replication Manager does *not* integrate in such a way that it runs the RMAN backup scripts; rather, it facilitates their usage once a replica has been mounted.

Cataloging with the RMAN repository can be done using a remote database, in which case the Recovery Catalog owner and related connection details will be required, or by the controlfile, in which case remote connection details are no longer applicable and the local controlfile of the mounted replica is used to store the RMAN repository.

Choosing the mount option **Catalog with RMAN** performs the following preparatory operations to facilitate the use of RMAN with Replication Manager replicas:

- ◆ Starts the Oracle instance on the mount host and brings the database to a mounted state
- ◆ Databases remains unopened
- ◆ Replication Manager automatically catalogs the components of the mounted Oracle replica either in the remote RMAN recovery catalog or in the local controlfile of the mounted replica.

Components that are cataloged include:

- Datafiles
- Backup control files
- Archive logs
- Flash Recovery Area contents (if enabled)

Also, the BCT file is included as part of the replica if the **Copy BCT File** option was selected during job creation. This file acts as a performance booster to improve RMANs incremental backup efficiency.

The BCT file and its use are only relevant when using a remote RMAN repository. Incremental backups are not applicable when using the controlfile based repository because a new controlfile is generated with each Replication Manager replica and consequently does not have the previous full and incremental backup history.

This cataloging allows administrators to utilize the following RMAN capabilities in conjunction with the mounted replica:

- ◆ View the contents of the cataloged replica(s) using RMAN commands such as:
list datafilecopy all;
- ◆ When mounting to the production host, RMAN cataloging of the replica's contents allows the following operations to be possible from the RMAN command line utility:
 - Perform RMAN individual file restore from the mounted replica. If a datafile has been lost or damaged, it can be restored using the following RMAN command:
restore/recover datafile X
 - Perform RMAN individual tablespace from the mounted replica. If a tablespace has been lost or damaged, it can be restored using the following RMAN command:
restore/recover tablespace X
 - Perform RMAN block-level recovery from the mounted replica. For example, a corrupt block can be recovered using the following RMAN command:
blockrecover datafile X block Y

Granular restores can be done both when the replica is cataloged in a remote recovery catalog or in a local controlfile. Using a local controlfile for RMAN catalogs allows an easy, on-the-fly cataloging of the datafiles and other artifacts of the replica into the RMAN repository of the production database without having to setup a remote RMAN catalog.

Consult the Oracle RMAN documentation for further help on commands and syntax. RMAN restore operations are possible for as long as the Replication Manager replica stays mounted/accessible to the production host.

Unmounting a replica cataloged with RMAN

Whenever Replication Manager unmounts an Oracle replica mounted using this option, the replica is uncataloged from RMAN to indicate that it is no longer available for the recovery operations.

Notes and restrictions when integrating with RMAN

The following restrictions apply to the Catalog with RMAN mount option:

- ◆ Replicas created without hot backup mode are not eligible for integration with RMAN.
- ◆ Replicas mounted with RMAN integration cannot be renamed using Replication Manager's database rename option.

Mounting and generating scripts for manual recovery

Choosing the option **Mounting and generating scripts for manual recovery**, formerly known as Prepare only, requires the user to complete additional fields. For specific information about these fields, refer to [“Additional fields required for some types of mounts” on page 377](#).

This recovery mode mounts the replica and prepares the environment so that the administrator can recover the database manually. EMC Replication Manager can import necessary files to the ERM_TEMP_BASE directory (or /tmp if ERM_TEMP_BASE has not been defined). Logs and metadata files are moved to the same directory as is defined for the control file on the production host.

Note: This type of replication and mount can be performed by users who do not have SYSDBA privileges, however if an sp-file is in use, Replication Manager provides a copy of the sp-file on Windows or a backup of the sp-file in UNIX environments. These files are not immediately usable in that format. Users must modify the sp-file prior to recovery. To avoid having to modify the sp-file, perform the replication and mount operations with a user that has SYSDBA privileges.

The Replication Manager recovery scripts are SQLPlus scripts that bring up the instance/database in a state ready to receive recovery. The RMAN scripts perform the actual recovery.

The SQLPlus script usage is as follows:

```
su - <oracleuser>
setenv ORACLE_SID <SID_name>
sqlplus "/" as sysdba" @/tmp/<SID_name>/<scriptname.sql>
where <scriptname.sql> is one of the following:
```

- ◆ **rm_recovery_startup_mount_script.sql** — Will start up the oracle instance
- ◆ **rm_recovery_drop_files_script.sql** — Will drop files that are not part of the replica (if the file does not contain any commands and it is a full replica of the database, it can be skipped)
- ◆ **rm_recovery_rename_paths_script.sql** — Will update the datafile and online redo logs file paths

The RMAN script usage is as follows:

```
rman target="/" cmdfile="/tmp/<SID_name>/<scriptname.sql>"
```

where *<scriptname.sql>* is one of the following:

- ◆ **rman_recovery_script.sql** — Will recover the database

rman_recovery_script2.sql — Is empty when using the "generate scripts" option, but otherwise contains the database open command:

```
sql 'alter database open [read only | resetlogs] ';
```

This type of mount can be performed as many times as you like as long as the database has not been recovered manually.

For replicas of Oracle 11gR2 ASM databases, mount in this recovery mode generates an `asm_steps.txt` file in the `/ERM_TEMP_BASE/<SID_name>/` location (where `<SID_name>` is the database SID) on the mount host. This file gives a list of manual steps to be performed on the mount host's ASM instance before database recovery can be done. In ASM RAC environments, an additional file (`asm_rac_steps.txt`) is generated in the same location. This file gives a list of manual steps to be performed on the mount host before database recovery can be done in RAC mode.

If the database was recovered on the mount host, the user should shutdown the database and perform the steps in the unmount section of the file, before unmounting the replica through Replication Manager GUI.

Mounting and recovering read only

Choosing the option **Recover and open Read only**, requires the user to complete additional fields. For specific information about these fields, refer to ["Additional fields required for some types of mounts"](#) on page 377.

This type of mount restores files and opens the database in *read-only* mode. Subsequently, the Database Administrator can manually apply archived redo logs accumulated since the mirror split. You can mount

and recover read-only as many times as necessary, because the control files that are changed in this mode are stored on the Replication Manager Server when the database is unmounted. Replication Manager uses those stored control files because they are guaranteed to be synchronized with the data from the replica.

In case of Oracle 11gR2 mounts with read-only recovery, Replication Manager connects to the mount host ASM instance using the user credentials entered in the Replication Manager mount options panel to perform ASM diskstring and diskgroup operations.

Mount and recover read only prevents flashback recovery operations.

Mounting and recovering read/write

Choosing the option **Recover and open Read/write** requires the user to complete additional fields. For specific information about these fields, refer to [“Additional fields required for some types of mounts” on page 377](#).

This type of mount recovers the Oracle database automatically and opens the instance. Users can mount the database using **Recover and open Read/Write** more than once because the control files that are changed during the first mount are stored on the Replication Manager Server when the database is unmounted. When future mounts occur, Replication Manager uses those stored control files because they are synchronized with the data from the replica.

In case of Oracle 11gR2, mounts with read-write recovery, Replication Manager connects to the mount host ASM instance using the user credentials entered in the Replication Manager mount options panel to perform ASM diskstring and diskgroup operations.



CAUTION

Except in the case of a replica created without hot-backup mode, Replication Manager imports control files to the mount host and saves these control files to the Replication Manager Server whenever the database gets unmounted from the mount host in Read/Write mode or Read-Only mode (even if reset logs are not applied).

Note: If you are using a replica created with consistent split, and you choose to mount and recover read only or read/write, the archive log of the production database should be part of the replica.

Mounting a Veritas Quick I/O replica

VERITAS Oracle Disk Manager and Veritas Quick I/O provide file system manageability with raw device performance. Replication Manager can support Oracle environments that have deployed ODM and Veritas Quick I/O on Solaris.

Additionally, the mount host may or may not be mounted with Veritas Quick I/O. During the mount operation Replication Manager identifies whether Veritas Quick I/O is enabled on the mount host or not and then mounts the replica with or without Veritas Quick I/O as appropriate.

Mounting to a Real Application Cluster

On the mount options panel you can select the **Mount as Real Application Cluster** checkbox, as long as the replica meets the following conditions:

- ◆ The database was created on asmlib devices.
- ◆ Replication Manager supports the use of two separate ORACLE_HOMES owned by distinct operating system users and groups.
- ◆ The application set includes a Linux-based ASM RAC database.
- ◆ The replica was created with the **Use consistent split** checkbox selected.
- ◆ The replica was created with the **Replicate archive log directory** checkbox selected.
- ◆ The mount setting **Generate scripts for manual recovery or Recover and open** is selected on the mount panel.

Note: If you mount to a Real Application Cluster using the **Generate scripts for manual recovery** setting, review [“Mounting and generating scripts for manual recovery” on page 371](#) for more information on critical recovery steps you must perform.

If you choose to mount to an Oracle RAC cluster the mount host you select must be a node in a target Oracle ASM RAC cluster. Additionally, you must provide the **Clusterware ORACLE HOME** for the Clusterware installed on the target ASM RAC cluster. The target ASM RAC cluster must be configured as defined in the section entitled [“Configuring ASM RAC mount to ASM RAC” on page 334](#).

Replication Manager automatically propagates the instances of the mounted database to all the nodes of the target ASM RAC cluster. Refer to the Replication Manager Online Help for more information on the mount panel options.

It is possible to mount a replica to an ASM RAC cluster, unmount the replica, and remount it again to a standalone environment.

Mounting replicas with Oracle ASM data

Replication Manager can mount replicas that originally resided in an Oracle ASM environment to an alternate host, or back to the production host, but the following prerequisites and restrictions apply to Oracle ASM mounts:

- ◆ As a prerequisite, after you install Oracle on the mount host, su to root and run the following command on the mount host once to start the `css` daemon that will enable ASM operations:

```
$ORACLE_HOME/bin/localconfig reset
```

Note: This does not apply to Oracle RAC.

- ◆ Production mounts are only allowed if you choose the disk group rename option so that two disk groups with the same name are not mounted on the same mount host simultaneously.
- ◆ For Oracle 11gR2 replicas of ASM databases to be mounted with read-only or read-write recovery, Replication Manager connects to the ASM instance on the mount host, as opposed to pre-11gR2 replicas, the ASM instance must be running ahead of time. The ASM instance user on the mount host must have SYSASM privilege.
- ◆ For mount hosts in environments where strict role separation is used between the Grid infrastructure home installation and the database home installation, an additional step is required for database instances to correctly communicate with the ASM instance and allow Replication Manager mounts to work properly. The group owner for the `$ORACLE_HOME/bin/oracle` executable of the database home needs to be changed to the `sysasm` OS group name of the Grid infrastructure installation.

For example, consider the two OS users: *oracle* (for the database installation) and *grid* (for the Grid infrastructure installation):

- **id oracle** — uid=7002(oracle) gid=7001(oinstall) groups=7002(dba),7004(asmdba),7006(oper),7001(oinstall)
- **id grid** — uid=7001(grid) gid=7001(oinstall) groups=7003(asmadmin),7004(asmdba),7005(asmoper),7001(oinstall)

When set up as such, the ownership of the oracle executable in the database oracle home after installation will be oracle:oinstall. This needs to be changed to oracle:asmadmin, where asmadmin is the group name corresponding to the SYSASM privilege on the Grid infrastructure installation. This does not indicate that the oracle user needs to be a member of asmadmin itself.

Refer to [Chapter 5, "Mount, Restore, and Recovery,"](#) for more information about possible mount scenarios.

Mounting replicas containing Oracle binaries

If you replicate and then mount Oracle binaries additional steps are required to allow those binaries to be used in their newly mounted location. These steps apply in any of the following situations:

- ◆ Application set contains Oracle binaries
- ◆ Oracle binaries mounted using an alternate path
- ◆ Changes to the users employed between the production and the mount host

If any of these situations apply to your use of Replication Manager with Oracle:

1. Create a callout script to run at level 600 on the mount host and make sure it has executable permission.
2. Include the following commands:

```
#!/bin/sh
ORACLE_HOME=<new Oracle binary home>
export ORACLE_HOME
```

If you are working in Oracle 10g, include the following additional lines:

```
cd $ORACLE_HOME/lib32
rm ldflags
ln -s $ORACLE_HOME/lib/ldflags ldflags
```

This is the end of script lines related to Oracle 10g.

```
cd $ORACLE_HOME/bin
su <oracle_unix_user_on_mount_host> -c "./relink all"
```

3. Name the callout script appropriately and add it to the correct directory on the mount host. For more information about creating, naming, and deploying callout scripts, refer to ["Using application callout scripts"](#) on page 235.

Additional fields required for some types of mounts

Users that choose mount modes that either recover the database or prepare for database recovery, must complete the following additional fields:

- ◆ Set the **Start instance as user** field to the username that is the owner of the Oracle binaries on the mount host and enter the corresponding password. (This username does not have to match the username from the production host.)
- ◆ Set the **ORACLE_HOME** field to include the appropriate value for the mount host.
- ◆ Optionally, type in a new **Database name** and/or **SID name**.
- ◆ Fail mount if SID exist.
- ◆ You must include a password in the field entitled **Set SYS password** so that the Replication Manager Oracle agent can create a new Oracle password file on the mount host.

Note: This does not apply to **Generate scripts for manual recovery** option.

- ◆ In 11gR2 environments during read-only or read-write recovery, Replication Manager offers the following fields to facilitate the connection to an ASM instance on the mount host:
 - **ASM instance name** — The name of the ASM instance on the mount host e.g. +ASM
In case of RAC to RAC mounts, the ASM instance should be the instance on the RAC node selected as mount host.
 - **Username** — The name of an ASM instance user with SYSASM privilege. Apart from SYS, any ASM instance user with SYSASM privilege is supported.
 - **Password** — The password of the specified ASM instance user.

Additional mount considerations

If you mount an Oracle database replica created without consistent split to an alternate host in Read/Write mode, Replication Manager:

1. Attaches the storage devices containing the Oracle data to the alternate host.

2. Moves the control file(s), init file, and archive logs associated with the replica to the alternate host in the directory specified by the `ERM_TEMP_BASE` environment variable.
3. Recovers the database and resets the logs. When the database is recovered and the logs are reset, Oracle marks the control files to indicate that the database has started a new log chain (new database incarnation).

If you are mounting a replica that was created using consistent split, then the control files and logs are part of the replicated data. If you are not using consistent split, these files are not located on the replicated devices, but they are copied via a network transfer to the mount host.

If you choose to mount an Oracle replica without recovering the database, Replication Manager does not reset the logs.

If the Oracle mount fails to create the password file on the mount host, check the value of `ORACLE_HOME` in the mount configuration. There should not be any leading or trailing spaces or extract characters following the actual `PATH` value for `ORACLE_HOME`.

Restoring Oracle replicas

When restoring an Oracle replica, you can choose to restore either the whole replica or a part of the replica. EMC recommends that the Database Administrator shut down the Oracle database prior to a restore to verify that there is no activity on the database during the restore operation. If the Oracle database is not shut down when you start a restore operation, Replication Manager shuts down the database automatically but does not restart the database automatically after the restore is complete. Restore operations also require a user with SYSDBA privileges.

Note: In certain configurations, when Oracle Restart is enabled, the database will be automatically restarted when shut down by Replication Manager. This prevents Replication Manager from properly shutting down the database during a restore. Therefore, Oracle Restart should be disabled prior to attempting a restore operation.

When you restore a partial database and the instance is up and running, the Oracle Agent takes the tablespaces themselves offline.

Note: Partial restores which include any of the core tablespaces (SYSTEM, USERS, TEMP, OEM_REPOSITORY, INDX, or UNDO) require a database shutdown because these tablespaces cannot be placed offline individually.

After the restore is complete, you can recover the tablespaces manually (if desired), subject to various limitations as described in [“Restore limitations for data in file systems” on page 599](#). Restore granularity for LVMs is at the volume group level. Refer to [“Restoring UNIX logical volumes” on page 597](#) for more information.

For Oracle restores, the Oracle Agent does not start the Oracle instance or recover the database. If the whole database is being restored, Replication Manager will shutdown the database.

Note: If you are using consistent-split technology, refer to [Chapter 6, “Using Consistent Split,”](#) for more information and considerations associated with restoring applications using that technology.

The Oracle Agent copies the archive logs to the archive log directory and the backup control file to a directory defined by the values of the ERM_TEMP_BASE environment variable and the name of the Oracle database as follows:

```
/ <ERM_TEMP_BASE> / <DB_name> /
```

The init file is restored to the ERM_TEMP_BASE directory of the Oracle production host, renamed to `init<SID_name>.ora.restored`.

Refer to “Controlling where temporary files are stored” on page 355 for more information.



CAUTION

If you are restoring a whole replica that was created using consistent split, any control files, archive logs, and flashback recovery area that was part of the replica will overwrite those on the production host, since they are part of the replicated data.

Post restore steps when restoring non consistent split replicas

Except in the case of a replica created with consistent split, control files and archive logs are moved to the directory specified by the ERM_TEMP_BASE environment variable on the production host because Replication Manager should not overwrite the existing control files or `init` files on the production host.

If necessary, the Database Administrator can retrieve these backup control files from that location to log in to the instance and manually execute the recovery. This is necessary if either of the following circumstances has occurred:

- ◆ The original archive logs and backup control files have been lost or corrupted.
- ◆ The replica you are restoring has been mounted in read/write mode, invalidating the existing control files.

In this case, the Database Administrator should follow these steps to complete the recovery:

Note: An example has been added to help illustrate the steps. In the example, a database is named TST4. Replication Manager has restored one archive log and three production controlfiles.

1. Review the contents of the temporary directory where the archive logs have been placed `<ERM_TEMP_BASE>/TST4`. This location contains the archive logs and control files that have been restored:

```
ls /tmp/TST4
```

```
1_5_623868257.dbf
```

```
ctrlfile_backup.ctl
```

2. Shutdown the database using SQLplus commands.
3. Overwrite the current controlfiles with the one from the ERM_TEMP_BASE location. In this example, use these commands:

```
cp /tmp/TST4/ctrlfile_backup.ctl
/oralg2/oradata/TST4/control01.ctl

cp /tmp/TST4/ctrlfile_backup.ctl
/oralg2/oradata/TST4/control02.ctl

cp /tmp/TST4/ctrlfile_backup.ctl
/oralg2/oradata/TST4/control03.ctl
```

4. Start the database and bring it to a mounted state using the following SQLPlus command:

```
startup mount;
```

5. Perform the recovery:

```
SQL> recover database using backup controlfile until
cancel;
```

```
ORA-00279: change 565593 generated at 05/29/2007
17:51:09 needed for thread 1
```

```
ORA-00289: suggestion : /oraflash2/1_5_623868257.dbf
```

```
ORA-00280: change 565593 for thread 1 is in sequence
#5
```

Transfer the archive log from the ERM_TEMP_BASE location because in this example, the archive log was no longer present in the archive log directory.

6. Replication Manager asks you to specify which log you want to apply. Enter the name of the log:

```
Specify log: {<RET>=suggested | filename | AUTO |
CANCEL}
```

```
/tmp/TST4/1_5_623868257.dbf
```

```
ORA-00279: change 565648 generated at 05/29/2007
17:51:10 needed for thread 1
```

```
ORA-00289: suggestion : /oraflash2/1_6_623868257.dbf
```

```
ORA-00280: change 565648 for thread 1 is in sequence
#6
```

```
ORA-00278: log file '/tmp/TST4/1_5_623868257.dbf' no
longer needed for this recovery
```

```
Specify log: {<RET>=suggested | filename | AUTO |
CANCEL}
```

CANCEL

Media recovery cancelled.

In the example, we applied log number

Oracle responds that log 5 has been applied previously and is no longer needed. Oracle suggests archive log 6. In the example, log 5 is the last log available to apply, so we terminate the recovery. More archive logs could be applied if available.

7. Once all the archive logs have been applied, open the database and reset the log sequence to 1. This is referred to as an incomplete point in time recovery, and a new incarnation of the database is created:

```
SQL> alter database open resetlogs;
Database altered.
```

Restoring a single Oracle table space

Replication Manager does not provide point-in-time recovery of a single Oracle table space as part of the restore process. Replication Manager will put the table space in the online state before the restore and will provide the necessary logs to recover the data, but you must manually perform the necessary recovery commands and log mining before recovery.

Shutdown Oracle 10.2 database on 32-bit Linux before restore

Before restoring an Oracle 10.2 replica on a host running 32-bit Linux, you must manually shutdown the database.

Restoring Oracle in an MSCS Cluster

Before Replication Manager restores to a Microsoft Cluster (MSCS), the Oracle database service resource must be offline. The procedure differs depending upon your cluster configuration.

Because the Oracle Agent was registered with the Replication Manager Server using the cluster's virtual IP address, you can restore to any active node, regardless of which node was active when the replica was created. This is true whether or not you are using Oracle FailSafe.

Perform these steps for all Oracle restores in an MSCS Cluster:

1. Go to Cluster Administrator and check whether the Oracle database resource is still offline. If the Oracle database resource is still online take it offline.

2. If Oracle Failsafe is in use, first use Oracle Failsafe `fscmd` to offline the resource.
3. Restore the selected Oracle database using the Replication Manager Console.
4. Recover the Oracle database using SQLPlus according to your DBA's requirements for recovery:
 - If consistent split without hot backup mode is used, bring the Oracle database resource online.
 - If consistent split with hot backup mode is used:
 - a. Make sure the Oracle service is running.
 - b. Connect to the database as SYSDBA using SQLPlus.
 - c. Run the following SQLPlus commands:


```
startup mount;
recover automatic database until cancel;
alter database open;
```
 - d. If using Oracle Failsafe, bring Oracle Failsafe service back online.
 - e. Bring the Oracle database resource online.

Restoring Oracle in a Solaris Sun Cluster

When Replication Manager restores to a Solaris Sun Cluster, it is necessary to suspend the cluster resource group prior to the restore operation to prevent the cluster from interfering with the Replication Manager restore operation. Follow these steps to manage the restore operation:

1. Move the Oracle resource group to suspended state. To achieve this, execute the following command as root or user with superuser privilege.

```
scswitch -s -g <resource_group_name>
```

Example: `scswitch -s -g oracle_res_grp`

2. Perform restore from Replication Manager.
3. Perform the database recovery, then open the database:

```
startup mount;
recover database;
alter database open;
```

Restoring Oracle RAC in an HP ServiceGuard environment

4. Remove the resource group from suspended state.

```
scswitch -r -g <resource_group_name>
```

Example: `scswitch -r -g oracle_res_grp`

To restore the Oracle RAC database in an HP ServiceGuard environment:

1. Shut down the Oracle RAC database on all nodes, either by logging in to each instance, or by using the following command:

```
srvctl stop database -d <database_name>
```

2. Restore the database on a given node in the cluster.
3. Change the owner and group of all the raw volumes associated with the Oracle database using the following command:

```
chown oracle:dba /dev/<vg_name>/r*
```

Note: Use the appropriate Oracle user and group.

4. Change the file permissions on each volume group directory:

```
chmod 777 /dev/<vgname>
```

5. Change the file permissions on all raw volumes associated with the Oracle database:

```
chmod 660 /dev/<vgname>/r*
```

6. Start up the database on all nodes in the cluster. If necessary, perform a database recovery on each node in the cluster.

Restoring Oracle RAC in a high availability multiprocessing (HACMP) environment

When you are restoring an Oracle database that resides on an HACMP/Oracle RAC cluster, you must perform some manual procedures in order to verify that the restore is successful:

Note: For the purposes of this discussion, the primary node is the node on which you plan to perform the restore and the secondary node(s) are one or more other nodes in the cluster.

1. Before you perform the restore, shut down Oracle on the secondary node(s).

On the primary node, perform step 2:

2. Perform the restore from Replication Manager. Refer to [“Restoring Oracle ASM disk groups” on page 385](#) for more information about working with ASM.

On the secondary node(s), perform step 3:

3. Through smit HACMP run HACMP Verification and Synchronization.

Perform step 4 on all nodes of the cluster:

4. Restart Oracle RAC.

Restoring Oracle ASM disk groups

When you create a replica of a database in an Oracle ASM environment, Replication Manager will not alert you if there is another database that you did not choose to replicate in the same disk group. If you later restore the replica, all databases in the disk group will revert to the point in time when the replica was taken. This may not be the desired effect. Additionally, if there are databases in the disk group that are not included in the replica, these will remain active and cause the disk group as a whole to remain active, which causes a restore of the replica to fail.

Replication Manager creates Oracle ASM replicas at disk group granularity, the replica includes both the database you chose to replicate and any other database that resides in the disk group. In other cases where this situation exists, Replication Manager warns you about the affected entity, but not in Oracle ASM environments at this time.

Replication Manager will take care of dismounting and remounting the appropriate ASM disk groups during the restore process unless you are running in a RAC-ASM environment on Linux. If that is your environment, refer to the next section for details.

Restoring Oracle RAC-ASM Replica on Linux

For Oracle running on Linux, Replication Manager cannot automatically unmount an ASM disk group globally to prepare for a restore when it is running in a RAC-ASM environment.

For example, if you have a three-node RAC database built on a shared disk group called DG1, Replication Manager can unmount DG1 from node 1, but not from node2 and node 3. Manually unmount the disk group as follows:

1. Manually shut down the RAC database instance on all nodes except the node on which Replication Manager will run restore: (In the example, RACDB2 is instance running on node 2, RACDB3 is instance running on node 3):

```
srvctl stop instance -d RACDB -i RACDB2
srvctl stop instance -d RACDB -i RACDB3
```

- If the ASM instances running on the RAC nodes only manage the disk groups involved in the restore, run these commands to shut down the ASM instances on the nodes to which data will not be restored (node 2 and node 3):

```
srvctl stop asm -n node2
srvctl stop asm -n node3
```

Note that ASM is not stopped on node 1 because node 1 is the target of the restore and ASM must be running on that node in order for Replication Manager to proceed.

If the ASM instances running on the nodes manage more disk groups than those involved in the restore, connect to each individual ASM instance (on node 2 and node 3) and manually dismount the affected disk groups:

```
alter diskgroup DG1 dismount;
```

Now only node 1 has DG1 mounted. If DG1 is mounted on any other node (besides the one on which the restore is occurring) the restore will fail with a message warning that the disk group needs to be dismounted from the other nodes of the RAC.

- Perform the Replication Manager restore.
- On node 1, perform the database recovery, then open the database on that node:

```
startup;
recover database;
alter database open;
```

- Restart ASM on the other nodes and start the RAC database globally:

```
srvctl start asm -n node2
srvctl start asm -n node3
srvctl start database -d RACDB
```

Restoring Oracle on clusters (general procedure)

When you are restoring an Oracle database that resides on a Cluster, you must perform some manual procedures in order to verify that the restore is successful.

If you have a standalone database in an Active/Passive Cluster that includes failover functionality, you should disable the failover mechanism prior to the restore and re-enable that mechanism once the restore is complete.

If you are running Oracle Real Application Clusters, you should:

1. Shut down the Oracle database on all nodes in the cluster via **srvctl** or individually on each node using SQLPlus.
2. Verify that the volume groups that contain the Oracle database volumes are offline.
3. Run the Replication Manager restore. Replication Manager automatically brings the volume groups back online, however, it does not bring them online in shared mode.
4. To remedy this, reissue the command to take the volume groups offline again.
5. Next, bring the volume groups back online, only this time, bring them online in shared mode.
6. Perform the tasks necessary to recover the database on each of the nodes. This might require you to restart the database using a command such as **srvctl** to verify that all the cluster functionality of the database is working properly.

Restoring Oracle from a RecoverPoint replica

Before you restore from a application-consistent (specific-point-in-time) RecoverPoint Oracle replica, the replica must be unmounted.

If you are planning on using the restore from a mounted crash-consistent (any-point-in-time) replica with RecoverPoint, note that the data being restored will be in the state it was in when the restore was initiated. For example, if the database is running on the mount host, this will result in restoring a crashed database over to the production host. You should shut down the Oracle database manually on the mount host before initiating the restore.

Mounting a Oracle point-in-time replica out-of-place may cause problems when mounting a subsequent any-point-in-time replica. If a problem does occur, remount the original point-in-time replica in-place, then try the any-point-in-time mount again.

Restore considerations on RAC aware standalone ASM databases

When you are restoring a standalone ASM database that resides on RAC aware Oracle, you must perform some manual procedures in order to initiate the restore operation.

For example, a standalone fail-over ASM database is created on shared disk groups, on node1. The other nodes in the cluster are node2 and node3. The following steps need to be performed for restore operation to succeed when initiated from node 1:

1. ASM instances running on all the other nodes should be shut down manually by running the command:

```
srvctl stop asm -n node 2
```

```
srvctl stop asm -n node 3
```

2. Perform Replication Manager restore.
3. On node 1, perform the database recovery by running the command:

```
recover database;
```

```
alter database open;
```

4. Restart ASM on the other nodes by running the command:

```
srvctl start asm -n node2
```

```
srvctl start asm -n node3
```

Using pre- and post-replication Oracle scripts

If you want to control how the database gets quiesced (into Online Backup mode) or shut down and restarted, you can write pre- and post-replication scripts to control these activities. If Replication Manager detects a pre-replication script, it leaves the task of quiescing the database to the script. The user-supplied pre-replication script must put the database in Online Backup mode or shut down the Oracle database.

If you specified a post-replication script, the script must either take the database out of Online Backup mode, or it must restart the database.

Online backups with user-supplied scripts

When the job specifies Online Backup mode, and there are user-supplied pre- and post-replication scripts, the Oracle Agent performs these steps:

1. Performs a log switch.
2. Runs the user-supplied pre-replication script.

The user-supplied script must put the database in online backup mode. When there are user-supplied scripts, the Oracle Agent will not put the database in online backup mode.

3. Checks that the database instance is in online backup mode; otherwise, the replication fails.
4. Backs up the control files.
5. Creates the replica.
6. Performs a log switch.
7. Runs the user-supplied post-replication script.

The user-supplied script should take the database out of Online Backup mode. When there is a user-supplied script, the Oracle Agent does not take the database out of Online Backup mode.

8. Saves the needed files and sends catalog information to the Replication Manager internal database.

Offline backups with user-supplied scripts

When the job specifies Offline Backup mode, and there are user-supplied pre- and post-replication scripts, the Oracle Agent performs these steps:

1. Runs the user-supplied pre-replication script.
The user-supplied pre-replication script must shut down the database or put selected tablespaces into Offline Backup mode. When there are user-supplied scripts, the Oracle Agent will not perform the shutdown.
2. Checks that the database instance is shut down or the tablespaces are offline before performing the replication. If the database instance has not been shut down, the replication fails.
3. Backs up the control files.
4. Creates the replica.
5. Runs the user-supplied post-replication script.
The user-supplied post-replication script should restart the database. When there is a user-supplied script, the Oracle Agent does not perform the restart.
6. Checks if the database has been started. If the database instance has not been restarted, the system logs an error and fails the replication.
7. Saves the needed files and sends catalog information to the Replication Manager internal database.

General guidelines for scripts

General guidelines for scripts are as follows:

- ◆ A script must be in an executable format (for example, .sh, .bat, .exe).
- ◆ Windows scripts must be either .bat or .exe format.
- ◆ UNIX scripts require their respective headers. For example, for sh, the script needs the header `#!/bin/sh` to work correctly.

- ◆ Suppress output from the script. The following techniques work for selected popular environments:

- **Windows** — Add @echo off to the first line of the .bat script.

- **UNIX sh** — Redirect output from the script as shown:

```
Prog.sh > /dev/null 2>&1
```

- **UNIX csh** — Redirect output from the script as shown:

```
Prog.csh >& /dev/null
```

Note: Suppressing script output prevents invalid characters from being added to replica history. Invalid characters in the history cause Replication Manager to stop responding when mounting or viewing a replica. If you need output from the script, set up and use a log file instead of directing output to standard out.

- ◆ Return a zero status to continue with the replication; return a nonzero status to fail the job.

Using pre- and post-replication scripts

To use a pre-replication and post-replication script:

1. Specify the name of the script and its full pathname when you configure the job in the console.
2. Set permissions on the script so that it is executable from the account you use to run Replication Manager.
3. Do not assume that any UNIX environment variables will be inherited by the script. Explicitly set environment variables such as PATH, ORACLE_SID, ORACLE_HOME, and LD_LIBRARY_PATH, and any other environment variables you might need.
4. Use the [“Sample Oracle pre-replication scripts \(offline\)” on page 391](#) and/or the [“Sample Oracle post-replication scripts \(offline\)” on page 393](#) as a model for your scripts.
5. Verify your scripts return a zero exit code. A nonzero exit code will cause the replication to stop with a status of failed pre- (or post-) script.

Sample Oracle pre-replication scripts (offline)

The following sample is a basic pre-replication script (with embedded SQL Plus commands) to bring the database instance offline in UNIX. Add the additional actions desired for your site:

```

#!/bin/sh
ORACLE_SID=jared
export ORACLE_SID
ORACLE_HOME=/jared
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:$PATH
export PATH
sqlplus /nolog << EOF
connect username/password as sysdba;
shutdown immediate;
startup restrict;
shutdown normal;
EOF
rc=$?
exit $rc

```

Note: If you run the Replication Manager Agent software in Secure mode, Replication Manager asks for user credentials (username and password) when you attempt to run a script. You should supply credentials for an Oracle user account that will be used to run all SQLPlus scripts. If you do not run it in Secure mode, the system runs the script as root.

You should return a nonzero exit code to instruct Replication Manager to fail the replication when errors occur during script processing. Use the next script to check for errors while running SQLPlus and to shut down the instance. The script:

- ◆ Connects, turns on spooling, shuts down the instance, and turns off spooling.
- ◆ Checks the exit code for the SQL command file, so if there is a failure in running SQLPlus, the replication also fails.
- ◆ Searches the spooled log file for error strings. If Replication Manager finds errors, you could add commands to the script that generate an SNMP event.

```

#!/bin/sh
ORACLE_SID=CAT815
export ORACLE_SID
ORACLE_HOME=/data/Oracle8.1.5/app/oracle/product/8.1.6
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:/usr/bin:$PATH
export PATH
rm -f /tmp/offline.log
sqlplus /nolog
@/data/Oracle8.1.5/app/oracle/product/8.1.6/scripts/offline_pre.sql
rc=$?

```

```

if [ $rc -ne 0 ]
then
##### Your choice if you want to abort the replication
because of the SQLPlus failure #####
##### Returning a non-zero exit code will cause the
replication to fail. #####
    exit $rc
fi
tmp=`egrep '^ORA-[0-9]*:' /tmp/offline.log`
if [ ! -z "$tmp" ]
then
    ##### Generate Your SNMP Event #####
fi
rm -f /tmp/offline.log
exit 0

```

Sample Oracle post-replication scripts (offline)

The following script in checks for errors running SQL Plus and in opening the database. The post-replication script:

- ◆ Calls a separate SQL command file to:
 - Connect to the Replication Manager Server.
 - Turn on spooling.
 - Start up the instance.
 - Turn off spooling.
- ◆ Checks the exit code for the SQL command file. If there is a failure in running Server Manager, you could choose to cancel the replication.
- ◆ Searches the spooled log file for error strings. If there are any, you could generate your own SNMP events.

```

#!/bin/sh
ORACLE_SID=CAT815
export ORACLE_SID
ORACLE_HOME=/data/Oracle8.1.5/app/oracle/product/8.1.6
export ORACLE_HOME
PATH=$ORACLE_HOME/bin:/usr/bin:$PATH
export PATH
rm -f /tmp/offline.log
sqlplus /nolog
@/data/Oracle8.1.5/app/oracle/product/8.1.6/
scripts/offline_post.sql
rc=$?
if [ $rc -ne 0 ]
then
##### Your choice if you want to abort the replication
because of
the SQL PLUS failure #####

```

```

exit $rc
fi
tmp=`egrep '^ORA-[0-9]*:' /tmp/online.log`
if [ ! -z "$tmp" ]
then
##### Generate Your SNMP Event #####
fi
rm -f /tmp/offline.log
exit 0

```

Sample Oracle pre-replication scripts (online)

The following sample is a basic pre-replication script (with embedded SQL Plus commands) to put the affected tablespaces in Online Backup mode. Add any actions desired for your site.

```

#!/bin/sh
connect <oracle_user>/<password>;
spool /tmp/online.log
alter tablespace <tablespace_name> begin backup;
alter tablespace <tablespace_name> begin backup;
alter tablespace <tablespace_name> begin backup;
spool off
exit 0

```

Sample Oracle post-replication scripts (online)

The following sample is a basic post-replication script (with embedded SQL Plus commands) to take the database instance out of Online Backup mode. Add any additional actions desired for your site.

```

#!/bin/sh
connect <oracle_user>/<password>;
spool /tmp/online.log
alter tablespace <tablespace_name> end backup;
alter tablespace <tablespace_name> end backup;
alter tablespace <tablespace_name> end backup;
spool off
exit 0

```

Oracle Real Application Cluster (RAC) commands can also be used in these Replication Manager scripts. For more information about RAC commands, refer to your Oracle documentation.

Using the root user to perform Oracle operations

EMC does not recommend that you use the root user to perform Oracle operations. However, to provide backward compatibility with previous versions of Replication Manager, this capability exists. If you decide to continue to use the root user, the following special configuration steps apply to you:

1. Locate (or create) an Oracle user for Replication Manager to use. The Replication Manager user should be granted `SYSDBA` and `DBA` privileges.
2. Verify that the root user has read access on the following directory to facilitate database rename operations:

```
$ORACLE_HOME/rdbms/lib
```

3. Verify that the root and daemon user have been added to the Oracle DBA group on the mount host.
4. Verify the `/etc/group` file has the daemon user and Oracle user listed in the DBA group. Also, verify that the Oracle user is the last user listed.

For example, if `oracle` is the user and `oradba` is the group, the `/etc/group` entry should look like:

```
oradba:!:<GID>:daemon,oracle
```

5. Verify that the Oracle binary on the mount host is set with the file permissions `6751`. For example, you can type the following command:

```
chmod 6751 oracle
```

6. Verify that all of the Oracle binaries are owned by the correct combination of Oracle user and group. In the example above, we would set the group to `oradba` and the user to `oracle`.

If the binaries are owned by the wrong owner or belong to the wrong group:

- a. Check whether `$ORACLE_HOME/rdbms/lib/config.c` file exists and verify that the owner listed in the `SS_OPER` entry of that file is the Oracle user defined in the `/etc/group/` file.
- b. If you change the user in the `SS_OPER` entry, you must generate a `config.o` file from the `config.c` file and relink the Oracle binaries with a **relink all** command.
- c. Verify that all directories that make up `$ORACLE_HOME` grant write permission to the Oracle user and group.
- d. For example, if `$ORACLE_HOME` starts with the directory `/u01/...` Use the following command to set the file permissions:

```
chmod -R 771 /u01
```

7. Verify that the Oracle daemon user has write access to `$ORACLE_HOME/dbs/`.

Storage Foundation for Real Application Clusters

This section describes Replication Manager support for Storage Foundation for Real Application Clusters (SFRAC). SFRAC is an integration of Veritas Cluster Services (VCS), Veritas Cluster File System (VCFS), and Oracle's Real Application Cluster (RAC).

In previous releases of Replication Manager, the replication of an SFRAC environment would be successful but mounts and restores would fail. Mounts failed because cluster file system attributes did not match the expected non-clustered file system attributes of the non-clustered mount host's file system. Restore failed for the same reason.

In this release, Replication Manager allows a mount of an SFRAC replica to a non-clustered host. Replication Manager filters out the cluster file system attributes to accomplish this. On restore, the attributes that have been preserved will match the environment to which restore is occurring, provided that Replication Manager is restoring to the same node from which the replica was created.

Prerequisite: SSH must be configured

SSH must be configured to communicate with other nodes during restore of a replica in an SFRAC environment. Verify that SSH is enabled and uses password-less authentication between production cluster nodes.

Environment variables for SFRAC

The following environment variable changes the default behavior that might be desirable for some customers. Set the variables in the Replication Manager Console, under Host properties, Advanced tab.

`ERM_VCSSFAC_RSH`

Use this environment variable on Solaris/VCS/SFRAC environments to use RSH for remote node operations. Without this environment variable Replication Manager will automatically use SSH for remote SFRAC node operations.

Set the variable to any value to enable RSH. Verify that RSH is enabled and does not require any password-based authentication when using this environment variable. Replication Manager also uses

the `erm_vxImportRestore` script as part of VG import operation. This script looks for this environment variable and executes remote commands using SSH or RSH.

Prerequisite: SCSI-3 PR must be enabled

For SFRAC support, SCSI-3 Persistent Reservation (PR) must be enabled. SCSI-3 PR lets multiple cluster nodes access a device while at the same time blocking access to other nodes. SCSI-3 PR reservations are persistent across SCSI bus resets or node reboots, and also support multiple paths from host to disk.

CLARiiON and VNX arrays use SCSI-3 PR as a default.

To enable SCSI-3 PR in a Symmetrix environment:

1. On the Solutions Enabler host, update the `configure.txt` file to include the following command:

```
set device 0090 attribute=SCSI3_persist_reserv;
```

2. From the command line, run these commands:

```
symconfigure -sid <sid_name> -f configure.txt preview  
symconfigure -sid <sid_name> -f configure.txt commit
```

3. Verify the change using the following command:

```
symdev -sid <sid_name> show 0090
```

4. Look for the text "SCSI-3 Persistent Reserve: Enabled" among the output.

Mount considerations in SFRAC environments

Replication Manager can create replicas of data protected by this configuration. Mount hosts must be standalone systems. Replication Manager does not support mounting an SFRAC replica to a cluster.

Restore considerations in SFRAC environments

In previous releases, Replication Manager supported VCS environments with Oracle RAC on raw volumes. Restoring in that environment required several manual steps and restores could only take place to the same host from which the replica was taken. In addition, that host was required to be the master node of the cluster.

Now Replication Manager supports a complete SFRAC environment, which extends raw volumes support with Veritas Clustered File System (VCFS). Replication Manager also automates much of the process, removing the manual steps and the restriction that required the restore to occur on the master node. Replication Manager continues to require SFRAC restores to occur on the same host from which the replica was taken.

When restoring a replica to an SFRAC environment, Replication Manager performs the following steps automatically:

1. Replication Manager shuts down the Oracle instance.
2. Replication Manager unmounts the file system on all nodes. If the restore occurs on a slave node, Replication Manager uses remote commands via SSH (by default) or RSH (if the appropriate environment variable is set).
3. Replication Manager deports the volume group on all nodes. If the restore occurs on a slave node, Replication Manager uses remote commands via SSH (by default) or RSH (if the appropriate environment variable is set) to communicate with the master node remotely and deport the volume group.
4. Replication Manager restores the replica.
5. Replication Manager imports the volume group on all nodes. If the restore occurs on a slave node, Replication Manager uses remote commands via SSH (by default) or RSH (if the appropriate environment variable is set) to communicate with the master node remotely and import the volume group.
6. Replication Manager mounts the file system on all nodes.

Replication Manager does not need to be installed on the remote node for restore to work. Replication Manager uses cluster commands to complete this work. Additionally, if the remote node happens to be down at the time, the process continues because by default, if the remote node is down there is nothing to unmount or deport.

If you mount with the Read-Write recovery or Database Rename option with an alternate path, the replica is altered even after unmount; such a replica should not be used for a restore operation.

Restore steps in SFRAC configuration

During a restore, you need to freeze the Oracle service group to prevent VCS from interpreting the shutdown as a database failure. The required steps are:

1. Freeze the service group.
2. Run Replication Manager restore and database recovery.
3. Once the restore is complete, unfreeze the service group.

4. Clear any failed resources.
5. Bring resources online:

`ERM_SFRAC_NOSLAVE`

Use this environment variable if Replication Manager should not continue with the VCS/SFRAC restore operation if the node where restore is done is a slave node. You can set this variable in the Replication Manager Console under Host properties, Advanced tab.

Without this variable, Replication Manager will detect the master node and perform VG deport and import operations on that node using remote commands.

If you do not want Replication Manager to take the default action, then set this environment variable to any value. This may help a situation where the Veritas disk names are different across master and slave nodes and remote operations are not likely to succeed, because Replication Manager does not have knowledge of remote Veritas disk names. Note that when this variable is set, it prevents restores of replicas unless the host is the master node.

Oracle troubleshooting

Certain situations can cause issues that can be remedied by making changes to the configuration. The following sections describe troubleshooting tips. Consult these sections before calling EMC Technical Support.

TNS permission denied at application set creation

If a “TNS permission denied” error is displayed during application set creation, it means that the permissions are not set correctly on the Oracle home directory structure. To correctly set the permissions:

1. Shut down the Oracle instances and ASM instances on the host.
2. Add root to the UNIX “dba” group.
3. Go to the top level of the Oracle directory structure and run the following command:

```
chmod -R 775
```

4. Run the following commands to set permissions on the binaries directory:

```
cd $ORACLE_HOME/bin/oracle  
chmod 6751
```

5. Log in as the oracle user (su - <oracle_user>).
6. Run **relink all**.
7. Restart the ASM and Oracle binaries.

Cannot find log archive destination

Verify that the log_archive_dest_<n> or log_archive_dest parameters point to a physical directory.

Note: From Replication Manager version 5.2.1 and later, archive log prefixing is supported.

Oracle 10g allows administrators to make use of an intermediary parameter to specify the archive log destination, as follows:

```
db_recovery_file_dest = "/recovery_area/"
```

```
log_archive_dest_1 = 'LOCATION=USE_DB_RECOVERY'
```

The final slash terminating the `log_archive_dest_<n>` or `log_archive_dest` parameters is no longer a requirement for Replication Manager environments. As of Replication Manager version 5.2, a slash terminated directory is no longer required

Unable to mount and recover Oracle replica on raw Linux devices

If you mount a replica of an Oracle database that was created on raw devices in No Recover mode, and subsequently mount the same database replica and attempt to recover the database, there is a chance that the second replica mount will fail with an error such as:

```
ERROR: Could not mount and recover the database.
```

To prevent this issue:

1. Mount the replica again in No Recover mode (instead of attempting to recover the database).
2. Make a note of all the names of raw data and redo devices listed on the Console (or in the Replication Manager client logs.) The logs contain the block devices that Replication Manager uses during the mount.
3. Run the utility **raw -qa** to check all the raw device bindings already on the mount host.
4. For each raw device referenced in the logs, and the associated block device, run the following command:

```
raw <raw_device_full_name> <block_device_name>
```

For example:

```
raw /dev/raw/raw1 /dev/sdab1
```

5. Perform the following steps within Oracle to mount the replica using the Replication Manager generated pfile in /tmp:
 - a. Log in to Oracle.
 - b. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables with the appropriate values.
 - c. Use SQLPlus to mount the replica by typing commands such as the following:

```
startup mount pfile=/tmp/init<SID>.ora
```

6. Rename the files as outlined in the client log. Search the log for instructions similar to the following:

```
11/2/04 8:39:31 AM You must create the Oracle
password file in /u01/app/oracle/product/9.2.0/dbs
to complete the recovery.
```

```
11/2/04 8:39:31 AM 000255 INFO: Use the following
init file to start your instance:
/tmp/initRACLN1.ora.11/2/04 8:39:31 AM You must
execute - alter database rename file /dev/raw/raw15
to /dev/raw/raw27 in order to complete the
recovery.
```

```
11/2/04 8:39:31 AM You must execute - alter
database rename file /dev/raw/raw16 to
/dev/raw/raw28 in order to complete the recovery.
```

```
11/2/04 8:39:31 AM You must execute - alter
database rename file /dev/raw/raw17 to
/dev/raw/raw29 in order to complete the recovery.
```

```
11/2/04 8:39:32 AM To ensure proper cleanup shut
down this Oracle instance when you are done working
with it.
```

```
11/3/04 4:29:23 PM Mount of the replica created at
Nov 2, 2004 11:51:22 AM, successfully completed at
time Nov 3, 2004 4:29:23 PM.
```

Now you should be able to manually open the Oracle database.

Database name in the control file does not match

Remounting a database replica that has previously been mounted and renamed requires that you use the database rename feature again.

Attempting a mount without the proper rename option produces an error similar to the following:

```
ERROR at line 1: ORA-01103: database name 'abc' in
control file is not 'xyz'.
```

Restores are not advisable after a database rename operation is completed on a replica, because the data being restored corresponds to an entirely different database.

Replication fails with a VSS error running as domain user

Under certain circumstances, Replication Manager replication jobs can fail with an error similar to the following:

```
027288 ERROR:A VSS_E_UNEXPECTED_PROVIDER_ERROR error
occurred while adding volumes to the VSS snapshot set.
The error code is: xxxxxxxxxxxx
```

This issue occurs if NTS authentication is used for the instance (that parameter is defined in `sqlnet.ora`). If that mode is used, then the operating system user running the Replication Manager Agent must be in the `ORA_DBA` group. If the NTS is not used, then this restriction does not apply.

Mounting and recovering an Oracle replica created on raw devices in Linux

If you mount a replica of an Oracle database that was created on raw devices in No Recover mode, and subsequently mount the same database replica and attempt to recover the database, there is a chance that the second replica mount will fail with an error such as:

```
ERROR: Could not mount and recover the database.
```

To prevent this issue:

1. Mount the replica again in No Recover mode (instead of attempting to recover the database).
2. Make a note of all the names of raw data and redo devices listed on the Console (or in the Replication Manager client logs.) The logs contain the block devices that Replication Manager uses during the mount.
3. Run the utility **raw -qa** to check all the raw device bindings already on the mount host.
4. For each raw device referenced in the logs, and the associated block device, run the following command:

```
raw <raw_device_full_name> <block_device_name>
```

for example,

```
raw /dev/raw/raw1 /dev/sdab1
```

5. Perform the following steps within Oracle to mount the replica using the Replication Manager generated pfile in `/tmp`:
 - a. Log in to Oracle.
 - b. Set the `ORACLE_SID` and `ORACLE_HOME` environment variables with the appropriate values.

- c. Use SQLPlus to mount the replica by typing commands such as the following:

```
startup mount pfile=/tmp/init<SID>.ora
```

6. Rename the files as outlined in the client log. Search the log for instructions similar to the following:

```
11/2/04 8:39:31 AM You must create the Oracle
password file in /u01/app/oracle/product/9.2.0/dbs
to complete the recovery.
```

```
11/2/04 8:39:31 AM 000255 INFO: Use the following
init file to start your instance:
/tmp/initRACLN1.ora.11/2/04 8:39:31 AM You must
execute - alter database rename file /dev/raw/raw15
to /dev/raw/raw27 in order to complete the
recovery.
```

```
11/2/04 8:39:31 AM You must execute - alter
database rename file /dev/raw/raw16 to
/dev/raw/raw28 in order to complete the recovery.
```

```
11/2/04 8:39:31 AM You must execute - alter
database rename file /dev/raw/raw17 to
/dev/raw/raw29 in order to complete the recovery.
```

```
11/2/04 8:39:32 AM To ensure proper cleanup shut
down this Oracle instance when you are done working
with it.
```

```
11/3/04 4:29:23 PM Mount of the replica created at
Nov 2, 2004 11:51:22 AM, successfully completed at
time Nov 3, 2004 4:29:23 PM.
```

Now you should be able to manually open the Oracle database.

Failure of Oracle mount with recovery on Linux, "Error in renaming log/datafiles" messages

If an Oracle mount with recovery on Linux fails, and messages such as "error in renaming log/datafiles" appear in the log, the problem may be caused when new bindings on the mount host have names that collide with the old bindings on the production host.

The workaround is to prebind the raw devices on the mount host using raw device numbers that are higher than the range used on the production host. For example, if the production host uses /dev/raw/raw1 to /dev/raw/raw100, the mount host should use /dev/raw/raw200 to /dev/raw/raw300.

This error condition exists only for Oracle and Oracle RAC when ASM is not in use.

Mount fails with error “Cannot identify control files”

This error occurs when the primary group for the OS username provided does not match on both production and mount host (GIDs must match). The username/UIDs can be different. If you get this error, try modifying the group in order to ensure that the GIDs match on the production and mount OS user.

Recover mount of Oracle replica fails when SGA is large

Recover mount of an Oracle replica can fail when the System Global Area (SGA) is very large. Refer to the Oracle MetaLink Doc ID 363178.1 for more information.

Replication Manager protects IBM Enterprise Edition DB2 Universal Databases (UDB) as well as single-node clusters that use the IBM Enterprise Edition DB2 Universal Database. Replication Manager protects these databases by creating and managing application sets that contain entire UDB databases. This appendix includes the sections that cover these subjects:

- ◆ Configuring the UDB environment (UNIX)..... 410
- ◆ Understanding UDB application sets and replicas 414
- ◆ Mounting and restoring UDB replicas 418
- ◆ Using pre- and post-replication UDB scripts 424
- ◆ UDB troubleshooting..... 428

Configuring the UDB environment (UNIX)

Before using Replication Manager with UDB, ensure that the UDB environment is configured properly. The following checklist can help you configure the UDB environment to work with Replication Manager:

- ❑ Collect important UDB information necessary to configure the system for use with Replication Manager. See the detailed instructions in [“Collecting important UDB information” on page 412](#).
- ❑ Ensure that the UDB agent successfully located the UDB package. If Replication Manager did not find the appropriate location, follow the steps in [“UDB agent unable to locate the UDB package” on page 428](#).
- ❑ The UDB agent requires that the UDB database being replicated have archive logging and roll forward recovery enabled. This can be done with LOGRETAIN set to RECOVERY and USEREXIT set to ON, or some other method.
- ❑ Ensure that all UDB 9.x databases being replicated were created using the AUTOMATIC STORAGE NO clause to ensure that the database does not have automatic storage characteristics. Replicas of databases with automatic storage characteristics experience mount issues during containers relocation.
- ❑ In UDB 9.x installations, Replication Manager requires /usr/local/bin/ to contain the db2ls file that the UDB installation creates by default. Failure to have db2ls installed in the default location will prevent Replication Manager from locating the UDB installation path and performing any UDB operation.
- ❑ Users of UDB 9.7 should set the environment variable URM_UDB_VERSION to 90700000 in /opt/emc/rm/client/bin/rc.irclient under the section “Set the environment variables as expected by IR”.
- ❑ Ensure that tablespaces, active logs, and the database directory are located on a supported storage array.
- ❑ Ensure that each mount host has been setup as described in [“UDB mount host setup” on page 418](#).

- ❑ If you run the `irccd` daemon in Secure mode, Replication Manager asks for user credentials (username and password) when you attempt to run a script. The credentials should match a UDB user account that will be used to run all scripts. If you do not run it in Secure mode, the system runs the script as root even though it asks for user credentials.
- ❑ You must perform special configuration steps if you plan to mount different replicas of the same UDB database to the same host simultaneously. Details of the steps to manage this configuration are outlined in [“UDB mount host setup” on page 418](#).

A typical UDB environment

Replication Manager can replicate IBM Enterprise Edition DB2 Universal Databases (UDB) that have a standard configuration. This section illustrates the supported configuration and highlights the components of such a configuration. [Figure 156 on page 411](#) shows a standard UDB configuration.

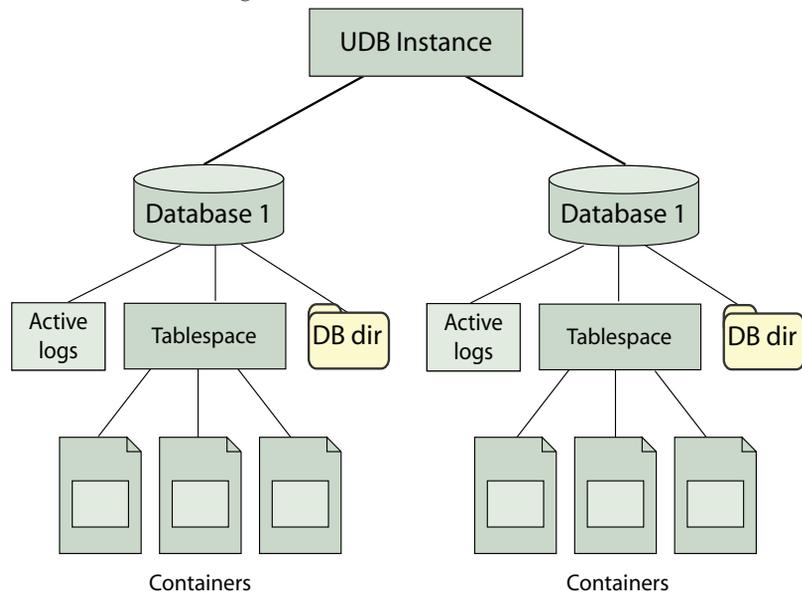


Figure 156 UDB database configuration

Replication Manager can create replicas of UDB databases and can mount and restore those replicas. Also, Replication Manager can rename an instance and/or a database when mounting the UDB replica.

Replication Manager requires all of the following objects to reside on a supported storage array in order to replicate the UDB information:

- ◆ All containers in each tablespace
- ◆ Active logs
- ◆ Database directory

All of these components are part of the replica.

Collecting important UDB information

Before you create a UDB replica, you must collect some important information such as the list of instances or databases. To collect the important UDB information needed:

1. Start a UDB DB2 CLI session.
2. Use the commands in [Table 28 on page 412](#).

Table 28 Accessing UDB information

Information needed	Sample command
DB Directory	list db directory for <db name>
LOGRETAIN (search for LOGRETAIN in output) LOGRETAIN should be set to RECOVERY	Get db cfg for <dbname>
USEREXIT (search for USEREXIT in the output, if LOGRETAIN is not set, this should be turned ON.)	
LOGPATH (search for 'Path to log files' in output)	
LOGARCHMETH1 and other log archive methods	
List instances	db2ilist
List databases	db2 list db directory

Table 28 **Accessing UDB information**

Information needed	Sample command
Show database configuration	db2 get db cfg for <i><database></i>
Show tablespace locations	db2 list tablespace containers for <i><id></i> show detail
Test username and password for an instance	db2 attach to <i><instance></i> user <i><username></i> using <i><password></i>

Understanding UDB application sets and replicas

When creating an application set for the UDB agent, users can specify only whole databases for replication. UDB replications copy all the tablespaces and containers within each selected database, and if tablespaces are added or deleted, the future replicas adjust appropriately.

Before each replication, Replication Manager:

- ◆ Discovers the location of the data to replicate
- ◆ Identifies the pathnames for all the containers in the requested databases
- ◆ Identifies the active log path
- ◆ Identifies the location of the database directory

All of these components must be located on supported storage arrays or the replication process fails.

IBM UDB instances that reside on Symmetrix systems can be included as part of an application set. Some restrictions apply and special configuration steps may be required. [Chapter 7, "Configuring Federated Data,"](#) provides full details on how to create IBM UDB federated application sets.

UDB online and offline replication

When configuring a job for a UDB application set, the user can choose between the following two consistency methods:

- ◆ Online by suspending writes to the database
- ◆ Offline by shutting down the database

The software adjusts its replication processing accordingly to stop the database (if necessary).

[Table 29 on page 415](#) and [Figure 156 on page 411](#) describe and illustrate UDB replication processing.

Table 29 UDB online and offline replication steps

Online replication	Offline replication
1. Automatically discovers the tablespaces, database directories, and active log locations for the databases you are replicating.	Automatically discovers the tablespaces, database directories, and active log locations for the databases you are replicating.
2. Establishes mirrors (if necessary).	Establishes mirrors (if necessary).
3. User-supplied pre-replication script runs (script must put the database into write-suspend mode). If there is no pre-replication script, Replication Manager puts the database into Write Suspend mode without running a script.	User-supplied pre-replication script runs (script must shut down the database). If there is no pre-replication script, Replication Manager shuts down the database without running a script.
4. Creates the replica of the tablespaces, database directory, and the active logs.	Creates the replica of the tablespaces, database directory, and the active logs.
5. User-supplied post-replication script runs (script must take the database out of Write Suspend mode). If there is no post-replication script, Replication Manager takes the database out of Write Suspend mode without running a script.	User-supplied post-replication script runs (script must restart the production database). If there is no post-replication script, Replication Manager restarts the production database without running a script.
6. Catalogs the following information: <ul style="list-style-type: none"> • DB directory location • Location of files in DB dir • Location of logs • User and group IDs 	Catalogs the following information: <ul style="list-style-type: none"> • DB directory location • Location of files in DB dir • Location of logs • User and group IDs

UDB replication steps

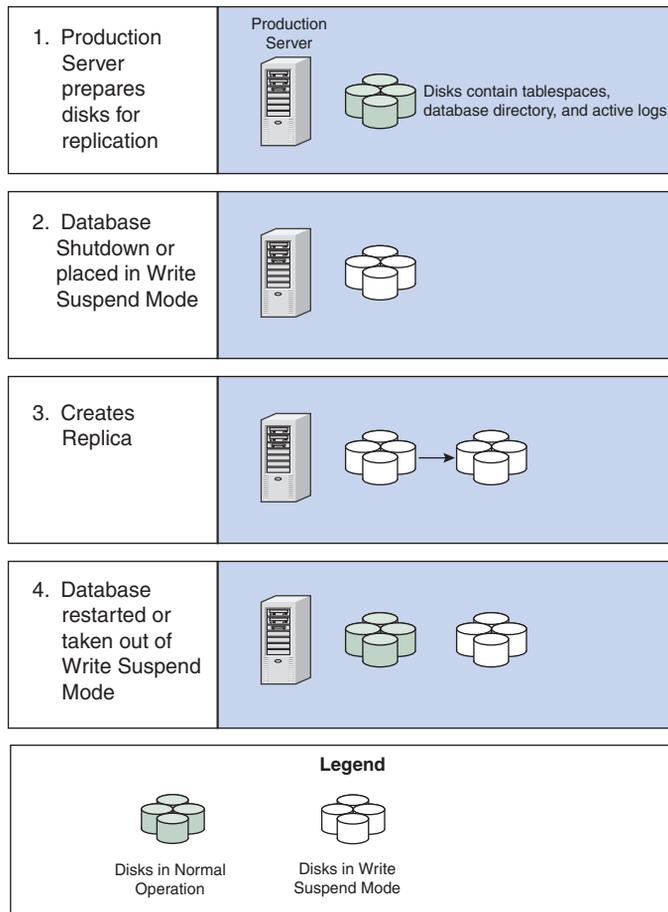


Figure 157 UDB replication steps

Figure 157 on page 416 illustrates the replication steps. Controlling the relocate configuration file you can set an environment variable to control where Replication Manager should put the relocate configuration file, if one is needed. The relocate configuration file is used to manage any changes to the location of various UDB components as a result of a mount. If you do not set the environment variable, the relocate configuration file is stored in the default temp directory.

To store the file elsewhere, set the following environment variable:

```
ERM_TEMP_BASE=<directory_name>
```

where *<directory_name>* is the base directory where the relocate configuration files will be stored. Replication Manager creates the relocate configuration file for each database in this directory.

Mounting and restoring UDB replicas

The UDB agent enables Replication Manager to initiate and control mounts and restores of UDB databases.

UDB mount host setup

The UDB agent can mount a replica to an alternate host as long as the following conditions are met:

- ◆ The mount host has at least one UDB instance installed and running.
- ◆ The mount host has the Replication Manager Agent software installed.
- ◆ The mount host is registered with the Replication Manager Server.
- ◆ The production and mount hosts have identical operating system, volume manager, file system, and application versions installed.
- ◆ The mount host has a UDB instance user ID and Group ID that matches the user ID and Group ID used in the UDB instance on the production server. Using an alternate instance or using the File system Mount option removes this restriction.
- ◆ If you are mounting to the production host, select alternate instance name and alternate path for mount.
- ◆ During the mount operation, the selected UDB instance has to be running.
- ◆ If you are mounting different UDB replicas of a single database, or two different UDB databases with the same name to the same host, you must specify a different instance for each database; otherwise, the mount of the second database will fail.
- ◆ If you change the name of the UDB instance on the mount host, keep in mind that the total UDB instance name cannot exceed 8 characters.

UDB mount types

If the UDB environment is set up as described above, Replication Manager can perform mount operations in the following ways:

- ◆ Immediate mounts on new replicas as part of a job.
- ◆ Mounts to alternate locations or mounts back to the production paths on a separate mount host.

- ◆ Mounts to alternate locations on the production host.
- ◆ Mounts of a single node of a UDB Enterprise Edition cluster on another nonclustered UDB instance (separate from the cluster).

Note: Specific information about how alternate paths and path mapping work can be found at “Mounting using alternate path” on page 174 or “Mounting using path mapping” on page 177.

The UDB agent allows users to mount a replica containing a UDB instance or database in one of the following three ways:

- ◆ File System mounts
- ◆ Prepare Only mounts
- ◆ Snapshot mounts

Figure 158 on page 419 shows the Mount Options panel.

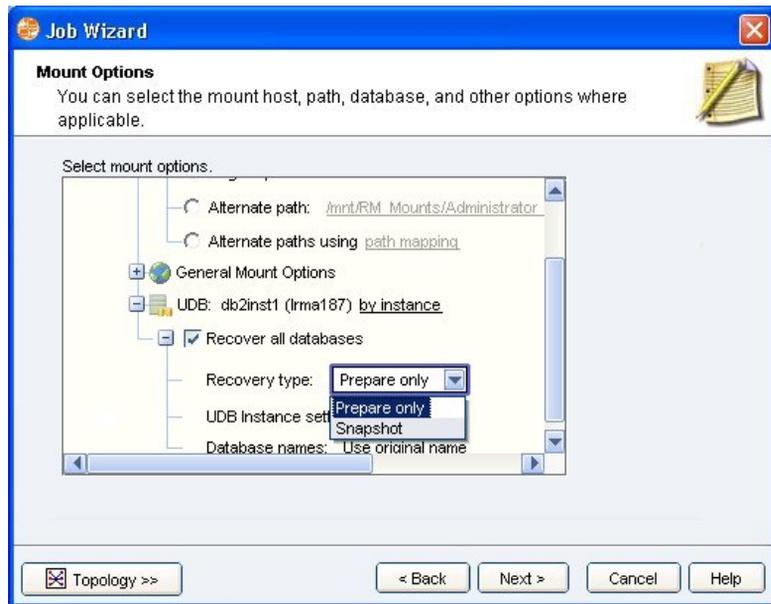


Figure 158 UDB database mount options

UDB File System mounts

Clear the Recover the database or Recover all databases checkbox to choose a File System mount. During a File System mount, Replication Manager performs the following steps:

- ◆ Makes all replica devices visible to the mount host
- ◆ Imports all volume groups
- ◆ Mounts all file systems associated with the replica

UDB Prepare Only mounts

Select the Recover the database or Recover all databases checkbox and choose the Prepare Only recovery type to perform a Prepare Only mount. Prepare Only mounts can be used for backup purposes or to create a separate copy of the database for repurposing or cloning.

During a Prepare Only mount, Replication Manager performs the following steps:

- ◆ Performs all the actions described above for File System mount.
- ◆ Creates the relocate configuration file, which includes the old and new locations of the file system paths, instance name, database name, log directory, and container paths. [“Sample relocate configuration file” on page 420](#) illustrates a sample relocate configuration file.
- ◆ Lists the steps (on the Replication Manager Console) that a user should perform in order to complete a manual recovery of the UDB database. These steps are also described in [“Manual recovery of a Prepare Only mount” on page 421](#). The relocate configuration file for a mount using path mapping is shown in [“Sample relocate configuration file” on page 420](#).

Sample relocate configuration file

```
DB_NAME = ORIG_DBNAME, NEW_DBNAME
INSTANCE = ORIG_INSTNAME, NEW_INSTANCE
DB_PATH = /ORIG_DBDIR, /NEW_DBDIR
NODENUM = 0
LOG_DIR =
/udbhome/db2inst1/NODE0000/SQL00001/SQLLOGDIR/, /alt/udbhome/db2inst1/NODE0000/SQL00001/SQLLOGDIR/
CONT_PATH=/udbhome/sysdb1, /altpath1/udbhome/sysdb1
CONT_PATH=/udb_on_syymm466/tempdb1,
/altpath2/udb_on_syymm466/tempdb1
CONT_PATH=/udb_on_syymm466/userdb1,
/altpath3/udb_on_syymm466/userdb1
```

Note: Prepare only mounts on a secured host require you to enter valid user credentials.

Prepare Only mounts can be performed as many times as you like as long as the database has not been recovered manually.

Note: Replication Manager does not save the files in the database directory during an unmount from a Prepare Only mount.

Manual recovery of a Prepare Only mount

To recover a Prepare Only mount manually, follow one of these procedures:

- ◆ *To recover a Prepare Only mount in Standby mode for backup, start a UDB Instance on the mount host, start a session and issue the following command:*

```
Db2inidb <db name> as STANDBY relocate using <RM generated relocate cfg file name>
```

- ◆ *To recover a Prepare Only mount in Snapshot mode for repurposing, start a UDB Instance on the mount host, start a session and issue the following command:*

```
Db2inidb <db name> as SNAPSHOT relocate using <RM generated relocate cfg file name>
```

Note: If you recover in Snapshot mode using this command, EMC recommends that you do not restore the replica because of the new log chain created.

UDB Snapshot mounts

Select the **Recover the database** or **Recover all databases** checkbox and choose the **Snapshot** recovery type to perform a Snapshot mount.

Snapshot mounts can be used to create a crash recovery copy of the database on the mount host. The database is fully available after the mount. Users can mount using Snapshot mode more than once because the files in the database directory that are changed during the first mount are part of the replica. When future mounts occur, Replication Manager uses those same files in the database directory because they are synchronized with the replica data stored on the storage array.

During a Snapshot Mount, Replication Manager performs the following steps:

- ◆ Performs all the actions described above for Prepare Only mounts.
- ◆ Automatically executes the following command to recover the database in snapshot mode:

```
Db2inidb <db name> as SNAPSHOT relocate using <RM generated relocate cfg file name>
```

Note: If you mount using Snapshot mode, EMC recommends that you not restore the replica because of the new log chain created.

UDB restores

When restoring a UDB replica, choose what database(s) to restore. When you restore the a database, the UDB agent shuts down the database prior to completing the restore. Refer to [Figure 159 on page 422](#).

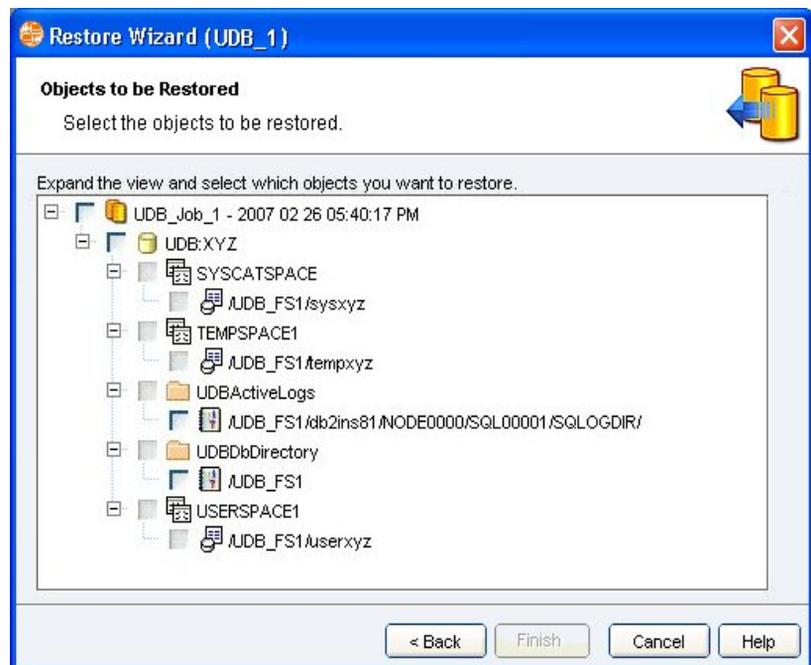


Figure 159 UDB restore wizard

Note: If a replica has been mounted using the snapshot mode EMC does not recommend that you restore that replica.

Users can use the replica to restore the production database and roll forward using the logs on the production system. Consider the following before performing a restore operation:

- ◆ Replication Manager works at the volume group level. By choosing an object for restore, all objects that share a volume group with the selected object will be restored to the production host.
- ◆ Although it is an available option, you should use care when selecting only the DB directory and active logs for the restore.
- ◆ Make sure that the DB directory and/or the active logs do not reside on the same volume group(s) as the tablespaces.
- ◆ If the DB directory and/or the active logs share one or more volume groups with the tablespaces, then save your original DB directory and active logs to a different location before restoring the replica.

Once the tablespaces are restored:

1. Start the database in MIRROR mode using the following command:

```
Db2inidb <db name> as MIRROR relocate using <RM  
generated relocate cfg file name>
```

2. This command brings up the database in Roll Forward pending state, Use the appropriate **ROLL FORWARD DATABASE** command to roll forward and recover the database to the desired point in time.

Using pre- and post-replication UDB scripts

If you supply pre- and post-replication scripts as part of your UDB job, the system works differently than if you omit those scripts. The following sections describe the behavior of Replication Manager when scripts are specified.

Offline replications with user-supplied scripts

When the job specifies offline replication, and there are user-supplied pre- and post-replication scripts, the UDB agent does the following:

1. If the user has supplied a pre-replication script, Replication Manager runs that script.

Any pre-replication script the user specifies must shut down the database. When there are user-supplied scripts, the UDB agent does not perform the shutdown. If the user does not specify a script, the UDB agent performs the shutdown automatically.
2. Checks that the database is shut down before performing the replication. If the database has not been shut down, the replication fails.
3. Creates the replica.
4. If the user has supplied a post-replication script, run that script. The script should restart the database. When there is a user-supplied script, the UDB agent does not perform the restart.
5. Checks if the database has been started. If the database has not been restarted, Replication Manager logs an error and the replication fails, Replication Manager cleans up the database and puts that database back into its original mode.

Using pre- and post-replication scripts

To use a pre-replication and post-replication script:

1. Specify the name of the script and its full pathname when you configure the job in the console.
2. Set permissions on the script so that it is executable by the user from which you run Replication Manager.
3. Do not assume that any UNIX environment variables will be inherited by the script. Explicitly set `PATH`, `DB2INSTANCE`, `DB2DIR`, and other UDB-specific environment variables. You should also set the `LD_LIBRARY_PATH` and any other environment variables you might need.

Suppress output from the script. The following techniques work for selected popular environments:

- **Windows** — Add `@echo off` to the first line of the `.bat` script.
- **UNIX sh** — Redirect output from the script as shown:

```
Prog.sh > /dev/null 2>&1
```

- **UNIX csh** — Redirect output from the script as shown:

```
Prog.csh >& /dev/null
```

Note: Suppressing script output prevents invalid characters from being added to replica history. Invalid characters in the history cause Replication Manager to stop responding when mounting or viewing a replica. If you need output from the script, set up and use a log file instead of directing output to standard out.

4. Sample scripts are provided in the following sections. Use these scripts as a starting point for site-specific customization. You need to define scripts only if you need to perform additional steps that Replication Manager does not perform automatically. Otherwise, you can choose to let Replication Manager shut down and start up the database automatically by omitting scripts completely.

Sample UDB pre-replication scripts (online)

This script is a basic pre-replication script to place a UDB database in Write-Suspend mode. Add the additional actions desired for your site.

```
#!/bin/sh
id
DB2INSTANCE=udbinst
export DB2INSTANCE
```

```
/udb/udbinst/sqllib/bin/db2 connect to MYUDB user udbinst
using password
/udb/udbinst/sqllib/bin/db2 set write suspend for
database
exit 0
```

Note: If you run the `ircdd` daemon in Secure mode, Replication Manager asks for user credentials (username and password). The credentials should match a UDB user account that will be used to run all scripts.

In addition to the standard script processing, you should ensure your script can handle errors that may occur while it is running. If errors occur in your script, you should fail the replication by returning a nonzero return value from the script.

**Sample UDB
post-replication
scripts (online)**

This script takes the database out of Write-Suspend mode.

```
#!/bin/sh
id
DB2INSTANCE=udbinst
export DB2INSTANCE
/udb/udbinst/sqllib/bin/db2 connect to MYUDB USER udbinst
using password
/udb/udbinst/sqllib/bin/db2 set write resume for database
exit 0
```

**Sample UDB
pre-replication scripts
(offline)**

The script is a basic pre-replication script to take a UDB database offline. Add any additional actions desired for your site.

```
#!/bin/sh
id
DB2INSTANCE=udbinst
export DB2INSTANCE
/udb/udbinst/sqllib/bin/db2 connect to MYUDB user udbinst
using password
/udb/udbinst/sqllib/bin/db2 force applications all
/udb/udbinst/sqllib/bin/db2 disconnect current
/udb/udbinst/sqllib/bin/db2 deactivate database MYUDB
user udbinst using password
exit 0
```

**Sample UDB
post-replication
scripts (offline)**

The script is a basic post-replication script to restart the database. Add any additional actions desired for your site.

```
#!/bin/sh
id
DBINSTANCE=udbinst
export DBINSTANCE
/udb/udbinst/sqllib/bin/db2 connect to MYUDB user udbinst
USING password
/udb/udbinst/sqllib/bin/db2 force applications all
```

```
/udb/udbinst/sql/lib/bin/db2 restart database MYUDB user  
udbinst USING password  
exit 0
```

General guidelines for scripts

General guidelines for scripts are as follows:

- ◆ Each script should be an executable file of some kind (for example, .sh).
- ◆ Shell scripts, whether Bourne shell (sh), C shell (csh), or Korn shell (ksh), require their respective headers. For example, for sh, the script needs the header #!/bin/sh to work correctly.
- ◆ Do not send more than 255 characters of output to either stdout or stderr. If you do, the output will be truncated. If more output or error data is required, set up and use your own log file.
- ◆ Return a 0 (zero) status to continue with the replication; return a nonzero status to instruct Replication Manager to fail the job.

UDB troubleshooting

Certain situations can cause issues that can be remedied by making changes to the configuration. The following sections describe troubleshooting tips. Consult these sections before calling EMC Technical Support.

UDB agent unable to locate the UDB package

The Replication Manager UDB agent uses certain commands that are part of the UDB installation. To access those commands, the UDB agent must identify the location where the UDB package is installed. These are the steps Replication Manager uses to identify where the UDB package resides:

- ◆ UDB agent searches the operating system packages to find out where the UDB product is installed on the system, for example UDB may be located at:

```
/opt/IBM/db2/V8.1
```

Solaris: UDB agent uses pkginfo to locate UDB.

AIX: UDB agent uses lspp to locate UDB.

- ◆ The bin and adm directories under the UDB installation are identified using the package found above. For example:

```
/opt/IBM/db2/V8.1/bin and /opt/IBM/db2/V8.1/adm
```

If the package installation is incomplete the UDB agent may not find the path in which the UDB product and its corresponding libraries are installed. In this case the environment variables specific to UDB instance can be set in the Replication Manager Console under Host properties, Advanced tab so the UDB binaries and libraries can be found.

The environment variables are:

```
ERM_UDB_INSTALL_PATH
```

If the UDB instance is db2inst1 and the UNIX home directory for the instance is /udbhome/db2inst1, set this environment variable to /udbhome/db2inst1.

```
ERM_UDB_VERSION
```

If the UDB installation path could not be found at the first location, it means Replication Manager could also not identify what the UDB

version was. This variable will set the UDB version for Replication Manager. For example, the value for UDB 9.0 fixpack 5 is 9050000.

Additionally, the library path in this case would have to be updated to point to the UDB libraries for this version. For example, on Solaris, the LD_LIBRARY_PATH should include /opt/IBM/db2/V9.5/lib32.

Workaround for UDB 8 restore failure

When you perform a restore of UDB 8 replica, existing database connections may prevent file systems from being unmounted cleanly, resulting in a restore failure. If this occurs, the DBA can manually force closure of all concurrent connections to the database to put database into a state to allow restore to succeed. The steps are:

```
force application all
db2stop
db2start
```

Workaround for UDB 8/UDB 9 coexistence

To replicate a UDB 8 database on a system where UDB 9 also is installed, use the following workaround:

1. Run **rc.irclient stop** to stop the **irccd** daemon.
2. Unlink /usr/local/bin/db2ls from its original location.
3. Run **rc.irclient start** to restart the **irccd** daemon.

UDB database is not accessible on the mount host

If a UDB database is not accessible on the mount host after a UDB SNAPSHOT mount, try this workaround:

1. Check the /tmp/relocatedb.cfg_*DBname* file.
2. Run the following command:

```
db2inidb DB_name as SNAPSHOT relocate using  
tmp/relocatedb.cfg_DB name
```

3. If the command gives you a syntax error, try to remove the CONT_PATH arguments from the file and run the command again.

Replication Manager can protect SQL Server databases by creating and managing application sets that contain either entire SQL Server databases or partial databases at the file group level. This appendix covers the specifics of the SQL Server support, including sections on the following:

- ◆ Configuring the SQL Server environment 432
- ◆ Understanding SQL Server application sets and jobs 439
- ◆ Dynamic Discovery of SQL Server databases..... 449
- ◆ Mounting and restoring SQL Server replicas 464
- ◆ Using SQL Server snapshot functionality 481
- ◆ Using pre- and post-replication SQL Server scripts..... 483
- ◆ Considerations for working with SQL Server in a cluster 485

Configuring the SQL Server environment

Replication Manager supports SQL Server 2000, SQL Server 2005, SQL Server 2008, and SQL Server 2012. Before using Replication Manager with SQL Server, ensure that the environment is configured properly. The following checklist can help you configure the environment to work with Replication Manager:

- ❑ For SQL Server 2005 and higher, the SQL Server Browser Service needs to be running when you:
 - ❑ Create the first application set for a SQL Server instance
 - ❑ Use dynamic ports
- ❑ Alternate mount hosts must have an installed SQL Server software instance unless you are mounting the replica without recovering the database (leaving the **Recover All Databases** checkbox empty). EMC recommends that you use the same version of SQL Server on both the production and alternate mount hosts.

Note: You cannot recover a SQL Server database to an older version of SQL Server. For example, you cannot recover a SQL Server 2008 database to a SQL Server 2005 or 2000 instance.

- ❑ Replication Manager does not truncate transaction logs so you should create a Database Maintenance Plan to coexist with your scheduled replications. The plan should back up and truncate the transaction logs unless your databases are using the Simple Recovery Model. The plan should span the intervals between replications to protect changes that occur between those replications. For more information about creating a Database Maintenance Plan, refer to your SQL Server documentation.
- ❑ Specific SQL Server Service Packs are required. Refer to the *EMC Replication Manager Support Matrix* for the latest support information. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.
- ❑ In SQL Server 2012, the default virtual account used in the service startup account of the database engine does not have the requisite file system permissions for accessing the mounted or restored database files. Therefore, recovery of SQL databases may fail. To

overcome this, you must change the service startup account for the SQL Server database engine to use a domain user account with appropriate privileges and permissions.

SQL Server prerequisites

Verify that your SQL Server configuration meets the prerequisites listed below. Refer to the *EMC Replication Manager Support Matrix* for updated information on required service packs and supported operating systems. Specific SQL Server Service Packs are required. Refer to the *EMC Replication Manager Support Matrix* for the latest support information. To access the *Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

- ◆ Replication Manager supports SQL Server 2000, SQL Server 2005, and SQL Server 2008 on various Windows operating systems. Refer to the *EMC Replication Manager Support Matrix* for specific SQL Server operating system requirements.
- ◆ When using CLARiiON or VNX Consistent Split, a SQL Server database and its transaction logs must be located on disks in the same storage array, since CLARiiON or VNX Consistent Split does not support multiple array datasets.
- ◆ Multiple SQL Server databases can exist on the same volume, or across multiple volumes.
- ◆ Replication Manager can support SQL Server 2000, SQL Server 2005, and SQL Server 2008 instances that coexist on the same server or cluster.
- ◆ The SQL Server database must be online during replication.
- ◆ The system databases (master, msdb, model, etc.) should not be located on the same volume as user databases. Microsoft SQL Server does not currently support using Virtual Device Interface (VDI) and snapshot technology to restore system databases.
- ◆ (SQL Server 2005 and SQL Server 2008 only) Any full-text catalogs associated with a filegroup will be automatically included as part of a replica of that filegroup. If the full-text catalogs are not located on supported storage, the job will fail.

Note: When using full-text catalogs, you should make sure that the storage device where the catalog is located does not include data that is not related to the database.

- ◆ Replication Manager will not discover SQL Server 2005 or SQL Server 2008 Database Snapshots since these snapshots cannot be replicated.

Note: In this context, the term "Database Snapshots" refers to the feature in SQL Server 2005 and SQL Server 2008 called database snapshot technology, not Replication Manager snapshots.

- ◆ (SQL Server 2008 only) Replication Manager automatically discovers filestream filegroups associated with the database, if any.

Required permissions and rights

Users require certain permissions and rights to configure application sets and run jobs in a SQL Server environment. The user account must be configured to use either SQL Server authentication OR Windows authentication.

When using Celerra storage in a Windows environment, the user must be a member of the local Administrators group.

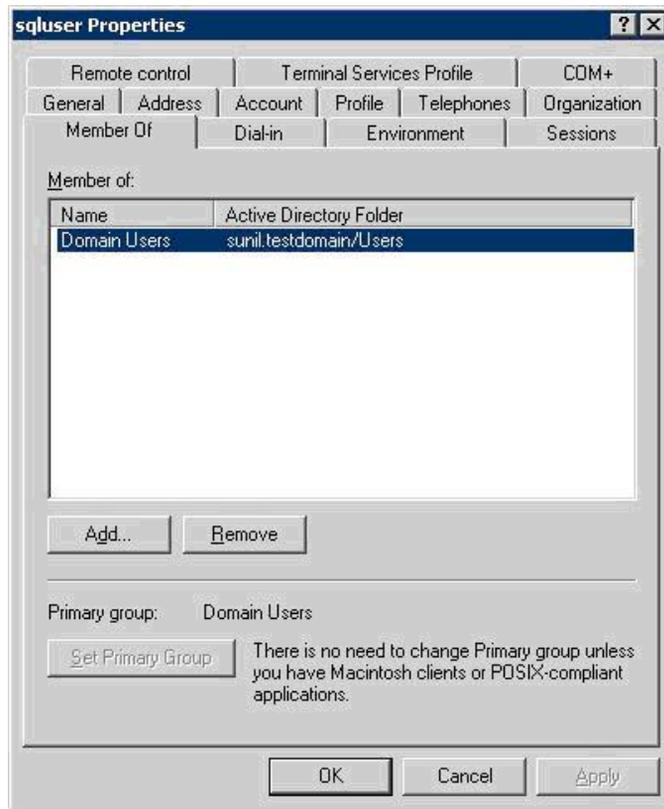
When using other types of storage, the Windows user account can either be a member of the local Administrators group or a non-Administrator account with the restrictions outlined next.

The user account running Replication Manager must have read and write permissions on the client\bin directory where Replication Manager is installed.

For non-administrators

If you use a domain user account that does not have local administrator privileges, the following pre-configuration steps are required:

1. Ensure that SQL Server's VDI API is able to support impersonation. The Microsoft hot fix QFE-934396 is required for this feature to work correctly on SQL Server 2000 and 2005.
2. Create a Windows domain user account (for example, sqluser) and make it part of the Domain Users group.



- To use callout scripts, use **regedt32** to grant Full Control Permissions to the Replication Manager registry key. On Windows 2012 Server, Windows 2008 Server, and Windows 2003 32-bit systems, the registry key is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\EMC ControlCenter
Replication Manager
```

On Windows 2003 64-bit servers, the registry key is:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\EMC\EMC
ControlCenter Replication Manager
```

- Create a new SQL Server login using the domain account and Windows authentication.
- Select **master** as the default database on the **General** tab.

6. In the **Server Roles** page of the **Login Properties** dialog box, assign System Administrator (sysadmin) and public server roles.
7. In the **User Mapping** page of the **Login Properties** dialog box, set the database role membership to public.
8. Add this user to all SQL instances to which this user would need access:
 - a. Log in to the domain controller machine and each host added to that domain that uses Replication Manager and set the Security policy:
 - On the domain controller:
Start > Programs > Administrative Tools > Domain Controller Security Policy
 - On the hosts added to the domain:
Start > Programs > Administrative Tools > Local Security Policy
 - b. Access security settings and allow login locally. **Security Settings > Local Policies > User Rights Assignment > Allow log on locally**. Add the user you created earlier.
9. Grant this user read and write permissions on the client directory where Replication Manager is installed.
10. Use this user from the Replication Manager Console to connect to SQL Server.
11. At the time of restore, if you select to back up the transaction logs to a file, the user must have rights to the target directory.

For non-administrator local user accounts

If you use a local user account that does not have local administrator privileges, the following pre-configuration steps are required:

1. Ensure that SQL Server's VDI API is able to support impersonation. The Microsoft hot fix QFE-934396 is required for this feature to work correctly on SQL Server 2000 and 2005.
2. Create a Windows user account (for example, sqluser) and make it part of the Users group.
3. Create a new SQL Server login for the user with Windows authentication.
4. In the Server Roles page of the Login Properties dialog box, assign System Administrator (sysadmin) and public server roles.

5. In User Mapping page of the Login Properties dialog box, set the database role membership to public.
6. Assign the user rights to log in to the machine:
 - a. Log in to the host that uses Replication Manager and set the Security policy.
On the domain controller:
Start > Programs > Administrative Tools > Local Security Policy
 - b. Access security settings and allow login locally **Security Settings > Local Policies > User Rights Assignment > Allow Logon locally**. Add the user you created earlier.
7. Grant this user read and write permissions on the client directory where Replication Manager is installed.
8. Use this user from the Replication Manager Console to connect to SQL Server.
9. At the time of restore, if you select to back up the transaction logs to a file, the user should have rights to that directory.

Support for upgrades to SQL Server 2005 and 2008

Replicas of a SQL Server 2000 database instance that were created using Replication Manager can be mounted and restored after an upgrade to SQL Server 2005 or SQL Server 2008.

Replicas of a SQL Server 2005 database instance that were created using Replication Manager can be mounted and restored after an upgrade to SQL Server 2008.

Understanding SQL Server application sets and jobs

Replication Manager uses application sets and jobs to define what data to add to a replica and the steps to take when creating the replica. The following sections describe topics you should consider when creating application sets and jobs for SQL Server replicas.

Microsoft SQL Server instances that reside on Symmetrix systems can be included as part of federated application sets. Some restrictions apply and special configuration steps may be required. [Chapter 7, "Configuring Federated Data,"](#) provides full details on how to create Microsoft SQL Server federated application sets.

SQL Server application sets

Users with the appropriate permissions can configure an application set to replicate a full SQL Server database or one or more file groups. The Application Set Wizard allows users to specify whether to replicate the entire database or only certain filegroups:

Database replication — Copies all the data and related transaction logs (for example, .ldf), which must also be located on disks in the same storage array as the data.

Filegroup replication — Copies one or more filegroups. Related active transaction log(s) are not replicated as part of a filegroup replication. This is configured by choosing a subset of filegroups instead of choosing the entire contents of the database. [Figure 160 on page 440](#) illustrates this.

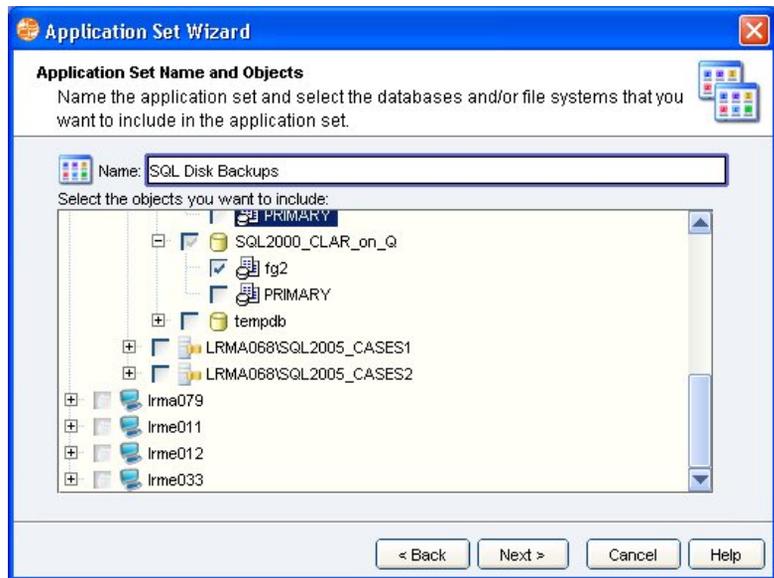


Figure 160 Configuring a filegroup replication

Note: Filegroup replications are not recommended unless the database is large and the storage resources are limited. Most replicas should use the database replication method.

If replicating at the filegroup level, make sure no other datafiles are on the volume with the selected filegroup's datafiles.

Transaction log backups are required to restore a filegroup replication.

You cannot use the **Recover a database** option when mounting a filegroup replication.

If you want to perform a filegroup restore sometime in the future, EMC recommends that you configure your application set to create a replica of the entire database and restore only the part of the database replica that you need. In that case, if something goes wrong, you always have the option to mount the complete replica to an alternate host and manipulate the data further.

SQL Server 2005/2008 consistency methods

When you are creating a job to take replicas of a SQL Server 2005 or SQL Server 2008 database, you must choose one of the following three consistency methods:

- ◆ **Online with limited recovery** — Replication Manager creates a replica and quiesces it with consistent split alone; Microsoft SQL Server VDI mode is not used. (available if consistent split was used.)

Note: If you need to restore a replica created using online with limited recovery mode, you can only restore to the point in time that the replica was created. In addition, replicas created with the **Online with limited recovery** option do not allow you to back up the transaction logs before restoring the database. This is because the transaction logs cannot be applied to this type of replica.

- ◆ **Full, Online with advanced recovery (using VDI)** — Replicates the database, and the active part of the transaction log. This replica type is typically used when the replica will be considered a backup of the database or when the replica will be mounted in order to use a third-party product to create a backup of the database. This type of replica allows you to restore transaction logs to bring the database forward to a point in time that is newer than the replica, assuming you have backed up those transaction logs. Replication Manager uses Microsoft SQL Server's VDI snapshot feature to create this type of replica.
- ◆ **Copy, Online with advanced recovery (using VDI)** — Replicates the database and the active part of the transaction log without affecting the sequence of backups. This provides DBAs with a way to create a replica without interfering with third-party backup applications that may be creating full and/or differential backups of the SQL Server databases. Replication Manager uses Microsoft SQL Server's VDI snapshot feature to create this type of replica.

Figure 161 on page 442 illustrates the consistency method portion of the tree where you can choose one of the options listed above.

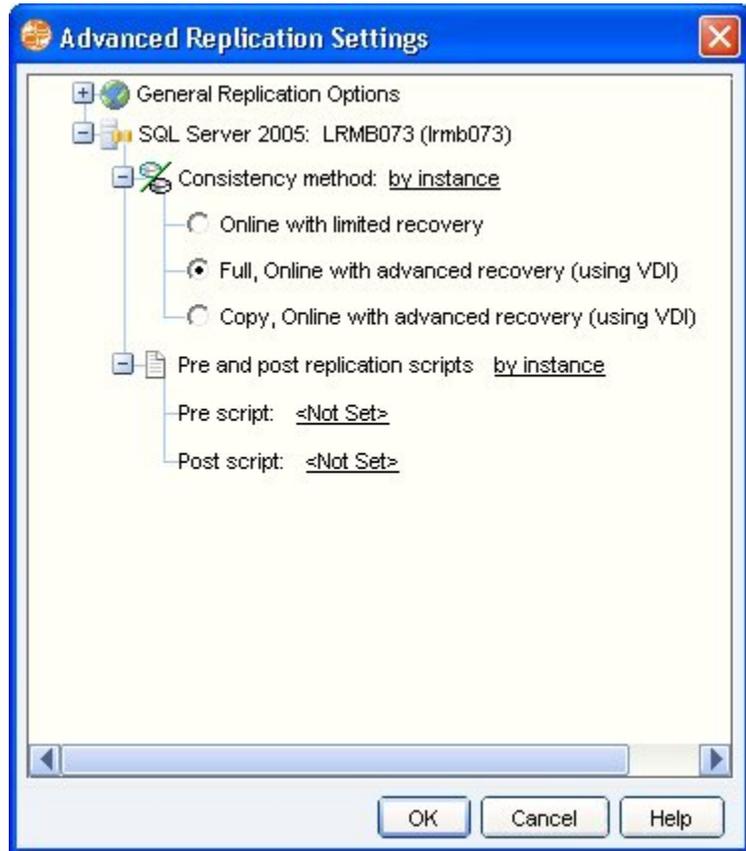


Figure 161 Advanced Replication Settings for SQL Server 2005/2008

SQL Server 2000 replication types

When you are creating a job to take replicas of a SQL Server 2000 database, you must choose between two replication types as follows:

- ◆ **Online With Limited Recovery** — Replication Manager creates a replica and quiesces it with consistent split alone; Microsoft SQL Server VDI mode is not used.

Note: If you need to restore a replica created using online with limited recovery mode, you can only restore to the point in time that the replica was created. In addition, replicas created with the **Online with limited recovery** option do not allow you to back up the transaction logs before restoring the database. This is because the transaction logs cannot be applied to this type of replica.

- ◆ **Online With Advanced Recovery** — See [“SQL Server online replication \(with or without VDI\)”](#) on page 443 for more information on advanced recovery.

SQL Server 2000 does not offer the Copy option.

SQL Server online replication (with or without VDI)

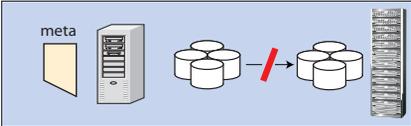
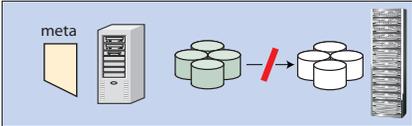
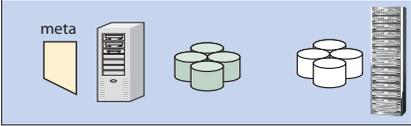
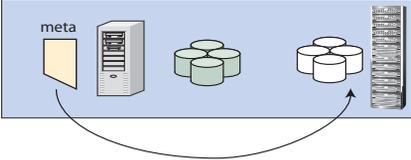
The steps in [Table 30 on page 444](#) illustrate SQL Server replication processing in Online mode. The left side describes **Online with Advanced Recovery (Using VDI)** while the right describes **Online with Limited Recovery (no VDI)**.

Replication Manager performs the steps outlined in [Table 30](#) on [page 444](#).

Table 30 Online replication steps (page 1 of 2)

Online with advanced recovery (using VDI)	Online with limited recovery (no VDI)
<p>1. Create metadata file.</p>	<p>Note: The metadata file is created only when using VDI. If you are using consistent split and have selected Online with Limited Recovery, then no metadata file is created.</p>
<p>2. Establish mirrors (if necessary), then Replication Manager SQL Server Agent creates a VDI snapshot.</p>	<p>Establishes mirrors (if necessary).</p>
<p>3. Use VDI to deactivate (freeze) input/output to the database.</p>	

Table 30 Online replication steps (page 2 of 2)

Online with advanced recovery (using VDI)	Online with limited recovery (no VDI)
<p data-bbox="325 302 515 326">4. Creates the replica.</p> 	<p data-bbox="753 302 1129 354">Creates the replica using consistent split to create a crash-consistent copy of the data.</p>  <p data-bbox="753 522 1158 631">Note: This means you can restore to the point-in-time when the replica was created, but you cannot apply transaction logs beyond that point-in-time.</p>
<p data-bbox="325 680 722 784">5. Uses VDI to reactivate (thaw) input/output to the database so the SQL Server database can resume normal operations, including writes to the database.</p> 	
<p data-bbox="325 982 686 1006">6. Copies the metadata file to the server.</p>  <p data-bbox="354 1225 729 1486">(Replication Manager SQL Server Agent catalogs the replica by copying all relevant information to its internal database; it catalogs the locations of the datafiles, transaction log files, and the metadata file required for a restore or mount. If you mount the replica to backup the data to long-term storage (such as tape media), remember to backup all of the metadata files generated for the replica, including this one.)</p>	

Note: SQL Server does not currently support using VDI and snapshot technology to restore system databases. The master, msdb, and model system databases should not be located on the same device or volume group as user databases.

The SQL Server database must be online during replication.

Use consistent split when you need the same point-in-time copy of data from multiple applications (for example, if you need a replica that includes a SQL database, an Oracle database, and a file system).

Backing up SQL Server including VDI metadata files

The SQL Server VDI metadata files are required for a complete backup of the SQL Server database. When you mount a replica, these files are transferred to the mount host so that you can create a backup to tape or other media.

Replication Manager can integrate with third-party backup software to create tape backups of SQL Server replicas. The following procedure assumes that you have already successfully created and run a SQL Server replication job:

1. In the content panel, right-click the replica that you want to back up to tape and select **Mount**.
2. In the **Replica Mount** screen in [Figure 162 on page 447](#), expand the **General Mount Options** and look for the field entitled **Copy Metadata Files to**. Select the checkbox next to that field and specify the directory path on the mount host to which you want to copy SQL Server metadata files.

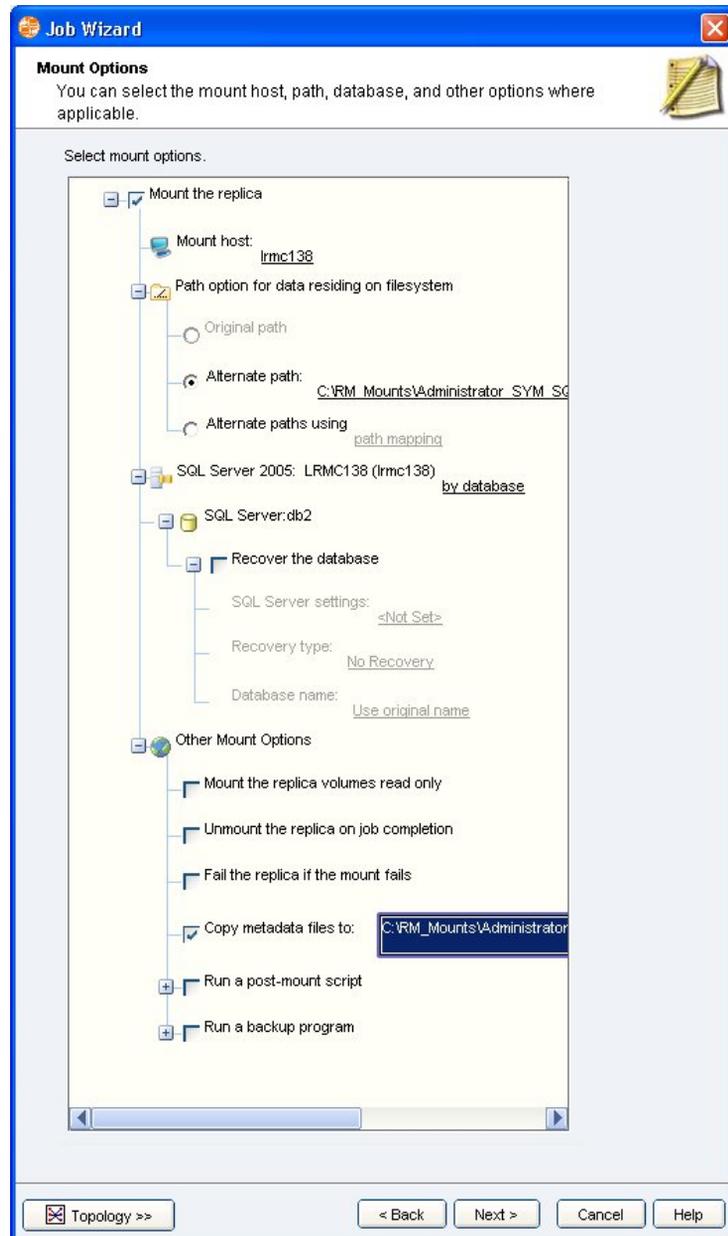


Figure 162 Mount and copy metadata files

3. Mount the SQL Server replica.

4. Using your third-party backup software, back up all volumes of the replica that reside on the mount host, along with all metadata files located in the directory path you specified in the **Copy Metadata Files to** field.

Replicating SQL Server 2005/2008 with mirror sessions

SQL Server 2005 and SQL Server 2008 can create application-based mirrors of databases independent of Replication Manager or other EMC technologies. This section describes how Replication Manager can interact with these existing mirror sessions.

When a database has an active mirror session each session would reside on a different server, to create a replica of these mirror sessions requires an application set for each server, one for principal and one for the mirror. The job can only run on the principal database.

Note: You cannot run a job to create a replica of the SQL Server mirror unless you first failover to the mirror, making it the principal. If you attempt to create a replica of a SQL Server mirror, the job fails with an error indicating that the database is currently acting as a mirror.

SQL Server 2008 filestream datatype

SQL Server 2008 includes a mechanism that makes it easier to store unstructured data by using a data type called filestream. The filestream data type offers rich streaming APIs that parallel the performance of direct file system writes, while at the same time maintaining a transactional consistency between this unstructured data and structured data. Additionally, the filestream datatype offers better security than other methods used to store unstructured data.

Replication Manager can create replicas of SQL Server 2008 databases that employ the filestream data type. It is possible to mount filestream replicas to the production host or to an alternate host.

Dynamic Discovery of SQL Server databases

You now have the option of replicating all user databases when creating a SQL Server application set. In the Application Set Wizard, there is now an option, **SQL Server Dynamic Discovery**. If you select this option, Replication Manager will dynamically discover and replicate all user databases of an instance at replication time, eliminating the need for modifying the application set and jobs.

Note: Prior to Replication Manager 5.3.2, the support for SQL Server was static, that is, Replication Manager would only replicate the databases included at the time the application set was created. If you added or deleted a database, Replication Manager did not detect these modifications automatically, and you were required to manually modify the application set and jobs accordingly or the job would fail.

Overview of dynamic discovery support

Note the following with regard to the dynamic discovery option:

- ◆ Dynamic discovery requires all Replication Manager components (the server, UI, and agent) to be at Replication Manager 5.3.2 or above.
- ◆ Dynamic discovery supports all versions of SQL Server that Replication Manager supports.
- ◆ Dynamic discovery is set at the time an application set is created and takes effect when a job is run.
- ◆ Dynamic discovery supports both federated and non-federated SQL Server application sets.
- ◆ Dynamic discovery is supported for user databases only. Refer to the following section, [“System and User databases” on page 450](#) for a discussion on system and user databases and how they differ.
- ◆ If dynamic discovery is selected for an instance, there is no “by database” support for SQL Server replication options, mount options, and restore options.
- ◆ If dynamic discovery is selected for an instance, all user databases must be in good state for replication, otherwise the replication fails.

- ◆ If dynamic discovery is selected for an instance, the instance must have at least one user database, otherwise replication fails with an error, “no databases discovered.”
- ◆ If dynamic discovery is selected for an instance, all user databases must be on the same type of storage array, for example, all on Symmetrix, or all on CLARiiON or VNX. Otherwise, replication fails as Replication Manager does not support mixed arrays in the same job.

System and User databases

The Replication Manager user interface now distinguishes system databases from user databases in the SQL server tree. [Figure 163 on page 450](#) shows the two nodes for user databases and system databases.

System databases are created by SQL Sever and include master, model, msdb, tempdb, resource, distribution, ReportServer and ReportServerTempDB. User databases include only the databases that you created. The dynamic discovery feature is only available for user databases.

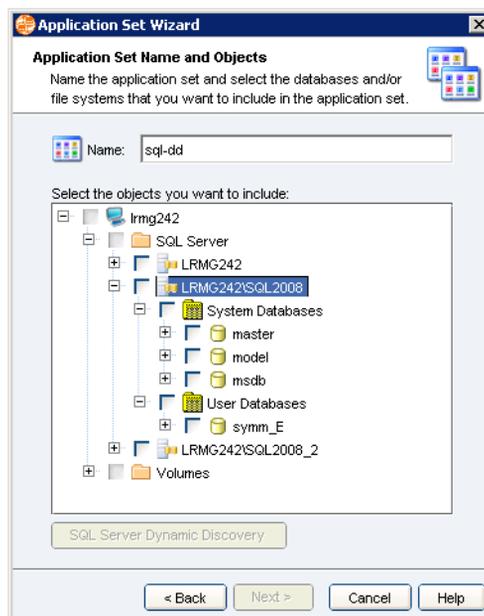


Figure 163 System and user databases

Creating application sets that use dynamic discovery

Dynamic discovery is an attribute that is set at the time an application set is created and takes effect when the job is run.



IMPORTANT

You can not modify the selection of dynamic discovery on the instances included in the application set. For example, if dynamic discovery is set on for an instance when the application set is created, the application set can not be modified afterwards to switch off dynamic discovery. Similarly, if an application set is created with dynamic discovery turned off, you can not turn it on by modifying the application set at a later date.

However, if you needed to change this you can remove the instances included in the application set and add them back with or without dynamic discovery as desired. Refer to [“Modifying application sets that use dynamic discovery” on page 461](#).

To use dynamic discovery:

1. Select all user databases for an instance, [Figure 164 on page 452](#).

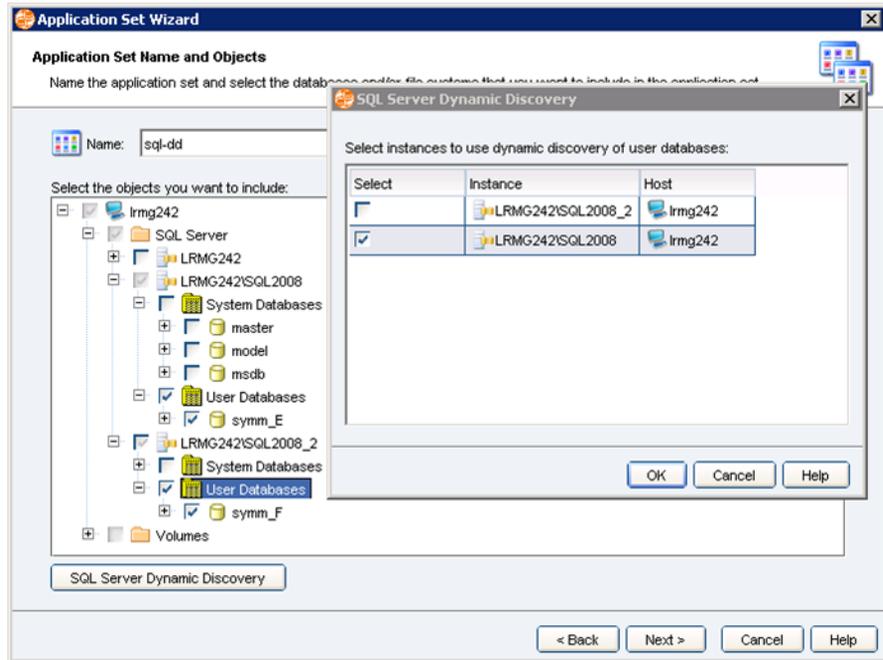


Figure 164 selecting all user databases.

Note: Initially, the SQL Server Dynamic Discovery button is disabled. It is enabled only when no system databases are selected and all user databases are selected for at least one instance.

2. Select **SQL Server Dynamic Discovery**.

The SQL Server Dynamic Discovery dialog displays all instances in which you can select dynamic discovery of user databases.

- In the SQL Server Dynamic Discovery dialog, select the instance that you wish Replication Manager to dynamically discover SQL Server user databases on as shown in [Figure 163 on page 450](#).

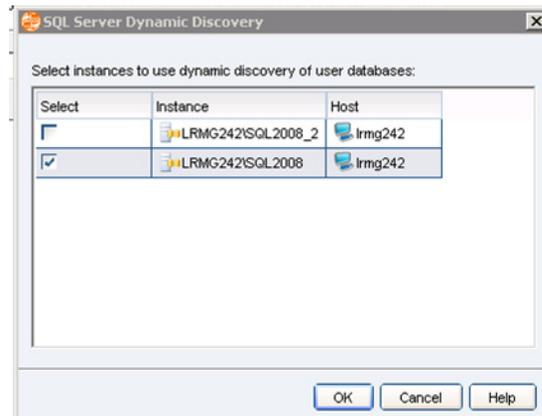


Figure 165 Selecting the instance for dynamic discovery.

Note: If dynamic discovery is selected, the application set summary and properties do not list the individual databases. Refer to the following section, [“Viewing application set properties” on page 453](#)

- Click **OK**.

The Application Set Wizard resumes for you to finish creation of the application set.

Viewing application set properties

You can check which SQL Server instance has dynamic discovery selected by viewing the properties of the application set.

To view the properties of the application set:

- Right-click the Application Set in the tree view and select **Properties**.
- In the Application Set Properties dialog, select the **Objects** tab.

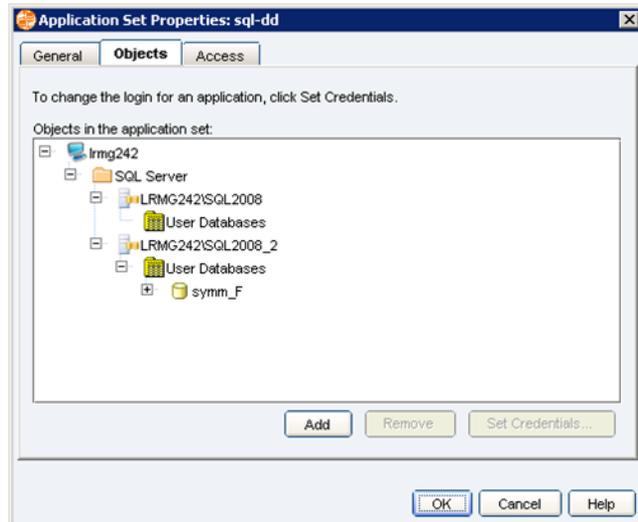


Figure 166 Application set properties

Refer to the two nodes in [Figure 166 on page 454](#):

- ◆ For SQL Server instances with dynamic discovery enabled, (LRMG242\SQL2008 in the figure above) only the User Databases node appears in the Application Set Properties display. You can not expand the node to see individual databases within it.
- ◆ For SQL Server instances without dynamic discovery enabled, (LRMG242\SQL2008_2 in the figure above) individual databases selected in the application set (in this case, symm_F) are listed under the User Databases node.

Creating jobs that use dynamic discovery

When you are creating jobs that use application sets that have dynamic discovery enabled, there are implications for the available replication options in the Advanced Replication Settings dialog and also for the mount options dialog.

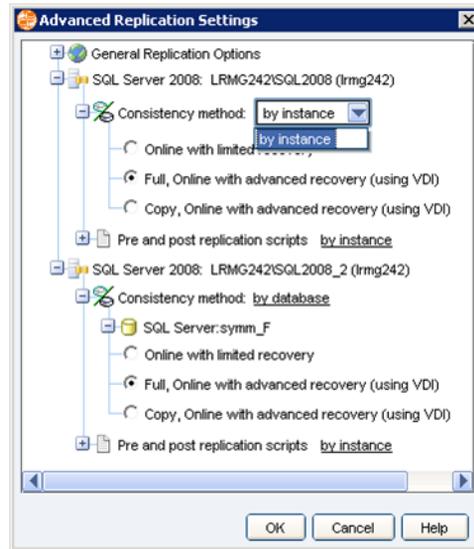


Figure 167 Replication options

Replication options

When you are creating jobs that use application sets that have dynamic discovery enabled:

- ◆ For SQL Server instances with dynamic discovery enabled, (LRMG242\SQL2008 in the figure above) only the “by instance” option is available in the drop-down menus for consistency method and scripts.
- ◆ For SQL Server instances without dynamic discovery enabled, (LRMG242\SQL2008_2 in the figure above) both the “by instance” and “by database” options are available in the drop-down menus for consistency method and scripts.

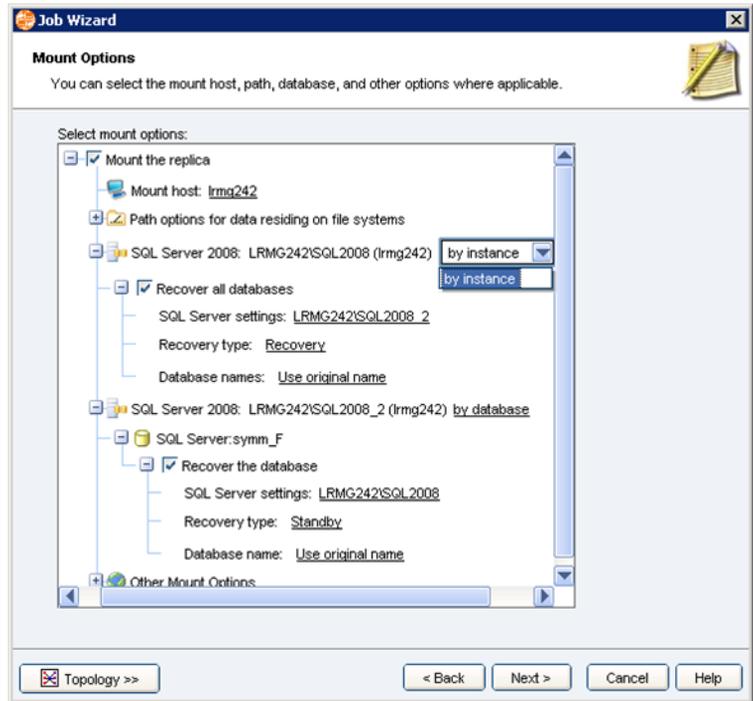


Figure 168 Mount options

Mount options

When you are creating jobs that use application sets that have dynamic discovery enabled:

- ◆ For SQL Server instances with dynamic discovery enabled, (LRMG242\SQL2008 in the figure above) only the “by instance” option is available in the drop-down menus in the Mount Options screen of the Job Wizard.
- ◆ For SQL Server instances without dynamic discovery enabled, (LRMG242\SQL2008_2 in the figure above) both the “by instance” and “by database” options are available in the drop-down menus in the Mount Options screen of the Job Wizard.

Note: These two points also apply when you do on-demand mounts, also.

Running jobs that use dynamic discovery

When you are running jobs that use dynamic discovery, Replication Manager dynamically discovers and replicates all user databases of a SQL Server instance if dynamic discovery is selected for that instance.

An example

The following example illustrates how Replication Manager discovers new databases. The figures depict two different replicas of the same job with different databases.

Figure 169 on page 457 is the History Log for the first run of a job. It shows which databases are being backed up (The red circle highlights the SQL backup commands completed for given databases.)

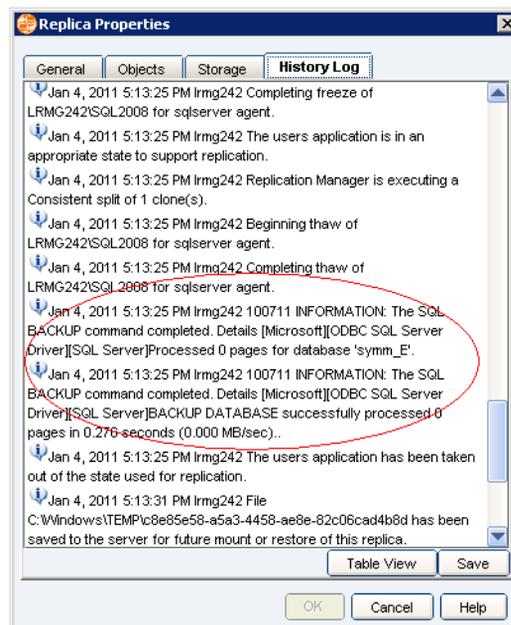


Figure 169 First run of a job — Replica properties, History Log tab

Figure 170 on page 458 is the Objects tab of the Replica Properties and shows what databases are included in the replica. In this example, Replication Manager discovered one database, symm_E.

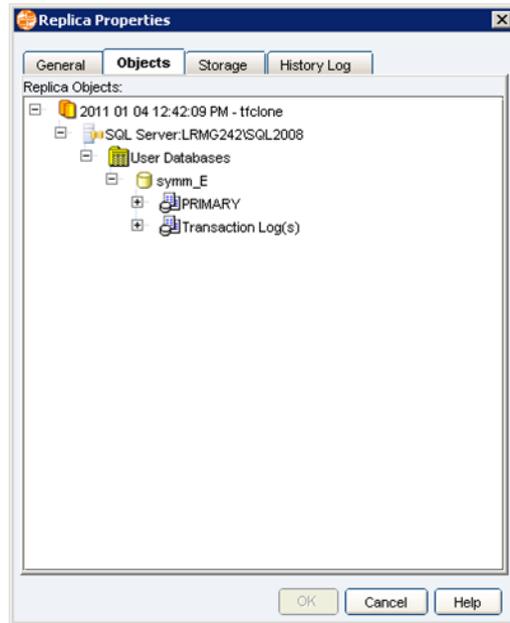


Figure 170 First run of a job — Replica properties, Objects tab

Figure 171 on page 459 is the History Log for the second run of a job. It shows which databases are being backed up (The red circle highlights the SQL backup commands completed for given databases.)

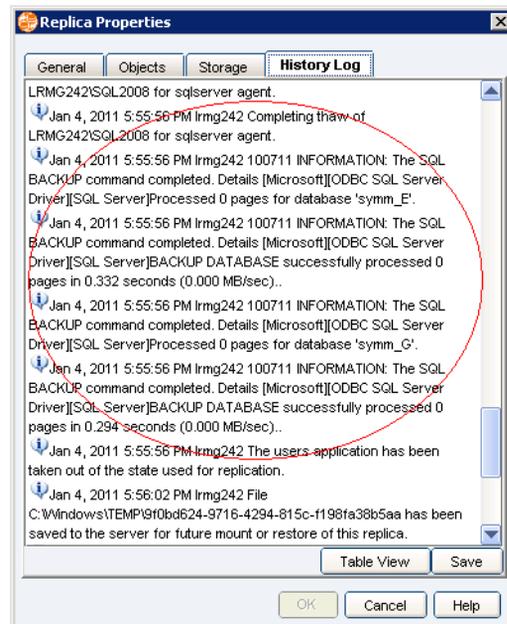


Figure 171 Second run of a job — Replica properties, History Log tab

Figure 172 on page 460 is the Objects tab of the Replica Properties and shows what databases are included in the replica. In this example, Replication Manager discovered two databases, `symm_E` and `symm_G`.

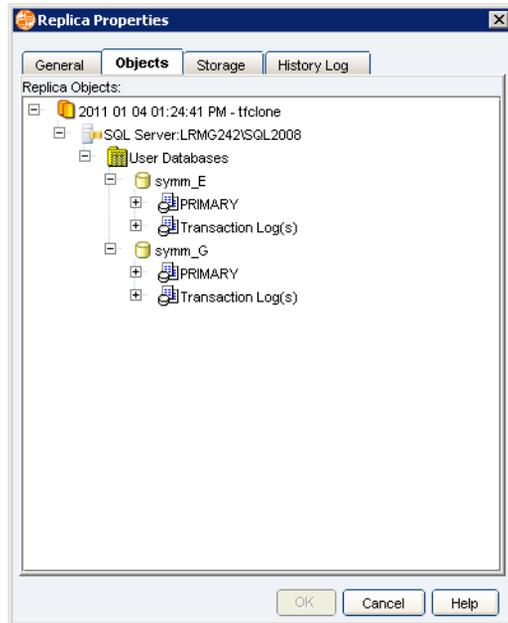


Figure 172 Second run of a job — Replica properties, Objects tab

Restoring replicas that use dynamic discovery

When you are restoring replicas:

- ◆ For SQL Server instances with dynamic discovery enabled, (LRMG242\SQL2008 in the figure below) only the “by instance” option is available in the Restore Options screen of the Restore Wizard.
- ◆ For SQL Server instances without dynamic discovery enabled, (LRMG242\SQL2008_2 in the figure below) both the “by instance” and “by database” options are available in the Restore Options screen of the Restore Wizard.

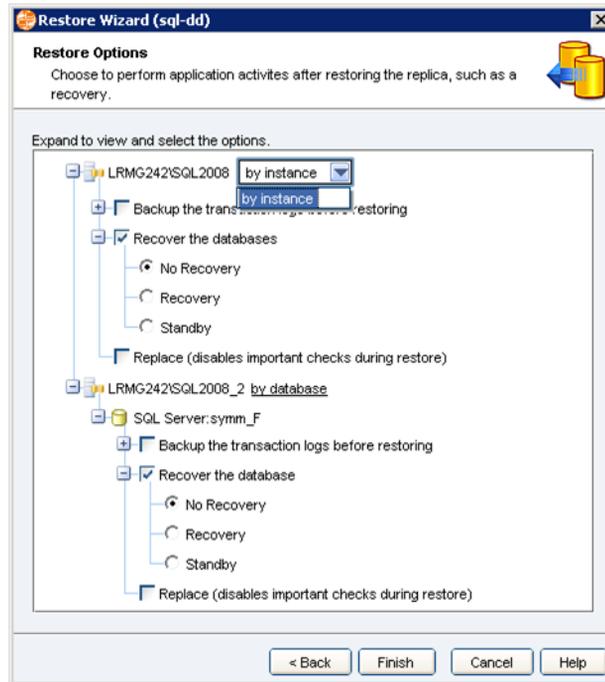


Figure 173 Restore replica

Modifying application sets that use dynamic discovery

When you modify application sets that use dynamic discovery:

- ◆ You can not change the selection of dynamic discovery for instances already included in the application set:
 - The instances with dynamic discovery enabled are grayed out (and therefore, unable to be modified as such).
 - The instances without dynamic discovery enabled are active and can add new databases, however, these instances will not be listed for dynamic discovery selection.
- ◆ You can add new instances with or without dynamic discovery selected.
- ◆ You can remove the included instances and add them back with or without dynamic discovery selected.

Note: If an instance was included in the application set without dynamic discovery and you removed the instance and added it back with dynamic discovery selected, you will not be able to select the replication and mount options by database in the jobs.

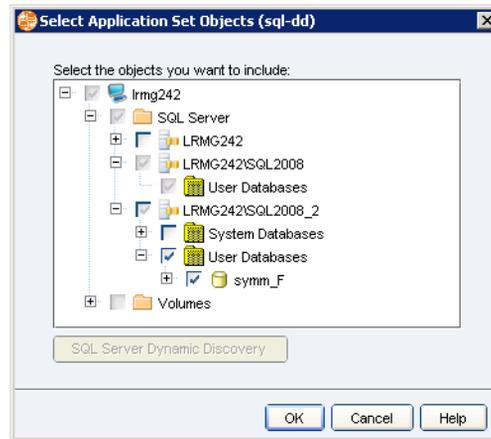


Figure 174 Modify application set

Upgrading implications for dynamic discovery

When you upgrade Replication Manager, take note of the following in regards to dynamic discovery feature:

- ◆ You must upgrade all Replication Manager components (the server, UI, and agent) to be at Replication Manager 5.3.2 or above for dynamic discovery support.
- ◆ If you upgrade the Replication Manager server, UI, and agent to Replication Manager 5.3.2 or above:
 - Application sets, jobs, and replicas created before the upgrade continue to behave as before the upgrade.
 - The new SQL Server tree view (with user database and system database nodes) is displayed when modifying the application set created before upgrade.
 - The SQL Server instances included in the application set created before the upgrade can not be modified to select dynamic discovery. New SQL Server instances can be added with or without dynamic discovery.

- The dynamic discovery option will not be available until Replication Manager detects the agent is at version 5.3.2 or above.
- ◆ If you upgrade the Replication Manager server and UI to Replication Manager 5.3.2 or above, but the Replication Manager agent is not upgraded:
 - Application sets, jobs, and replicas created before the upgrade continue to behave as before the upgrade.
 - The new SQL Server tree view (with user database and system database nodes) is displayed when modifying the application set created before upgrade.
 - Dynamic discovery remains disabled due to retention of the older agent.

Mounting and restoring SQL Server replicas

The Replication Manager SQL Server Agent can:

- ◆ Mount a SQL Server replica to an alternate host.
- ◆ Mount a SQL Server replica to the production host.
- ◆ Restore a replica over the original database.

Note: During a mount to an alternate host, EMC recommends that your mount host meet certain configuration guidelines. These guidelines are described in [“Mounting SQL Server replicas” on page 464](#).

EMC recommends that you choose to mount all SQL Server replicas as part of the process of creating them.

Mounting SQL Server replicas

The Replication Manager SQL Server Agent can mount a replica to an alternate host, as long as that host has:

- ◆ SQL Server and the agent software installed

Note: If you do not intend to recover the database, SQL Server software is not required.

- ◆ Identical version of Replication Manager Agent software on the production and mount hosts
- ◆ Registered with the Replication Manager Server
- ◆ Identical operating system¹, volume manager, file system, HBA drivers, and application versions as the production host

Replication Manager can perform mounts:

- ◆ On an alternate mount host to the same location as the production host.

1. However, for Windows 2003 and Windows 2008, Replication Manager supports differences in operating systems between the mount host and production host. Replicas created on Windows 2003 can be mounted to another Windows 2003 server, even if the platform is different. The same applies to Windows 2008. For example, a replica created on Windows 2003 x86 can be mounted to a Windows 2003 x64 server and vice-versa; a replica created on Windows 2003 IA64 can be mounted to a Windows 2003 x64 server and vice-versa; a replica created on Windows 2003 IA64 can be mounted to a Windows 2003 x86 server and vice-versa. The same examples apply for Windows 2008. Windows 2012 is x64 only.

- ◆ On an alternate mount host to a new location (determined by adding an alternate mount path to the pathname).
- ◆ On an alternate mount host to a new location (determined by path mapping). To the original production host in a new location (determined by adding an alternate pathname).
- ◆ To the original production host in a new location (determined by path mapping).
- ◆ Using an alternate SQL Server instance (selected from a list) to prevent instance collisions when mounting to the same mount host.
- ◆ Renaming the database to prevent the same database name from running on the same host.

Note: Specific information about how alternate paths and path mapping work can be found at [“Mounting using alternate path” on page 174](#) or [“Mounting using path mapping” on page 177](#).



CAUTION

After a database is mounted to an alternate location on either the production or the mount host using its original name, it cannot be remounted to the same SQL Server instance unless the replica is first unmounted. However, if you apply a different database name, no such restriction applies. For more examples and information on alternate mounting, refer to the online help.

Mounting SQL Server by instance or by database

When you mount a SQL Server replica, you can specify how you want to set recovery options for SQL Server. Options are as follows:

- ◆ **By instance** — Allows you to set options for all databases within the instance.
- ◆ **By database** — Allows you to set options for each database individually.

Choose the appropriate option depending upon how you want to recover your SQL Server environment.

SQL Server mounts to an alternate location

When you mount the SQL database to an alternate location on an alternate mount host, you select the alternate mount host and the new path where the data should be mounted on that mount host.

Replication Manager mounts the data to the alternate mount host in the new location and depending on the SQL Server mount options,

Replication Manager can use the VDI metadata (which was created at replication time and retrieved from the server) to recover the database on the mount host.

Mounts to alternate locations can allow you to mount several copies of SQL Server data from different servers onto the same mount host, even if the servers have an identical structure on each production host. In SQL Server, there is no restriction on database name because the mounts may use different instances of SQL Server.

SQL Server mounts to the production host

When you choose to mount SQL Server replicas to the production host, select another instance of SQL Server. Then, Replication Manager will not overwrite the production database on that host. Replication Manager uses the VDI metadata file to mount the database to an alternate location on the production machine. The ability to mount onto the same production server can reduce the overall number of servers needed to review data. No extra mount server is necessary. Data can be mounted to an alternate location by changing the root path or by using substitution tables.



CAUTION

When you mount a replica of a SQL Server database to the production server, do not mount it using the same instance of SQL Server that the production database is using. You must use a different instance of SQL Server.

Refer to [Chapter 5, "Mount, Restore, and Recovery,"](#) for more information about possible mount scenarios.

Mount recovery modes with limited recovery

This section describes the mount options for SQL Server replicas. [Figure 175 on page 467](#) illustrates the mount options available with limited recovery.

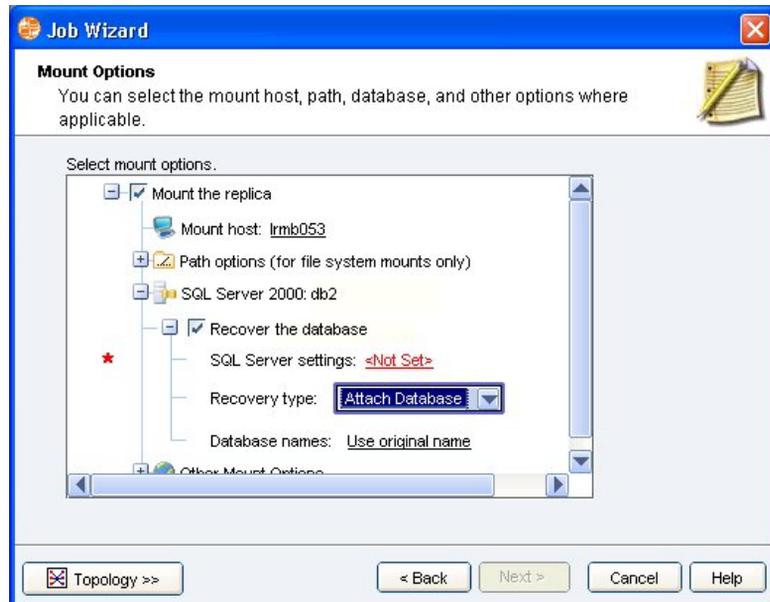


Figure 175 Limited recovery mount options

If the replica was created using consistent split technology with limited recovery (no VDI), there are two options:

- ◆ **Attach Database** — Instructs Replication Manager to mount the file system on which the database files are located, and then attach the database to SQL Server. The **Attach Database** option is only available for full replicas created using consistent-split technology (non-VDI) because all the data necessary to attach the database is part of the replica. Use this option for repurposing; refer to the Caution on [page 468](#).

Note: If UAC is enabled on a Windows Server 2008 or Windows Server 2012 mount host, attach may fail with an 'Access is denied' error. Consider disabling UAC on the mount host or using a SQL Server login to work around the error.

- ◆ **File System Mount** — (Occurs when you clear the **Recover all databases** checkbox). Instructs Replication Manager to mount the file system on which the database files are located. File-system mounts allow users to access the files on that file system, without SQL Server having exclusive access to them. Use this option when

Mount recovery modes with advanced recovery

you want to back up the database, log, and metadata files to tape. This option also allows mounts of databases with filestream data to an alternate location. These options are the only ones available when mounting a crash-consistent RecoverPoint replica.

Replicas created with Advanced Recovery can be mounted as just a file system or VDI can be used to recover the database on the mount host for repurposing. [Figure 162 on page 447](#) illustrates the mount options available with advanced recovery.

In that case the options are:

- ◆ **No Recovery** — Instructs the restore operation not to roll back any uncommitted transactions. When in No Recovery mode, the database is unusable. This option is useful when the Database Administrator needs to restore one or more transaction log backups.
- ◆ **Recovery** — Instructs the restore operation to roll back any uncommitted transactions. After the recovery process, the database is ready for use.
- ◆ **Standby** — Restores files and opens the database in read-only mode. Subsequently, the Database Administrator can manually apply additional transaction log backups taken after the mirror split.

Note: If you are restoring a database from an older version of SQL Server onto a newer SQL Server version, do not use standby mode. If you use standby, the upgrade to the newer version cannot happen and that will result in a failure of the operation.

- ◆ **File System Mount** — Occurs when you clear the **Recover all databases** checkbox. Instructs Replication Manager to mount the file system on which the database files are located. File-system mounts allow users to access the files on that file system, without SQL Server having exclusive access to them. Use this option when you want to back up the database, log, and metadata files to tape. This option also allows mounts of databases with filestream data to an alternate location.



CAUTION

If you plan to remount or restore the SQL Server database in a replica at some time in the future, mount using the File System mount option only. All other options modify the database.

Recovering a mounted database changes the replica and limits the future mount and restore options. You will not be able to apply transaction logs. You will only be able to restore to the point in time of the mount, and will not be able to use Replication Manager's advanced options to recover the database.

SQL Server restrictions on mount options

There are restrictions that prevent you from choosing certain mount options in some environments. This section outlines those restrictions.

RecoverPoint crash-consistent mount restrictions

The following SQL Server mount options are not supported for mounting RecoverPoint crash-consistent replicas:

- ◆ Recovery
- ◆ No Recovery
- ◆ Standby

Effect of SQL replication database option when mounting a replica

SQL Server replicas can be created and mounted while the SQL Server database has the replication database option set, however when mounting using the Standby recovery mode, Replication Manager is unable to change the SQL Server replication flags. In that case, the replication flags can prevent Replication Manager from successfully unmounting the SQL Server database. In order to unmount the database, the **sp_replicationdboption** flag must be cleared manually by the user or subsequent mounts will fail.

Effect of password change on mounted replicas

Before changing the password that Replication Manager uses to access SQL Server, unmount any SQL Server replicas containing databases that have been recovered. If you mount and recover a database, Replication Manager uses the user account and password that was provided during the mount activity to detach the database when performing an unmount. If you change the password for that account while a replica is mounted, the unmount activity will fail.

Transparent data encryption

In SQL Server 2008, transparent data encryption (TDE) provides real-time encryption of data and log files. Data is encrypted before it is written to disk and decrypted when it is read from disk. Replication Manager can replicate SQL Server databases that are TDE enabled. When mounting or restoring such replicas, the TDE certificate and private key must be copied to and made available on the appropriate hosts in order for Replication Manager to recover the

databases. Otherwise, Replication Manager can mount the replica only as filesystems on in No Recovery mode.

Restoring SQL Server replicas

The following section describes the SQL Server modes for restoring replicas:

- ◆ **Database Restores** — Allow you to restore the entire database instead of a subset of the database: This includes the data, log files, and for SQL Server 2008, all full-text catalogs. Of course this option is only available if the replica includes the entire database.
- ◆ **Filegroup Restores** — Allow you to restore a subset of the database at filegroup granularity. However, if more than a single filegroup resides on the same volume, the restore operation will restore all file groups on that volume, so the user should select all filegroups on the volume. Also, for SQL Server 2005 or SQL Server 2008, any full-text catalogs or indexes associated with the filegroup are restored.

EMC recommends that the Database Administrator detach the database before performing a restore of an entire database. If the database is not detached, Replication Manager will detach it to ensure that the database is no in use during the restore.

For SQL Server 2008, restore of just a filestream filegroup is also supported.

- ◆ **Replace Restore option** — (SQL Server 2000 and SQL Server 2005/2008 with replicas created using VDI only) The REPLACE restore option causes SQL Server to skip certain important safety checks that the restore procedure would normally perform. Because these checks are not performed, the REPLACE restore option should be used only by an experienced database administrator:
 - REPLACE does not check if another database already exists with the same name. So, the existing database is deleted.
 - REPLACE does not check whether a log backup has occurred prior to the restore.
 - REPLACE does not check whether the restore is overwriting files of the wrong type.

SQL Server restore prerequisites

When restoring SQL Server replicas, you must meet the following prerequisites:

- ◆ For a filegroup restore, you must put the database in Restrictive mode and back up the transaction log. Replication Manager can perform these tasks if you choose.

Note: You cannot perform a transaction log backup for a database using the simple recovery model; therefore, filegroup backup and restores are not supported for databases using a simple recovery model. Only a full restore is allowed with these databases.

- ◆ (SQL Server 2005 only) If you restore one or more filegroups, all full-text catalogs associated with the restored filegroups will also be restored, providing they reside on supported storage.
- ◆ (SQL Server 2008 only) If you restore one or more filegroups, all full-text indexes associated with the restored databases will also be restored. If you restore one or more databases, any filestream filegroups associated with the restored databases will also be restored.
- ◆ You should detach the SQL Server database to perform a full restore (if the database is not detached when you request a restore, Replication Manager detaches the database automatically).
- ◆ Close all applications that might access the database before you start the restore operation.
- ◆ You should always back up the transaction log before performing any restore operation.

Note: Replicas created with the **Online with limited recovery** option do not allow you to back up the transaction logs before restoring the database. This is because the transaction logs cannot be applied to this type of replica.

For replicas created with the Advanced Recovery Option, Replication Manager supports Recovery, No Recover, Standby, and File System restore modes when restoring a full database.

Note: Remember that File System restore is accomplished by clearing the checkbox entitled **Recover all databases**. This option is not explicitly listed. Manually attach the database after the restore is completed.

Only No Recovery mode is supported when restoring a filegroup or set of filegroups. When in No Recovery mode, the database is unusable because it is in an intermediate and nonrecovered state. No Recovery mode is useful when the Database Administrator needs to restore one or more transaction log backups.

If you want to perform a filegroup restore, EMC recommends that you start by replicating the database and restoring only those filegroups that contain the erroneous data. In that case, if something goes wrong, you always have the option to perform a full restore or to mount the complete replica to an alternate host and manipulate the data further.

In SQL Server, it is possible to restore many transaction log backups one after the other. The intermediate restores are performed in No Recovery mode and the last restore must be done in Recovery mode to make the database usable again.

The last restore must be done in Recovery mode to make the database usable again. If the database has been restored in Replication Manager using No Recovery or Standby mode, then the Database Administrator must recover the database manually after the restore is complete. The restore process differs depending on whether you used consistent split when you created the replica.

If you are using consistent-split technology, refer to Chapter 6 for more information and considerations associated with restoring applications using that technology.

If the replica was created with consistent-split technology and limited recovery (no VDI), file system restore is the only recovery option. A user with appropriate permissions must attach the database manually after the restore is complete.

Celerra restrictions

When restoring part of a SQL Server replica (a filegroup or one database of many), all subsequent replicas are partially marked unrestoreable. The part of the replica that is restored is marked unrestoreable in any newer replicas in the application set.

**SQL Server recovery
(without transaction
logs)**

In SQL Server environments, users must explicitly create transaction log backups. In the absence of transaction log backups, users can recover the database (after a mount or restore) by running the following SQL query in the SQL Server Query Analyzer:

```
RESTORE DATABASE <dbname> WITH RECOVERY
```

where <dbname> is the name of the database.

Note: The command shown above returns the database to the point in time when the replication was taken. After the database is recovered, it is impossible to apply any log backups to roll forward the database.

**SQL Server recovery
(with transaction logs)**

If you have transaction logs backups that you want to apply to the database, restore the database with No Recovery, then restore the transaction log backups.

One way to restore a transaction log is to run the following command in the SQL Server Query Analyzer:

```
RESTORE LOG <dbname>  
FROM DISK = '<backup_filename>'  
WITH NORECOVERY;
```

After the transaction log backups are restored, run the restore as described in “SQL Server recovery (without transaction logs)” on page 473.

**SQL Server recovery
for file system restores**

If you restore a database using the **File System Restore** option, you will have to manually attach the database to SQL Server. How this is done depends on the version of SQL Server and the number of files in the database. If there are fewer than 17 files, use SQL Server Enterprise Manager or Management Studio to attach the database. Otherwise, use the **CREATE DATABASE TSQL** command with the **FOR ATTACH** option.

**Using the rsqlrestore
utility**

Replication Manager includes a SQL Server restore utility called **rsqlrestore**. The **rsqlrestore** utility lets you restore individual SQL Server databases from a tape backup or mounted replica without reverse-syncing the target device over the source device. It can restore a database, filegroup, or file. The utility can restore to the original database or to a new database. The SQL Server VDI metadata that was created as part of the replication activity is required to restore a database using **rsqlrestore**.

Restoring an individual database from a mounted replica is especially useful when you need to recover only one database and do not want to overwrite an entire device, as would happen with a normal Replication Manager restore.

How to run `rmsqlrestore`

The `rmsqlrestore` utility is a command line interface that you run from a command prompt window on the Replication Manager client. `rmsqlrestore` is installed on the client as part of the product installation.

Steps for restoring with `rmsqlrestore`

The exact steps you need to take may differ from the following, but the basic steps are:

1. Log in to the SQL Server system as a user with Administrator rights.
2. Back up the SQL Server transaction log.
3. If restoring a database, take the target SQL Server database offline. (For file or filegroup restore, the database must be online.)
4. If restoring a database, restore the database files (.ldf, .ndf, and .mdf) from tape, or copy them from a mounted replica. You can copy them over the original files or to a new location. (Does not apply to file or filegroup restore.)

Open a command prompt window and cd to:

```
C:\Program Files\EMC\rm\client\bin  
or
```

```
C:\Program Files (x86)\EMC\rm\client\bin
```

5. Run the `rmsqlrestore` command. Complete syntax and sample commands are given in the syntax section following. The basic command syntax is:

```
rmsqlrestore -s <SQLservername> -d <databasename>  
-f <metadata file> -r <recovery_type>
```

6. If necessary, apply transaction logs and recover the database.

Performing the file or filegroup restore

Be sure you understand how restores of files and filegroups work in SQL Server before proceeding:

Note: You cannot use the **rmsqlrestore** utility to restore a SQL Server filegroup if the filegroup name contains non-ASCII characters.

1. Log in to the SQL Server system as a user with Administrator rights.
2. Be sure the target SQL Server database is online.
3. Be sure that the transaction log is backed up.
4. Open a command prompt window and cd to:

```
C:\Program Files\EMC\rm\client\bin
or
C:\Program Files (x86)\EMC\rm\client\bin
```

Run the **rmsqlrestore** command. To restore two files, for example, run:

```
rmsqlrestore -s <SQLservername> -d <databasename>
-f <metadatafile> -lf <logical_filename1>
-lf <logical_filename2> -r norecovery
```

To restore two filegroups, run:

```
rmsqlrestore
-s <SQLservername>
-d <databasename>
-f <metadatafile>
-lf <logical_filename1>
-fg <logical_filegroupname1>
-fg <logical_filegroupname2>
-r norecovery
```

Do not use quiet mode for a file or filegroup restore. You can use **-lf** and **-fg** in the same restore command.

5. When **rmsqlrestore** displays the restore command that it is about to run, verify with **Y** if it is correct.
6. When **rmsqlrestore** prompts, restore the files you are recovering, enter **Y** to continue.

The `rmsqlrestore` command syntax

Table 31 on page 476 lists the command syntax for the `rmsqlrestore` command.

Table 31 The `rmsqlrestore` command options (page 1 of 2)

Option	Description
Required	
-s	SQL Server name including instance name (host\instance).
-f	Metadata filename and location.
-d	Database name.
Connection Types (-E or -U)	
-E	Trusted connection (default).
-U	SQL Server login ID.
-P	Clear text password (used with -U option).
-p	Encrypted password (used with the -U option).
Optional	
-r	Recover option - RECOVERY, NORECOVERY (default), or STANDBY.
-u	Undo filename, required for STANDBY.
-m	Move file. Option has two parameters: <code>logical_file_name</code> and <code>operating_system_file_name</code> . Pathnames must exist. Repeat option for each file, including log files or full text catalog files. If you are restoring to a new database name, use the -m option so you do not overwrite the original files. For example: <pre>-m logicalfilename S:\existingdir\newfilename.mdf</pre>
-fg	Filegroup to restore. Repeat option for each filegroup.
-lf	Logical file to restore. Repeat option for each logical file.

Table 31 The `rmsqlrestore` command options (page 2 of 2)

Option	Description
<code>-e</code>	Displays encrypted password. Not used with other parameters.
<code>-v</code>	Verbose mode.
<code>-q</code>	Quiet mode. Will not ask questions.
<code>-l <log_dir></code>	Creates log files in the specified directory.
<code>-h</code>	Help.

Examples

To restore without applying logs using a trusted connection:

```
rmsqlrestore -s sql1\instance1 -d custinfo
-f "C:\sqlmounts\RMSQLMETADATA_lrma047 test db on R"
-r RECOVERY
```

To restore to a new database name and move files using a SQL login and encrypted password:

```
rmsqlrestore -s sql1\instance1
-d custinfoTest
-f "C:\sqlmounts\RMSQLMETADATA_lrma047 test db on R"
-r RECOVERY
-m custinfo_Data S:\custinfoTest.mdf
-m custinfo_Log T:\custinfoTest.ldf
-U sa -p 1EMC_4roJdyU5;x
```

To get the encrypted password:

```
rmsqlrestore -e unencrypted_password
```

Note: Command options are case-sensitive.

Refer to the SQL Server books online for a description of the **T-SQL RESTORE** command and its options.

Considerations when restoring SQL Server 2005/2008 replicas



In SQL Server 2005 and SQL Server 2008, these additional recommendations and restrictions apply:

- ◆ SQL Server 2005 and SQL Server 2008 require a log backup before any restore (except when using the **REPLACE** option). In other versions of SQL Server this is recommended; with SQL Server 2005/2008 it is required. This restriction does not apply for databases using the Simple Recovery Model.

CAUTION

SQL Server 2005/2008 normally detects if the logs are backed up, however because Replication Manager uses array-based technology to restore the database and log files, SQL Server cannot detect whether the logs have been backed up. It is very important to back up the transaction logs *before* restoring the database.

- ◆ Restoring a replica of a SQL Server 2005/2008 database that is synchronized to another database via a SQL Server mirror session may require special procedures. Refer to [“Restoring SQL Server 2005/2008 mirrored databases” on page 479](#) for more information.
- ◆ Database Snapshots associated with a database must be deleted before Replication Manager can restore a replica to the SQL Server 2005/2008 database. The product checks for database snapshots before running the restore and fails the restore operation if it finds them.

Note: In this context, the term *Database Snapshots* refers to the Microsoft SQL Server 2005/2008 feature called database snapshot technology, not Replication Manager snapshots.

Restoring SQL Server 2005/2008 mirrored databases

To restore a Replication Manager replica in a SQL Server 2005/2008 environment that includes SQL Server mirror sessions, it is important to remember the following important information about how Replication Manager interacts with SQL Server 2005/2008 mirror sessions:

- ◆ Replication Manager cannot restore a replica to a principal or mirror target while the mirror session is active. Stop the mirror session prior to restoring.
- ◆ Replication Manager cannot restore to a different database instance regardless of the state of the mirror, for example, Replication Manager cannot restore to a new principal established after a cluster failover. See [Figure 176 on page 479](#) and [Table 32 on page 480](#) for more information.

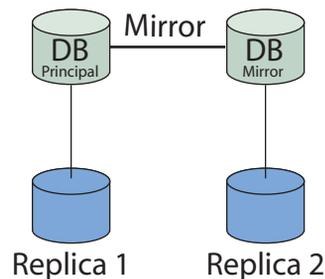


Figure 176 Replication Manager and SQL Server 2005/2008 mirroring

Table 32 on page 480 shows which replicas can be restored to which source locations from Figure 176 on page 479.

Table 32 Restore restrictions when restoring SQL Server 2005/2008 mirrored drives

Replica	Source	Restore comments
Replica 1	DB (Principal)	This can be restored to the source if the mirror session is broken first.
Replica 2	DB (Principal)	This restore is not allowed regardless of the state of the mirror.
Replica 1	DB (Mirror)	This restore is not allowed regardless of the state of the mirror.
Replica 2	DB (Mirror)	This can be restored to the source if the mirror session is broken first.

After restoring the principal database, re-create the mirrored database as follows:

1. Mount the principal database's replica onto the mirror server
2. Use **rmsqlrestore** to recreate the mirrored database.

Refer to "Using the rmsqlrestore utility" on page 473 for more information.

Using SQL Server snapshot functionality

SQL Server includes snapshot functionality. (In this context, snapshot refers to the Microsoft SQL Server VDI snapshot mechanism, not SnapView snapshots or TimeFinder/Snaps). Using snapshot functionality, SQL Server can quickly create a point-in-time backup of a database, which can be moved to alternate storage. When you are using the snapshot mechanism, you should back up the transaction logs as part of your regular maintenance tasks. Snapshot functionality allows you to restore the database to the point in time when the replica was taken. Transaction log restores allow you to roll forward to any point in time.

Restore steps in detail

To restore a database:

1. Restore a replica containing a snapshot backup in No Recovery mode.
2. Apply transaction log backups, if they exist, and recover the database manually. Transaction log backups are created separately by the SQL Server database administrator. These backups are sent to disk or tape and are not backed up by Replication Manager.

Note: You should create a Database Maintenance Plan to coexist with your scheduled replications. The Database Maintenance Plan should back up transaction logs only. The plan should span the intervals between replications to protect changes that occur between those replications. For more information about creating a Database Maintenance Plan, refer to your SQL Server documentation.

3. All but the last transaction log restore must be in No Recovery or Read-Only mode. Recover each of the transaction logs separately through Microsoft SQL Server Enterprise Manager.
4. If you do not implement a maintenance plan that creates transaction log backups, you can recover the database to a point in time that coincides with a replication using the command line interface described in [“SQL Server recovery \(without transaction logs\)”](#) on page 473.

Warm standby server

You can use the snapshot to initialize a secondary database on a standby server. After the standby server has been initialized, it can be maintained using the concept of log shipping. To initialize a standby server:

1. Mount the database to an alternate host in No Recovery or Standby mode.
2. Maintain the database using SQL Server log shipping. Refer to the documentation for SQL Server for more information about log shipping.

Using pre- and post-replication SQL Server scripts

The Replication Manager SQL Server Agent performs certain default actions on the production database before and after splitting the mirror. You can add customized actions with your own user-supplied scripts.

Using pre- and post-replication scripts

To use a pre-replication and post-replication script:

1. Name the script and its location (by specifying the full pathname) while configuring the job in the console.
2. Ensure that the script will *not* take the database offline.
3. Ensure the script is owned by the SQL Server user and is executable by the SQL Server user and group.

General guidelines for SQL Server scripts

The following guidelines can help you produce appropriate scripts for SQL Server:

- ◆ Scripts must be in .bat or .exe format.
- ◆ Suppress output from the script. The following techniques work for selected popular environments:

- **Windows** — Add @echo off to the first line of the .bat script.
- **UNIX sh** — Redirect output from the script as shown:

```
Prog.sh > /dev/null 2>&1
```

- **UNIX csh** — Redirect output from the script as shown:

```
Prog.csh >& /dev/null
```

Note: Suppressing script output prevents invalid characters from being added to replica history. Invalid characters in the history cause Replication Manager to stop responding when mounting or viewing a replica. If you need output from the script, set up and use a log file instead of directing output to standard out.

- ◆ To run **transact-sql** against the database, ensure it is in the following format within the .bat file:

```
osql -E -S <instance_name> -Q "<sql_query>"
```

where <instance_name> is the name of the SQL instance to run against and <sql_query> is the SQL statement in transact SQL.

Sample replication scripts

The sample in [Figure 177 on page 484](#) is a basic pre-replication script that prints the database statistics before the replication occurs. Your scripts can add any additional actions that you want to perform before replication.

```
osql /E /S <server_name> /d "master" /Q "sp_helpdb
'<database_name>' /o "c:\prescript.dat" /w 600 /b
if ERRORLEVEL 1 goto FAILED
osql /E /S <server_name> /d "<database_name>" /Q
"sp_helpuser" /w 600 /b >> "c:\prescript.dat
if ERRORLEVEL 1 goto FAILED
goto SUCCESS
:FAILED
echo An error occurred >"c:/error.log"
exit
:SUCCESS
echo The script finished successfully.
```

Figure 177 Sample replication script

In the script shown, replace *<server_name>* with the name of the host and *<database_name>* with the name of the database that you want to gather information about.

Using callout scripts in a SQL Server environment

If you want to utilize callout scripts in a SQL Server environment you must meet one of the following conditions:

- ◆ Run Replication Manager with a user that has Administrator privileges to the SQL Server instance.
- ◆ Ensure that the non-admin local or domain user that is running Replication Manager has “Full Control Permissions” on the Replication Manager registry key.

Considerations for working with SQL Server in a cluster

Replication Manager supports SQL Server installed in a cluster running Microsoft Cluster Service (MSCS):

- ◆ To create application sets, register the virtual server network name for the SQL Server instance on the Replication Manager Server.

For Windows 2003, use the Cluster Administrator to discover which resource group is associated with SQL Server:

- a. View the Properties to discover the network name and corresponding IP address.
- b. Register the host corresponding to that network name as a Replication Manager host.
- c. Select that host when you configure the application set for SQL Server replicas.

For Windows 2008 and Windows 2012, use the Failover Cluster Management MMC snap-in to get the network name or IP address of the SQL Server virtual server:

- a. Expand the cluster node.
- b. Expand Service and Applications.
- c. Select the SQL Server resource group.

The network name and IP address are displayed under the **Server Name** heading in the pane on the right.

- ◆ Mounting back to the same cluster is not supported. You need a mount host to mount the replica for backup.

See the *EMC Replication Manager Administrator's Guide* for more information about how to set up Replication Manager in a cluster environment.

The following procedure describes the restore of a SQL Server replica in an MSCS environment on Windows Server 2003 that does *not* have the Microsoft Extended Maintenance Mode hotfix installed.

To restore Windows 2003 MSCS Cluster without Extended Cluster Maintenance Mode:

1. Using the Enterprise Manager or Query Analyzer, detach the database.

2. Using Cluster Administrator, take the SQL Server Resource Group offline.
3. Remove the disk dependencies for the database and log files that you are restoring. To do this, open the correct resource group and modify the resource with the Resource Type of SQL Server:
 - a. Right-click the SQL Server resource, and then click **Properties**.
 - b. On the **Dependencies** tab in **Properties**, click **Modify**.
 - c. Select all disk resources that you plan to restore and click the ← (left arrow) button to remove the resource dependencies.
 - d. Click **OK** twice to save the changes.
4. Delete the same physical disk resources from the SQL Server resource group.
5. Bring the remaining Physical Disk resources and Network Name resources online. Do not bring any other resources online.
6. Start SQL Server as a nonclustered application:
 - a. Select **Run** from the **Start** menu, type **cmd** in the dialog box and click **OK**.
 - b. Type:

```
net start MSSQLSERVER  
or  
net start MSSQL$<instancename>
```
7. Using the Replication Manager Console, restore the database in the appropriate recovery mode.
8. Use the Enterprise Manager to verify that the database has been restored properly.
9. Stop the SQL Server service:
 - a. Select **Run** from the **Start** menu, type **cmd** in the dialog box and click **OK**.
 - b. Type:

```
net stop MSSQLSERVER  
or  
net stop MSSQL$<instancename>
```
10. Using the Cluster Administrator, add the Physical Disk resources back into the SQL Server resource group.

11. Add the disk dependencies back for the database and log files that you are restoring. To do this, open the correct resource group and modify the resource with the Resource Type of SQL Server:
 - a. Right-click the SQL Server resource, and then click **Properties**.
 - b. On the **Dependencies** tab in **Properties**, click **Modify**.
 - c. Select all disk resources that you plan to restore and click the → (right arrow) button to add the resource dependencies.
 - d. Click **OK** twice to save the changes.
12. Using the Cluster Administrator, add the Physical Disk resources back into the SQL Server resource group.
13. Add the disk dependencies back to the SQL Server Resource Type resource.
14. Recover transaction logs as needed.

Microsoft Exchange Procedures

This appendix covers the specifics of the Microsoft Exchange support, including the following sections:

- ◆ Setting up Exchange hosts 490
- ◆ Exchange 2010/Exchange 2013 and Replication Manager 495
- ◆ Importing a replica from a backup 516
- ◆ Exchange 2007 and Replication Manager 520
- ◆ Item level restore 553
- ◆ Using pre- and post-replication Exchange scripts 555
- ◆ Considerations for Exchange in a cluster 556
- ◆ Troubleshooting Exchange issues 558

Setting up Exchange hosts

Replication Manager can replicate data from any of the following Microsoft Exchange environments:

- ◆ Microsoft Exchange 2010/Exchange 2013 Database Availability Groups (DAGs) including:
 - Native Exchange DAGs with active and passive copies
 - Third-party replication enabled DAGs using EMC Replication Enabler for Exchange 2010 (REE) and CLARiiON or VNX MirrorView/S
- ◆ Microsoft Exchange 2010/Exchange 2013 standalone servers
- ◆ Microsoft Exchange 2007 high availability:
 - Single Copy Cluster (SCC) environments can be replicated.
 - Cluster Continuous Replication (CCR) both active and passive copies of the storage group can be replicated.
 - Local Continuous Replication (LCR) active copies of the storage group can be replicated.
 - Standby Continuous Replication (SCR) active copies of the storage group can be replicated.
 - SCR targets cannot be replicated.
- ◆ Microsoft Exchange 2007 standalone servers

This section describes how to setup your Exchange configuration to meet the prerequisites necessary to use Replication Manager on Exchange production and mount hosts. It also discusses permissions that you need to grant in order to run Replication Manager successfully.

Note: Refer to the *EMC Replication Manager Support Matrix* for updated information on required service packs and supported operating systems. Specific SQL Server Service Packs are required. Refer to the *EMC Replication Manager Support Matrix* for the latest support information. To access the *EMC Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

Setting up Exchange production hosts

Exchange production host setup differs depending upon what version of Exchange you are running. [Table 33 on page 491](#) describes setup requirements and highlights the version of Exchange to which each requirement applies.

Table 33 Exchange production host setup prerequisites by Exchange version (page 1 of 2)

Production host setup requirement	Exchange 2007	Exchange 2010	Exchange 2013
Partitioned disks are not supported	X	X	X
Do not use nested mountpoints; for example, if your logs are on L: and the database is on the mountpoint called L:\SG1DB1, Replication Manager will not be able to mount or restore the replica.	X	X	X
The following files must be stored on a supported storage array to ensure that they will be replicated: EDB files, log files, system files, and streaming files (if applicable).	X	X	X
Circular logging must be disabled. For more information, refer to “Disabling circular logging” on page 521 .	X	X	X
For better restore flexibility, Exchange databases should not share volumes.	X	X	X
For better restore flexibility, Exchange storage groups should not share volumes.	X		
For Exchange in a RecoverPoint environment, restore granularity is at the consistency group level.	X	X	X
Volumes that contain Exchange databases should contain no other data.		X	X
Volumes that contain Exchange storage groups or databases should contain no other data. This includes the SMTP Queue.	X		
Microsoft Exchange is not supported on VMware virtual disks (VMDKs).	X	X	X
The database file and transaction logs for an Exchange database can reside on the same volume if there is more than one copy of the database.		X	X

Table 33 Exchange production host setup prerequisites by Exchange version (page 2 of 2)

Production host setup requirement	Exchange 2007	Exchange 2010	Exchange 2013
Transaction log and system files must be on different volumes from the databases.	X		
The system path location and transaction log location must be the same for each storage group.			
EDB and STM files must be stored in the same directory on the same volume to ensure that they will be replicated at the same time.			

Setting up Exchange mount hosts

This section helps you verify that your Exchange configuration meets the prerequisites necessary to use Replication Manager on a mount host. The most common reasons to mount an Exchange replica are:

- ◆ Check the consistency of the Exchange replica.
- ◆ Back up an Exchange replica.
- ◆ Recover a mailbox from an Exchange replica.

[Table 34 on page 493](#) describes the setup requirements for Exchange in a Replication Manager environment.

Table 34 Exchange mount host setup prerequisites by Exchange version

Mount host setup requirement	Exchange 2007	Exchange 2010	Exchange 2013
The version of Windows on the mount host must match the version of Windows on the production host. For example, you cannot mount a replica created on Windows Server 2008 to a Windows Server 2003 host and vice versa. However, you can mount a replica created on an x64 edition of the operating system to a server running the x86 edition.	X	X	X
Only one version of the Exchange Management Tools can be installed.	X	X	X
The Exchange Management Tools on the mount host must be the same version and service pack level as the production host (except in the case outlined below).	X	X	X
A single mount host with Exchange 2010 Management Tools can be used to run consistency check for both Exchange 2010 and Exchange 2007.	X	X	
A single mount host with Exchange 2013 Management Tools can be used to run consistency check for Exchange 2007, Exchange 2010 and Exchange 2013.	X	X	X
If you plan to recover mailboxes, the Exchange Server must be installed on the mount host. In that case, the Exchange Server installed on the mount host must be in the same administrative group as the production server.	X	X	X
The mount host must have visibility to the same storage array from which the original replica was created. In the case of a DAG or CCR environment, you may need more than one mount host if your replicas are created from different copies of the database.	X	X	X
In the case of native DAG environments, Replication Manager can mount a replica to a server that is part of the production DAG.		X	X
For DAG with REE enabled, mounts back to the production DAG are not supported.		X	

Setting up Exchange permissions for Exchange 2010 / Exchange 2013

The following permissions are required for accounts that Replication Manager uses:

1. Perform the steps outlined in the first two bullets below, or alternatively perform the step in the third bullet:
 - On Exchange 2010/Exchange 2013 stand-alone servers: Domain user account with the Databases role.
New-ManagementRoleAssignment -Role "Databases" -User <user account>
 - On Exchange 2010/Exchange 2013 DAG servers: Domain user account with the Database and Database Copies Group roles.
New-ManagementRoleAssignment -Role "Databases" -User <user account>
New-ManagementRoleAssignment -Role "Database Copies" -User <user account>

or

 - The customer can add the user to the Server Manager group instead of adding the Roles.
2. On Exchange 2010/Exchange 2013 mount hosts, provide a Domain user account that is a member of the local Administrators group.

Setting up Exchange permissions for Exchange 2007

During installation with Exchange 2007, you need to specify credentials for domain account that is both a member of the local Administrators group and the Exchange Server Administrator role.

Exchange 2010/Exchange 2013 and Replication Manager

Replication Manager supports Exchange 2010/Exchange 2013 in standalone, Microsoft's native DAG or EMC's array-based DAG (with third-party replication enabled). If you are operating in an Exchange 2010/Exchange 2013 environment, refer to this section for coverage of the following topics:

- ◆ Preparing the Exchange 2010/Exchange 2013 environment for Replication Manager
- ◆ Creating Exchange 2010/Exchange 2013 replicas
- ◆ Mounting Exchange 2010/Exchange 2013 replicas
- ◆ Restoring Exchange 2010/Exchange 2013 replicas
- ◆ Considerations when working in a DAG environment

For information on troubleshooting in an Exchange environment, consult [“Troubleshooting Exchange issues”](#) on page 558.

Preparing the Exchange 2010/Exchange 2013 environment

Before you can use Replication Manager in the Exchange 2010/Exchange 2013 environment, you must perform some preparatory steps.

- ◆ Configure hosts and permissions
- ◆ Disable circular logging
- ◆ Enabling Windows authentication is not required

Refer to the following sections for details.

Configure hosts and permissions

To configure hosts and permissions, refer to [“Setting up Exchange production hosts”](#) on page 491, [“Setting up Exchange mount hosts”](#) on page 492, and [“Setting up Exchange permissions for Exchange 2010 / Exchange 2013”](#) on page 494 for information on preparing your environment to use Replication Manager with Exchange 2010/Exchange 2013.

Disabling circular logging

Circular logging is off by default for Exchange 2010/Exchange 2013 databases. To verify that circular logging is disabled:

1. Open the Exchange Management Console.
2. Navigate to **Microsoft Exchange > Microsoft Exchange On-Premises > Organization Configuration** and select **Mailbox**.

RecoverPoint restrictions

3. Open the properties of a database and select the **Maintenance** tab.
4. Verify that the **Enable circular logging** checkbox is not selected.

Note the following restrictions related to restoring RecoverPoint replicas:

- ◆ Replication Manager does not support RecoverPoint and DAGs using EMC Replication Enabler for Exchange 2010/Exchange 2013 (REE).
- ◆ RecoverPoint crash-consistent restores of Exchange 2010/Exchange 2013 are not supported. Exchange 2010/Exchange 2013 restores must be performed from an application-consistent replica.

Despite this inability to restore a crash-consistent replica, these replicas can still be used to minimize data loss.

To use log files from a crash-consistent replica, follow these steps:

1. Restore a database from an application-consistent replica without recovering it.
2. Use Replication Manager to mount a replica from a newer point in time.
3. Copy the newer log files to the production log volume.
4. Use **ESEUTIL /k Enn** (Enn is the log prefix for the database) to check the logs, then recover and mount the database.

Creating Exchange 2010/Exchange 2013 replicas

Replication Manager has specialized agents that enable it to:

- ◆ Create online full and copy replicas of Exchange databases and logs using Microsoft Volume Shadow Copy Services (VSS).
- ◆ Create replicas of Exchange databases protected as part of a native DAG or a DAG with REE enabled. Native DAG replicas can be created whether or not the source is an active or passive copy of the database, while replicas of DAG with REE enabled can only be created if the source is an active copy of the database.
- ◆ Check the consistency of replicated data.

Before each replication occurs, Replication Manager:

- ◆ Discovers the location of the data to replicate.
- ◆ Identifies pathnames for the database and for log files.

Support for VSS

Microsoft Volume Shadow Copy Service (VSS) coordinates with business applications, backup applications, and storage hardware to enable application-aware data management. VSS is the infrastructure that enables Replication Manager to create application-aware replicas. During replication, Replication Manager coordinates with VSS and Exchange 2010/Exchange 2013 to create a shadow copy, which is a point-in-time copy of the volumes that contain the data, logs, and system files for Exchange 2010/Exchange 2013 databases.

Replication Manager coordinates with VSS and Exchange 2010/Exchange 2013 to quiesce the input/output to the databases during replication, and then resume the flow of data after the replication is complete. During a restore, Replication Manager coordinates with VSS and Exchange to recover the point-in-time shadow copy.

Exchange data objects in the replica

The Exchange Administrator can configure an application set to replicate one or more database from the Exchange 2010/Exchange 2013 Server. Exchange 2010/Exchange 2013 can have up to 100 databases per mailbox server.

When you restore a Microsoft Exchange 2010/Exchange 2013 replica, you can choose to restore one or more database(s), with or without logs, from the replica. You choose what to restore by selecting the appropriate components from a tree as you move through the Restore Wizard.

In Exchange 2010/Exchange 2013, it is best to arrange the data so that volumes used for Exchange 2010/Exchange 2013 data do not share physical volumes with other data that is not associated with the database. Isolating Exchange 2010/Exchange 2013 data prevents potential problems when you restore data from a replica to the production Exchange server.

For Exchange 2010/Exchange 2013 environments with more than one copy of the Exchange database, you can choose whether or not to allow the database and logs to reside on the same volume. The option to allow this can be configured when creating or modifying application sets.

Replicating Exchange 2010 /Exchange 2013 with VSS

If you are creating replicas of Exchange 2010/Exchange 2013 on Windows 2008 or Windows 2012, Replication Manager uses VSS to perform a consistent online replication. You can use the Advanced options screen of the Job Wizard to select a replication option:

- ◆ **Online Full** — Replication Manager replicates the databases, transaction logs, and checkpoint files, and then runs a consistency check to verify the consistency of the databases and logs. If the consistency check completes successfully, Replication Manager instructs Exchange to truncate the logs so that only changes that are uncommitted to the database remain.
- ◆ **Online Copy** — Replication Manager replicates the databases, transaction logs, and checkpoint files in the same way as it does during a Online Full option, however, it does not truncate the logs. Online Copy replications are often intended for testing and diagnostic purposes only.

Limitation in the Microsoft VSS architecture

A limitation in the VSS architecture prevents successful mounts and restores when the root drive letter has mount points on it and they are all included in the job. For instance, if the log and system files are on L:\ and the mailbox stores are on L:\SG1DBMP (where SG1DBMP is a mount point), mounts and restores will fail.

Replicating Exchange 2010/Exchange 2013 on VMware VMFS

Replication Manager supports VMware's use of VSS with VM snapshots when VSphere 4.1 or later is installed and VMware Tools are present on the virtual machine on the VMFS you are replicating. When Replication Manager replicates a VMFS on which Exchange resides, the VM snapshot in the VMFS replica is consistent. From an Exchange point of view, however, it is a VSS copy replica, so logs are not truncated, a consistency check is not run, and the database is not marked as backed up.

Replicating Exchange 2010/Exchange 2013 high availability environments

Replication Manager supports replication of Exchange databases that are part of a native DAG (Exchange 2010/Exchange 2013) or a DAG with REE enabled (Exchange 2010). This functionality protects Exchange databases by using a subset of Windows Failover Clustering technologies. Failover management for native DAG occurs entirely within Exchange, while failover for DAGs with REE is managed by REE. There is no shared storage in a native DAG environment. [Figure 178 on page 499](#) illustrates a native DAG environment, while [Figure 179 on page 500](#) illustrates DAG with REE enabled.

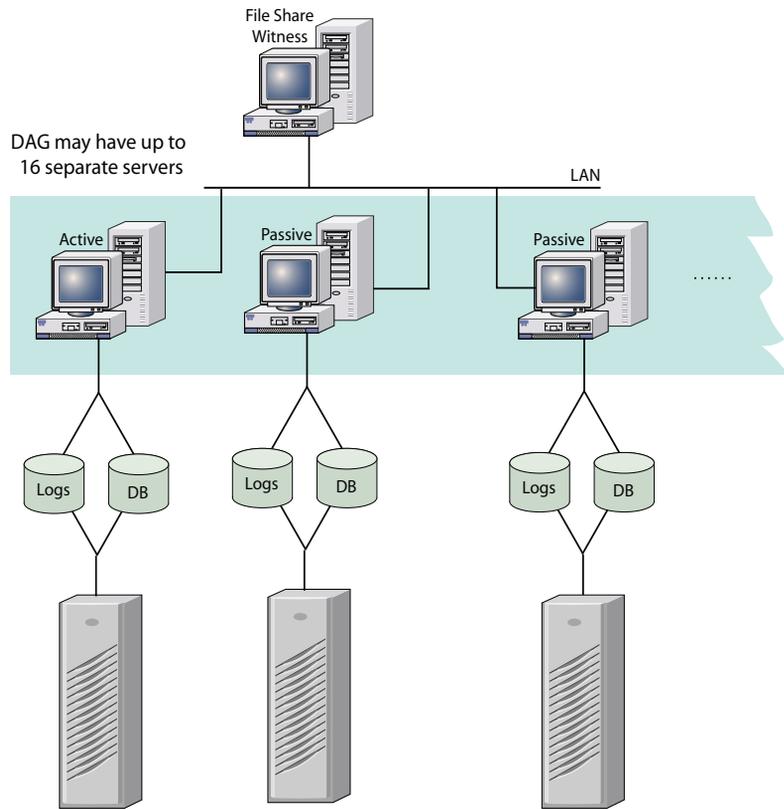


Figure 178 DAG (native environment)

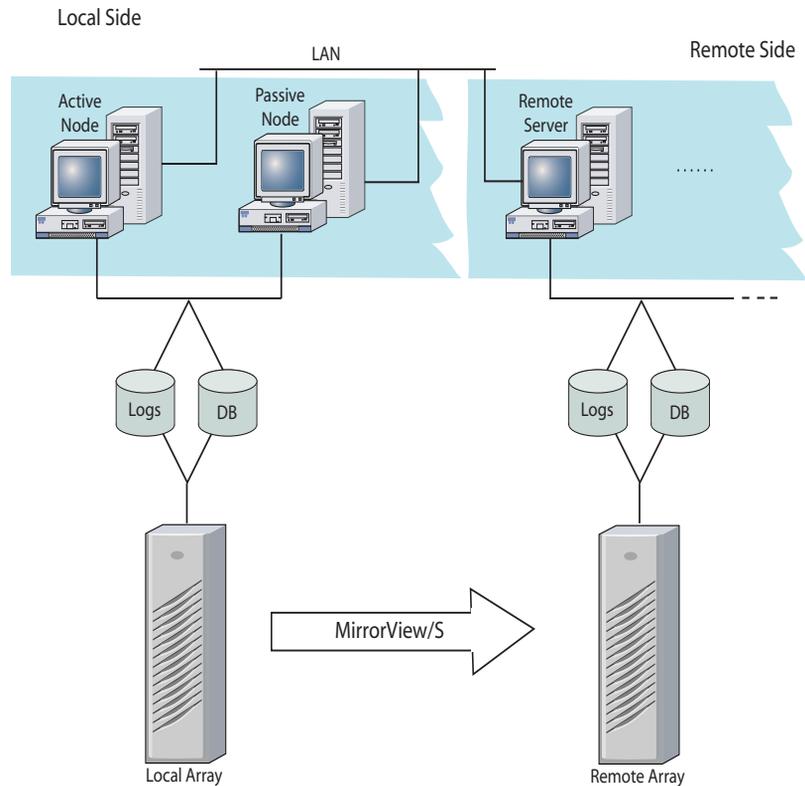


Figure 179 DAG with REE enabled

DAG replication techniques

If you are replicating data protected by a DAG (native environment), follow these guidelines when configuring your environment for use with Replication Manager:

- ◆ Choose a server within the DAG from which to create the replica.
- ◆ Choose one or more databases with a copy on that server to include in the replica.

In native DAGs, Replication Manager creates a replica of each selected database, regardless of its state (database may be active or passive during the replication). Replication Manager's DAG support does not currently offer you the ability to specify a state in order to limit whether a database copy is replicated.

Regardless of the state of the database when it was replicated, a restore can only occur to an active copy of the database. Refer to

Replication techniques for DAG with REE enabled

“Restoring Exchange 2010/Exchange 2013 replicas” on page 507 for additional information about how to restore replicas to a DAG.

Replication Manager supports DAGs with REE in one of the following ways:

- ◆ **Using a dedicated server** — In this configuration, Replication Manager always attempts to run the job from a specific server and if the database is not active on that server the job fails.
- ◆ **Using dynamic jobs** — Replication Manager determines which server is currently hosting the selected database and automatically runs the job from that server.

In the REE environment, Replication Manager can only create replicas or restore from a server hosting the active copy of the database.

If you are replicating data protected by a DAG with REE enabled, follow these guidelines when configuring your environment for use with Replication Manager:

- ◆ For information on configuring REE, consult the following documents entitled *EMC Replication Enabler for Microsoft Exchange Server 2010 Installation and Configuration Guide*.
- ◆ For dynamic job support install the Replication Manager agent on all servers that are members of the DAG (both local and remote).
- ◆ For dynamic job support register all of the servers in the DAG as Replication Manager hosts.
- ◆ Register the virtual server for the DAG, (the DAG name). This server is sometimes used to find the server that currently hosts the active copy of a database. This is required for dynamic job support.
- ◆ If one server is registered using a fully qualified network name, all servers in the DAG must be registered with FQN names. IP addresses are not supported with the dynamic job feature. They are only supported when using a static host to run jobs.
- ◆ Register both the local and remote storage array with Replication Manager.
- ◆ Limit application sets to only one Exchange 2010 database per application set because failover is at database granularity.

- ◆ When REE is enabled, Replication Manager does not support mounting a replica back into the DAG where it was created. You can mount to a server in another DAG with REE enabled or a stand-alone server.
- ◆ Replication Manager supports the creation of clones and snaps from primary or secondary MirrorView/S LUNs.
- ◆ Replication Manager can restore clone or snap replicas created from MirrorView primary LUNs.
- ◆ Replicas created in a DAG with REE enabled cannot be used as the source for a copy job.

In DAGs with REE enabled, passive databases are not on mounted filesystems, and therefore cannot be accessed by Replication Manager for any operation. DAGs with REE enabled operations can only occur on active databases.

Replication Manager can only restore replicas of MirrorView/S primary images, not MirrorView/S secondary images. Refer to [“Restoring Exchange 2010/Exchange 2013 replicas” on page 507](#) for additional information about how to restore replicas to a DAG.

Creating and running dynamic jobs in an REE environment

The Job Name and Settings panel in the job wizard has been changed to support DAGs with REE enabled. The new options are:

- ◆ **Replication Target:** A DAG with REE enabled has two arrays, one local and one remote. There is a MirrorView/S session between the arrays that is used to keep the database and log LUNs up to date at the remote, secondary site. When you create a job, you select the target array instead of a replication source. The replicas are always created in the replication target array.
- ◆ **Replicate using Exchange Mailbox Server:** This option is used to specify which host the job will run on. As explained in [“Replication techniques for DAG with REE enabled” on page 501](#), for dynamic jobs, select the Any server hosting the active database option. To run a job on a dedicated server, pick the server that was used when you created the application set.
- ◆ Snap and clone are the only types of replicas that you can select

When you run a dynamic job, the job is always run on the server that hosts the active copy of the database and the replica is always created in the array that was selected in the job. Where the job runs can change, but where the replica is created does not.

When a dynamic job runs, Replication Manager first looks for the active copy of the database. If the database has moved to another server, Replication Manager updates the server name in the application set and job; the job will then run on the new server. Next Replication Manager checks to see if it needs to change the replica source type. This happens because the personality of the LUNs in the target array changes when the active copy of a database moves from one array to another. The primary images become the secondary images, and vice versa.

When you run a job after the database has moved to the other array, the replication technology will change. Replication Manager changes the job to reflect the personality swap; the change can be seen in the Console under the job and replica properties.

If you select the *local* array when creating a job, at run time the job will:

- ◆ Create a local clone or snap when the active copy of the database resides on the local array. In the Replication Manager Console, the Technology property of the replica and job reads SnapView Clone or SnapView Snap. Replication Manager creates a replica of the primary LUNs in this case. These are the LUNs that the active database currently resides on.
- ◆ Create a remote clone or snap when the active copy of the database has failed over to the remote array. The Technology of the replica and job will say Remote MirrorView Clone or Remote MirrorView Snap. The replica type changes because Replication Manager always uses the same target array to create replicas, and that target array is now remote from the array that hosts the active copy of the database.

When the database fails over to the remote array, the LUNs in that array are promoted to be the primary images; the LUNs in the local array become the secondary images. Replication Manager still creates the replica in the local array which was selected in the job -- in this case it creates a replica of the secondary LUNs. VSS is used to freeze and thaw the database I/O to the primary LUNs, so this type of replica is still considered application consistent. However, it is not restorable by Replication Manager.

If you select the *remote* array when you create a job, at run time the job will:

- ◆ Create a local clone or snap when the active copy of the database resides on the remote array. In the Replication Manager Console, the Technology of the replica and job will read SnapView Clone or SnapView Snap.
- ◆ Create a remote clone or snap when the active copy of the database resides on the local array. In this case the Technology property of the replica and job will read Remote MirrorView Clone or Remote MirrorView Snap.

Replica expiration is still done at the job level, even if the job runs on different servers.

Managing Exchange errors

Exchange logs the following errors in the Application event log when they occur:

- ◆ **-1018** — The database tried and failed to verify information about a particular page in the database.
- ◆ **-1019** — This is similar to a -1018 error but indicates that the accessed page has returned an invalid page number (usually all zeros) rather than an invalid checksum.
- ◆ **-1022** — This is indicative of major hardware problems, particularly disk subsystem problems. If the database engine requests a page from disk but instead receives an error from the I/O subsystem, a -1022 error results.
- ◆ **447** — This is indicative of corruption in the logical database structure. This accompanies a message stating that the information store terminated abnormally.
- ◆ **448** — This error denotes an inconsistency or corruption in a table in the jet database. This accompanies a message stating that an information store data inconsistency has been detected in a table.

Running a consistency check

Exchange 2010/Exchange 2013 jobs use the Consistency Check API to check the database and log files. If you select an Online Full replica that truncates transaction logs, you should mount the replica as part of the job, and run a consistency check. If the consistency check or mount fails for any reason or device errors occur, the logs are not truncated and the replica fails.

To mount the Exchange database for consistency checking only, you need to install the appropriate tools on the mount host. When using Exchange 2010/Exchange 2013, install the Microsoft Exchange

2010/Exchange 2013 Management Tools. A single mount host with the Exchange 2010/Exchange 2013 Management Tools can be used to run the consistency check for both Exchange 2010/Exchange 2013 and Exchange 2007.

Consistency check advanced features

Replication Manager offers some advanced features that change how consistency checks are executed. Enabling these features can impact performance. For most users the default settings are sufficient. These advanced features include:

- ◆ **Minimize log checking** — Choosing this option from the mount options panel of the Job Wizard speeds up the log checking by instructing the consistency checking software to check only those logs that are required to recover the database. Selecting this option improves the performance of the consistency check.

If you clear that checkbox, then consistency check will be performed on all of the database's logs.

This command instructs Replication Manager to only check a subset of the Exchange logs that are included in the replica. If your backup window is small, you may find this option useful. However, the replica contains logs that have not been checked for consistency. If you attempt to restore the log volume, you may find that some log files are corrupt or the log sequence is not complete. Before restoring the log volume, you should mount the replica and run **eseutil /k Enn** against the log path.

For maximum protection, clear **Minimize log checking**. For maximum performance, select it.

You must also set the Working directory for Exchange 2010/Exchange 2013 replicas. The working directory is the directory to which the required log files will be copied in order to check them.

The **Minimize log checking** option is not available when the consistency method is Online - Differential.

- ◆ **Parallel consistency checks** — Consistency checks can run against the databases in parallel (all databases for a storage group at the same time) or sequentially. If the databases all reside on the same LUN, use the sequential option.
- ◆ **I/O throttling during consistency checks** — In Exchange 2010/Exchange 2013, consistency checks can be paused to slow down the I/Os during the consistency checking operation.

Exchange 2010/Exchange 2013 throttling options allow you to specify the number of I/O's after which to pause and the duration of the pause.

- ◆ **Do not truncate logs if mount or consistency check fails** — When this option is checked, and mount or consistency checks fail on the replica, Replication Manager does not truncate the transaction logs from the production database. If you clear this checkbox, Replication Manager truncates the logs even though consistency check or mount fails.

This option is not available for on-demand mount and linked copy jobs. It is available only if **Replica Consistency Method** is Online - Full.

If **Fail the replica if the mount fails** mount option is selected, then this option is selected (and made unavailable) by default.

Mounting Exchange 2010/Exchange 2013 replicas

Replication Manager can:

- ◆ Mount a replica to an alternate host using the same file paths as on the production host.
- ◆ Mount a replica on an alternate host in a new location (determined by adding an alternate path to the start of the path).
- ◆ Mount a replica on an alternate host in a new location (determined by path mapping). Refer to the following note.
- ◆ Mount a replica to the production host in a new location determined by an alternate path or path mapping. Refer to the following note.

Note: Specific information about how alternate paths and path mapping work can be found at [“Mounting using alternate path” on page 174](#) or [“Mounting using path mapping” on page 177](#).

Replication Manager can mount a replica to an alternate host that may or may not have Exchange Server installed.

It can perform:

- ◆ Mounts of new replicas as part of the replication job
- ◆ On-demand mounts of existing replicas

Replicas created in a native DAG can be mounted:

- ◆ to an alternate host
- ◆ to a server in another DAG (native or with REE enabled). DAG REE is supported only on Exchange 2010.
- ◆ to another server in the same native DAG

Replicas created in a DAG with REE enabled (on Exchange 2010) can be mounted:

- ◆ to an alternate host
- ◆ to a server in another DAG (native or with REE enabled)

Restoring Exchange 2010/Exchange 2013 replicas

When you restore a Microsoft Exchange full or copy replicas, you can choose to restore any of the following:

- ◆ One or more databases which include the logs.
- ◆ One or more database files (.edb).

- ◆ One or more active or passive databases or database files, if the server is a member of a native DAG.
- ◆ One or more active database or database files, if the server is a member of a DAG with third-party replication enabled.

“Restoring in an Exchange 2010/Exchange 2013 DAG environment” on page 511 provides more information about restoring to a DAG.

When you are restoring just a database file, verify that the transaction log files needed for recovery are present. An unbroken sequence is required. To determine the minimum required range of logs, run the following command against each database after the restore and before running recovery:

```
ESEUTIL /mh <database name>
```

If the database is the active copy, it must first be dismounted in order to run the **ESEUTIL** command successfully.

Look for the **Log Required** information in the **ESEUTIL** output.

If you choose to restore one or more database(s), Replication Manager restores the databases, transaction logs, and checkpoint files that make up the database(s).

When you restore a Microsoft Exchange database, the restored files include a checkpoint file (.chk). The checkpoint file records the location in the transaction log files of the last complete transaction that Exchange wrote to the database. If you do not select the Recover and Mount option when restoring, it may be necessary to delete this checkpoint file before recovering the database(s).



CAUTION

If you choose to restore a database or transaction log volume, the restore will overwrite any logs created since the replica was created. That means after the restore, your database reflects the point in time when the replica was created.

If you want to preserve logs created since the replica, you should choose to restore only the database(s), preventing Replication Manager from restoring older logs over your newer logs, or make a copy of the current log files on another volume.

Restoring with VSS

Replication Manager coordinates with VSS and Exchange 2010/Exchange 2013 to restore the volumes that contain the databases and logs that you have selected to recover. The database

and log file paths must be an exact match in order for this operation to be successful. When restoring a VSS replication, you must restore it to the same location on the server from where it originated. The Exchange 2010/Exchange 2013 VSS Writer is always involved in the restore and recovery of the Exchange databases. Therefore, Replication Manager selects the **Recover and mount databases** checkbox by default.

Restoring with Recover and Mount Database option selected

When the **Recover and mount databases** checkbox is selected for an Exchange 2010/Exchange 2013 restore, the VSS Writer performs the following tasks:

- ◆ Verifies that all of the required logs are available
- ◆ Verifies that there are no gaps in the log sequence
- ◆ Deletes the checkpoint file (as needed)
- ◆ Recovers the databases

Restoring without Recover and Mount Database option

If the **Recover and mount databases** checkbox is not selected, the Exchange 2010/Exchange 2013 VSS Writer will not delete the checkpoint file. You will be responsible for the following tasks:

- ◆ Deleting the checkpoint file (Enn.chk)
- ◆ Deleting the restore.env file (EnnRESTORE.env)
- ◆ Recovering the databases

To recover Exchange databases after this type of restore:

1. Use Replication Manager to restore the replica.
2. Run the **ESEUTIL** command as follows for each database restored:

```
eseutil /r E<nn> /l <logpath> /s <chkpt file path> /d  
< database path>
```

3. To recover Exchange 2010/Exchange 2013 databases, use the Exchange Management Console to mount all the restored databases.

Selecting full or partial restores

The first panel of the Restore Wizard allows you to choose all or part of a replica to restore. This section describes what will be restored if you make certain selections on that panel.

To restore the entire contents of the application set, select the top node of the Replica tree.

Restoring datafiles in Exchange 2010/Exchange 2013

To restore datafiles for one or more databases, select the individual database file.

Restoring log files in Exchange 2010/Exchange 2013

Logs can only be restored as part of a full database restore.

Restore considerations

If the logs are selected for restore as part of a database restore, they will overwrite newer logs created since the replica was created. That means your database will be rolled back to the point in time when the replica was created.

Note: If you want to preserve logs as they exist on the production host, you should choose to restore only the database files to prevent Replication Manager from restoring older logs over newer logs.

If you want to restore individual databases, each database must be stored on a separate physical volume.



CAUTION

If data other than that associated with the database resides on the same physical volumes, you may inadvertently restore data that you did not intend to restore or overwrite data that you did not intend to overwrite.

Celerra restrictions

Restoring a local SnapSure Celerra iSCSI replica prevents a restore of any newer replicas of the same volume. If you have a replica of an Exchange database and its logs on Celerra iSCSI storage, and restore just the database file, Replication Manager prevents restore of the database file in any newer replica.

Restore considerations for public folder mailboxes in Exchange 2013

In an Exchange 2013 environment, you must take into account the following considerations for successful restore operations:

- ◆ If the database you are restoring hosts a primary hierarchy mailbox, restore is not allowed to continue for the database that contains primary hierarchy. Move the primary hierarchy mailbox to another database prior to restore.
- ◆ Replication Manager requires that the user for configuring the Exchange Interface Service should have an additional Exchange role: View-only Organization Management.

Restoring in an Exchange 2010/Exchange 2013 DAG environment

If the additional role is not configured, Replication Manager displays a warning that it is unable to determine the database hosting the primary hierarchy public folder mailbox. You can continue with the restore operation after you acknowledge this warning message

In a Exchange 2010/Exchange 2013 DAG environment, replicas can be restored as long as the following requirements are satisfied:

In native DAG environments:

- ◆ Restore target is the same server from which the replica was originally created.
- ◆ Restores target database is the active copy of the database or the **Activate databases before restore** checkbox is selected.

In DAGs with REE enabled (only for Exchange 2010):

- ◆ Restore target is the same server from which the replica was originally created.
- ◆ Restore target database must be the active copy. If the server is not hosting the active copy of the database, the user must manually move the database to another server, use **Move-REEActiveMailboxDatabase**. For example, to move a database and mount it use the following command:

```
Move-REEActiveMailboxDatabase -Identity <dbname>
-Mount -MailboxServer <targetserver> -verbose
```
- ◆ Replication Manager can only restore clone or snap replicas created from MirrorView primary LUNs. Remember to fracture the MirrorView session before performing the restore and synchronize the MirrorView session after the restore.

In native DAG environments, choose the **Activate databases before restore** checkbox in the Restore Options screen of the Restore Wizard, [Figure 180 on page 512](#) illustrates this screen.



Figure 180 Native DAG restore options

Restore considerations in a native DAG environment

Before activating a native DAG database, Replication Manager verifies that the passive copy of the database is healthy and the `CopyQueueLength` is 0. If the passive copy is not healthy or the queue length is greater than zero, the restore fails. In this case, manually move the active copy to the proper server using the advanced options of the following Exchange Management PowerShell cmdlet:

Move-ActiveMailboxDatabase

Options such as `MountDialOverride` and `SkipHealthChecks` can be used if necessary. Refer to Exchange 2010/Exchange 2013 documentation for further information on how to activate an unhealthy passive copy of a database.

Replication Manager suspends database replications on all passive copies during the restore to the active copy.

When Replication Manager restores just the database file from the most current replica and existing logs are preserved and replayed during recovery, you can resume replication using the following command:

Resume-MailboxDatabaseCopy

When Replication Manager restores the database and logs, the active copy becomes older than the passive copies. In this case, you must manually reseed the passive database copies using the following command:

Update-MailboxDatabaseCopy

This requires manual deletion of files on the passive copy or use of the `-DeleteExistingFiles` option to remove the logs, checkpoint, and database files at the target location.

Restore considerations in a REE DAG environment

Similarly, in a DAG with REE enabled, additional steps may be necessary when restoring databases. Before performing these steps, it is necessary to import the REE cmdlets using the following Exchange Management PowerShell command:

Import-Module reecli.base

Once these cmdlets have been imported, you can view the entire list of REE cmdlets as follows:

Get-Command -Module reecli.base

Backing up Exchange 2010/Exchange 2013 replicas

Use the **Move-REEMailboxDatabase** cmdlet to move the active database to the appropriate server before attempting a restore in an REE environment.

You should check the health of the database after the move by using the **Get-MailboxDatabaseCopyStatus** cmdlet. The state of the database should be mounted and the content index should be healthy.

Remember to fracture the MirrorView session before performing the restore and synchronize the MirrorView session after the restore.

Replication Manager can integrate with third-party backup software to create backups of Exchange replicas. The following procedure assumes that you have already created an Exchange replication job that successfully creates a replica and mounts that replica to the desired mount host.

To back up a replica:

1. In the content panel, right-click the replica that you want to back up and select **Mount**.
2. Select the replica you want to mount and click **Next**.
3. Select the checkbox next to **Mount components for host: xxxxxx**
4. Expand **General** mount options and look for the **Copy Metadata Files to** field. Select the checkbox next to that field and specify the directory path on the mount host to which you want to copy Exchange metadata files.
5. Mount and copy metadata files.
6. Mount the Exchange replica.
7. Using your third-party backup software, back up all volumes of the replica that reside on the mount host, along with all metadata files located in the directory path you specified in the **Copy Metadata files to** field.

Restoring a deleted Exchange 2010/Exchange 2013 database

There are additional steps that are required if you plan to restore a database that has been deleted from the Exchange 2010/Exchange 2013 Server. This section outlines how to complete such a restore:

1. Before the restore, use the Exchange Management Console to recreate the database(s) that you want to restore. Be sure to use the exact same database name(s) and the same path(s) to the database and logs as the database(s) stored in your replica.
2. Choose the **Recover and Mount** option during the restore to instruct Replication Manager to recover the database and mount it automatically.

Alternatively, you may choose not to recover the database automatically. If you do not use the option above, perform the following additional steps during the manual restore:

1. Copy the required logs from `_restoredLogs` directory to the directory where the current logs reside.
2. If the log file prefix changed, rename the required log files to use the new prefix.
3. Delete the `E<nn>restore.env` file.
4. Run the **ESEUTIL** command as follows for each database restored:

```
eseutil /r E<nn> /l <logpath> /s <log_path> /d <database_path> for recovery
```

5. Delete the `_restoredLogs` directory which should be empty after the database is recovered.

Note: In an Exchange 2010/Exchange 2013 DAG environment, recreate only the active database. As described above, the database name(s) and path(s) should be exactly the same as those in your replica. After the database has been restored and recovered, the user can recreate the DAG passive copies.

Importing a replica from a backup

Replication Manager includes special procedures that create a replica from an existing backup of Exchange 2010/Exchange 2013 data. This is referred to as a *replica import* because you are importing an existing backup for use as a replica.

The original, backed-up replica can be created with snap or clone technology.

To import a backup as a replica:

1. Create a temporary replica. The requirements of the temporary replica are:
 - Generated from an application set whose name exactly matches the name of the application set that created the backed up replica.
 - Generated from a job whose name exactly matches the name of the job that created the backed up replica.
 - The Exchange database names, number of devices, and paths must be the same as those on the replica that was backed up.
 - The application set that creates the temporary replica cannot be a composite that includes other data such as file systems or other non-Exchange databases.
 - Must reside on Symmetrix or CLARiiON or VNX storage.
 - Must be created by a replication technology that creates a full copy of the production volume, such as SnapView Clone or TimeFinder/Clone. Do not create the temporary replica using (COFW) technology, such as SnapView Snap, TimeFinder/Snap, and Celerra SnapSure. Note that this restriction applies to the temporary replica only; the original backed-up replica can be created with snap or clone technology.
 - In the case of a Native Exchange 2010/Exchange 2013 DAG replica, the state of all the databases (active or passive) must match the state of the databases at the time the backup was created.

Note: If the production databases you are importing no longer exist in the production environment, then it is necessary to recreate the database(s) with the same name(s) as those in the backup, then create a replica of those dummy databases to which you can import the backup.

2. Next, determine which type of restore you plan to perform. Are you recovering one database or a set of databases? Are you restoring to the point in time of the backed-up replica, or are you planning to roll forward to another point in time?
 - If you plan to restore to the point in time of the backup, delete any log files that exist in the log path of the mounted replica.
 - If you plan to roll forward to another point in time, save the log files on the mounted replica. You will need these files during the recovery phase.
3. In your third-party backup software, perform an in-place restore of all volumes of the replica from the backup to the mount host.

Note: If volumes are mounted as mount points, consider restoring files for the volume rather than the mount point, because the backup software may not be able to restore the mount point correctly.

4. In your third-party backup software, restore the metadata files to the mount host.
5. Right-click the temporary replica in the Replication Manager Console's content panel and select **Import Replica**. The **Import Replica** dialog box appears. Refer to [Figure 181](#) on page 517.

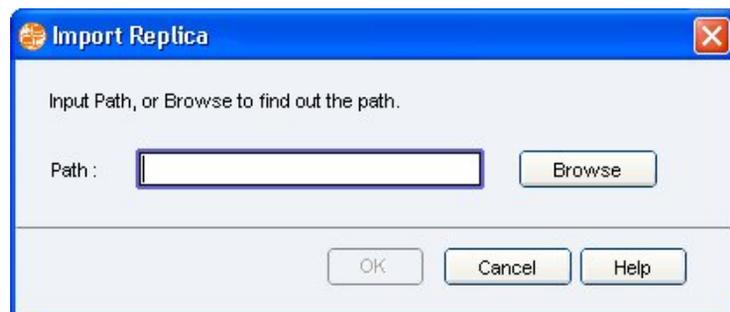


Figure 181 Import Replica

6. In the **Path** field, enter the location to which the Exchange metadata files were restored. Click **OK** to start the import process.

After Replication Manager successfully imports the replica, the timestamp of the replica changes to the date when the replica was backed up.

7. Run the following **ESEUTIL** command against the database and streaming files:

```
ESEUTIL /mh <database_name>
```

8. Examine the output of the **ESEUTIL** command to determine which log files you need to recover the Exchange databases. Look for the Logs Required line in the output, which should resemble the line in the following example:

```
Log Required: 66-68 (0x42-0x44)
```

To recover the database to the point in time of the replica, you will need log files Enn00000042.log through Enn00000044.log (Enn is the log prefix for the storage group). To roll forward, you will need the complete sequence of logs beginning at Enn00000042.log through the highest generation currently on the production machine. You may have to restore log files from more than one backup to complete the range.

Note: You should also verify the database and log signatures. Refer to Microsoft Exchange documentation for more information.

9. Right-click the replica and select **Unmount**.
10. The next step depends on your individual restore needs, as outlined in step 1:
 - If you are restoring database and log files to the point in time of the backup, right-click the replica in the content panel and select **Restore**. Select the database you want to restore and then proceed.
 - If you want to restore one database or mailbox to the point in time of the replica, but there are other databases on the LUN, you will need to use a Recovery Database. The procedure for using a Recovery Database is explained in white papers available from Microsoft and EMC.

- If you want to restore a database to a point in time beyond the backup, right-click the replica in the content panel and select **Restore**, and restore only the database file. Then, copy or restore the necessary log files to the log path used by the database(s).

For example, if you determined in [Step 8](#) that the starting log file is E0000100042.log and the ending log file is E0000100100.log, you will need to restore every log file in the sequence from E0000100042.log through E0000100100.log and if the current log generation that you want to recover is E0100140.log, you will need to recover the sequence of logs from E0000100100 through E0000100140:

- a. After you have recovered all of the log files, run the following command to verify that all of the log files are present:

```
ESEUTIL /ml <Enn>
```

where *<Enn>* is the log prefix for the database.

- b. Use the **ESEUTIL** utility to recover the database files, or use Exchange System Manager to mount and recover the database files.

Exchange 2007 and Replication Manager

Replication Manager supports Exchange 2007 environments. If you are operating in an Exchange 2010 environment, refer to [“Exchange 2010/Exchange 2013 and Replication Manager”](#) on page 495 for more information. This section covers the following topics:

- ◆ Preparing the Exchange 2007 environment
- ◆ Creating Exchange 2007 replicas
- ◆ Mounting Exchange 2007 replicas
- ◆ Restoring Exchange 2007 replicas
- ◆ Considerations when working in LCR/CCR/SCR environments

For information on troubleshooting in an Exchange environment, consult [“Troubleshooting Exchange issues”](#) on page 558.

Preparing the Exchange 2007 environment

Before you can use Replication Manager to replicate the Exchange 2007 environment, you must perform some preparatory steps:

- ◆ Configure data within storage groups
- ◆ Disable circular logging

Configuring Exchange storage groups

If you want to restore at the storage-group level, Microsoft Exchange storage groups must be arranged on physical volumes so that each storage group uses separate physical volumes for the Exchange data and the logs. If you want to restore individual databases, each database must be stored on a separate physical volume.

Note: Replication Manager checks to ensure that the Exchange data and logs are located on separate physical volumes. If not, the system issues an error and will not replicate the data.

It is also best to arrange the data so that volumes used for Exchange data do not share physical volumes with other data that is not associated with that storage group. Isolating Exchange data prevents potential problems when you restore data from a replica to the production Exchange server.



CAUTION

If data other than that associated with Exchange resides on the same physical volumes, you may inadvertently restore data that you did not intend to restore.

Disabling circular logging

Replication Manager will not replicate a storage group that has circular logging enabled.

To ensure that circular logging is disabled in Exchange 2007:

1. From the Exchange Management Console, select the server.
2. Right-click each storage group and select properties.
3. Make sure the **Enable Circular Logging** checkbox is cleared.

RecoverPoint restrictions and behaviors

Note the following restrictions and behaviors related to restoring RecoverPoint replicas of Exchange 2007:

- ◆ The **Online - Differential** option is not available for RecoverPoint jobs.
- ◆ RecoverPoint with Replication Manager in an Exchange CCR configuration is not supported.
- ◆ Partial restore of an Exchange 2007 replica can fail if RecoverPoint is not configured according to the best practice of no more than one Exchange storage group per consistency group. Be aware that, in such a configuration, restore is at the consistency group level. If you have more than one storage group in a consistency group, Replication Manager will let you select one storage group for restore, however the restore will fail because other storage group is part of that consistency group.
- ◆ RecoverPoint crash-consistent restores of Exchange 2007 are not supported. Exchange 2007 restores must be performed from an application-consistent replica.

Despite this inability to restore a crash-consistent replica, these replicas can still be used to minimize data loss.

To use log files from a crash-consistent replica:

1. Restore a storage group from an application-consistent replica without recovering the databases.
2. Use Replication Manager to mount a replica from a newer point in time.

3. Copy the newer log files to the production log volume.
4. Use `ESEUTIL /k Enn` (Enn is the log prefix for the storage group) to check the logs, then recover and mount the databases.

Creating Exchange 2007 replicas

Replication Manager has specialized agents that enable it to:

- ◆ Create online replicas of Exchange storage groups using Microsoft Volume Shadow Copy Services (VSS). The section entitled [“Support for VSS” on page 497](#) provides more information about how Replication Manager leverages VSS.
- ◆ Create online differential replicas of Exchange 2007 transaction log volumes.
- ◆ Check the consistency of replicated data.

Before each replication occurs, Replication Manager:

- ◆ Discovers the location of the data to replicate.
- ◆ Identifies pathnames for all the data in the requested storage groups.

Exchange data objects in the replica

The Exchange Administrator can configure an application set to replicate one or more storage groups. Exchange 2007 can have up to 50 storage groups or databases.

Each storage group can contain between one and five Exchange databases and the accompanying logs for those databases. With Exchange 2007, you can choose to create a differential backup and replicate just the volume(s) containing the transaction logs and system files (.chk file).

Replicating Exchange 2007 with VSS

If you are creating replicas of Exchange 2007 in Windows 2003 or Exchange 2007 SP1, Replication Manager will not use hot-split technology to quiesce the data. Instead, Replication Manager uses VSS to perform a consistent online replication. In an Exchange 2007 VSS environment, you can use the Advanced options screen of the Job Wizard to select a replication option:

- ◆ **Online Full** — Replication Manager replicates the storage group(s) (or databases), transaction logs, and checkpoint files, and then runs a consistency check to verify the consistency of the databases and logs. If the consistency check completes successfully, Replication Manager instructs Exchange to truncate the logs so that only changes that are uncommitted to the database remain.
- ◆ **Online Copy** — Replication Manager replicates the storage group(s) (or databases), transaction logs, and checkpoint files in the same way as it does during a Online Full option, however, it does not truncate the logs. Online Copy replications are often intended for testing and diagnostic purposes only.
- ◆ **Online Differential** — (Exchange 2007 only) Replication Manager replicates only the transaction logs since the last full or incremental backup. A full backup of the selected storage group must exist or the replication fails. If the system files are on another volume, it is also replicated. The transaction logs are not truncated on completion of the backup.

Exchange 2007 Replication when databases are offline

The Exchange 2007 VSS Writer does not back up databases if they are offline. If any of the databases are online, Replication Manager will run the job. It will generate a warning for the offline database. Transaction logs will not be truncated for storage groups that contain an offline database. The resulting replica is called a partial replica.

When you restore a partial replica, any databases that were offline and on their own separate volume, will not be part of the replica.

Therefore, do not select the storage group that contains these offline databases when restoring from a partial replica. If the offline database is on the same volume as an online database, it will be part of the replica. If it is part of the replica and you want to recover it, it may require additional log files.

To determine the minimum required logs, run the following command against the database after the restore:

```
ESEUTIL /mh <database name>
```

Look for the **Log Required** information in the **ESEUTIL** output.

Exchange on VMFS

Replication Manager supports VMware's use of VSS with VM snapshots when VSphere 4.1 or later is installed and VMware Tools are present on the virtual machine on the VMFS you are replicating. When Replication Manager replicates a VMFS on which Exchange resides, the VM snapshot in the VMFS replica is consistent. From an Exchange point of view, however, it is a VSS copy replica, so logs are not truncated, a consistency check is not run, and the database is not marked as backed up.

Replicating various Exchange environments

In Exchange 2007 environments, Replication Manager can replicate both active and passive copies of storage groups in an Exchange Cluster Continuous Replication (CCR) environment, or active copies in a Local Continuous Replication (LCR) environment. Replication Manager is also tolerant of Exchange Standby Continuous Replication (SCR).

Overview of Exchange SCR environments

SCR is designed for scenarios that use standby recovery servers. SCR extends the existing continuous replication features and enables new data availability scenarios for Exchange 2007 (starting with SP1) Mailbox servers. SCR uses the same log shipping and replay technology as local continuous replication (LCR) and cluster continuous replication (CCR) to provide added deployment options and configurations. Exchange SCR tolerance is available when you are replicating Exchange 2007 SP1.

Overview of Exchange LCR/CCR environments

Exchange CCR is a high availability feature of Exchange 2007 that combines asynchronous log shipping and replay features of Exchange 2007 with the failover and management features provided by the Microsoft Cluster Service (MSCS).

Exchange LCR is similar to Exchange CCR but it is a method to provide a highly available solution for the databases with a single server solution without the need to perform clustering.

The Exchange CCR cluster is a two-node Majority Node Set (MNS) failover cluster with a file share witness. The active node hosts the Clustered Mailbox Server that serves Exchange data from the active copy of the storage groups. The passive node hosts a hot standby or passive copy of the storage groups and Exchange CCR continuously ships the logs from the active to the passive copy and replays the logs into the database that resides in the passive copy of the storage group. [Figure 182 on page 526](#) illustrates a standard Exchange CCR environment.

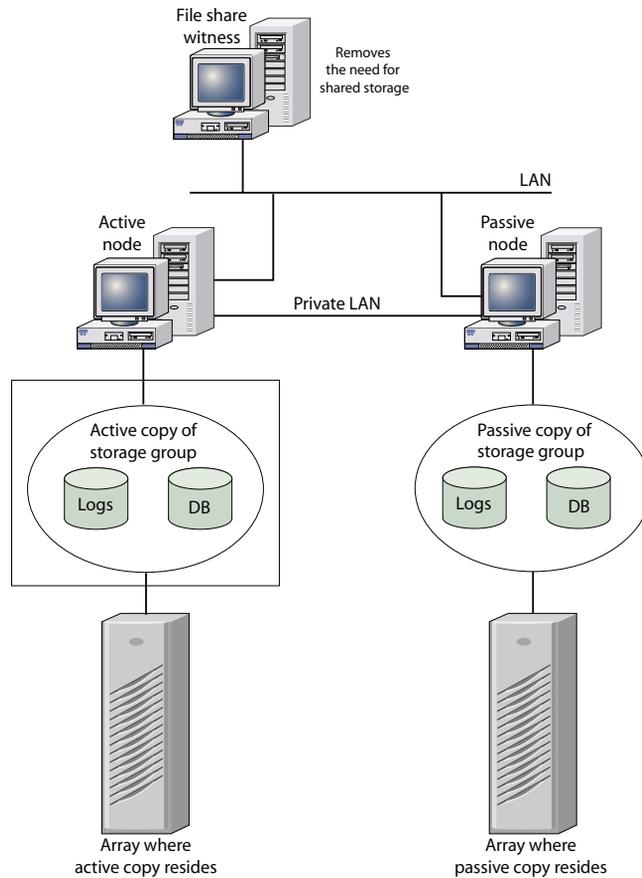


Figure 182 Exchange CCR environment

Note that the storage array where the active copy resides and the storage array where the passive copy resides do *not* have to be the same type of storage array, but EMC recommends the use of the same type of storage array for both copies for best results. For example, active copies may be stored on Symmetrix and passive copies stored on CLARiiON or VNX. This affects what replication technologies will be available to create replicas on each node of the cluster.

Exchange CCR storage group requirements

To replicate, mount, and restore Exchange data in CCR environments, there are a couple of storage group requirements as follows:

- ◆ In Exchange CCR environments, you can create only one database per storage group.
- ◆ Do not attempt to replicate passive and active copies of the same Exchange CCR storage group simultaneously as this will cause one of the jobs to fail.

Exchange CCR replication techniques

Replication Manager can create Exchange CCR replicas using either of the following techniques or a combination of the two:

- ◆ **Role-based replication** — Creates a replica of either the active copy or passive copy of an Exchange CCR storage group, regardless of which physical node currently hosts that role.
- ◆ **Node-based replication** — Creates a replica based on the particular physical node where the data is hosted, regardless of whether that node is hosting an active or passive copy of the storage group.
- ◆ **Combination replication** — Creates a replica from a specific physical node, but considers whether the node is currently hosting an active copy or a passive copy of the storage group.

Note: Exchange 2007 requires that you only restore to the active copy of the storage group. If the target node for a restore is not active, failover the cluster to make that node active before you restore. Do *not* use the cluster administration tool. Use the Move-ClusteredMailboxServer Exchange PowerShell cmdlet.

Configuring role-based replicas

Note: This scenario requires two licenses, one for each node.

Replication Manager can configure role-based replicas as follows:

1. Make sure that you have registered the following entities in the Replication Manager Console:
 - Both Exchange CCR physical nodes
 - Exchange Clustered Mailbox Server (CMS)

Figure 183 on page 528 shows a screen with all cluster nodes and the Exchange Clustered Mailbox Server registered.

Note: When adding the nodes of an Exchange CCR cluster, be sure to use the network name, not the IP address.

Name	Platform	Agent Version
lrmg047	Windows 2003	5200
lrmg055	Windows 2003	5200
lrmg190	Windows 2003	5200

Figure 183 Register physical nodes and the Exchange CMS

2. Create an application set and select the storage groups that you want to replicate from the *Exchange Clustered Mailbox Server*. Refer to Figure 184 on page 529 for more information. (Replication Manager maps both physical nodes as well so that a specific physical node can be selected as the source of the job later on.)

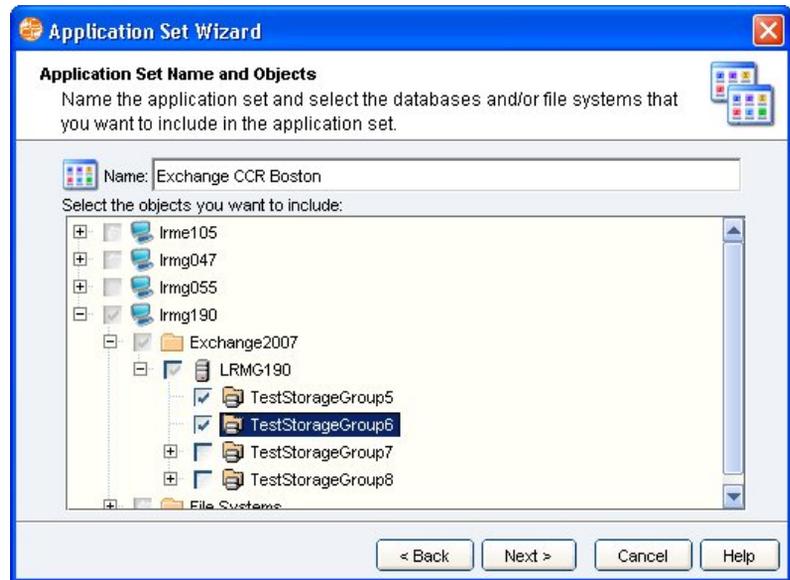


Figure 184 Create Exchange CCR Application Set from the CMS

For more information about configuring an application set, refer to the EMC Replication Manager online help.

3. Create the *first of two linked jobs* from the Exchange application set that you created above. “[Understanding link and copy jobs](#)” on [page 137](#) explains the concept of link and copy jobs.

In the first job, choose the following attributes:

- Choose one of the two *physical nodes* in the **Source** field of the **Replication Options** panel of the Job Wizard.
- Set **Replicate Copy When** in the **Job Name and Settings** panel to the storage group role to replicate (either *Active* or *Passive*). In the example in [Figure 185 on page 530](#) we are replicating the passive copy of the storage groups on the physical node lrmg047.
- Always create and schedule the job for the node that normally hosts the role you are replicating first. Then create a job for the other node. For example, if you want to replicate the passive copy, create a job for the node that normally hosts the passive copy first.

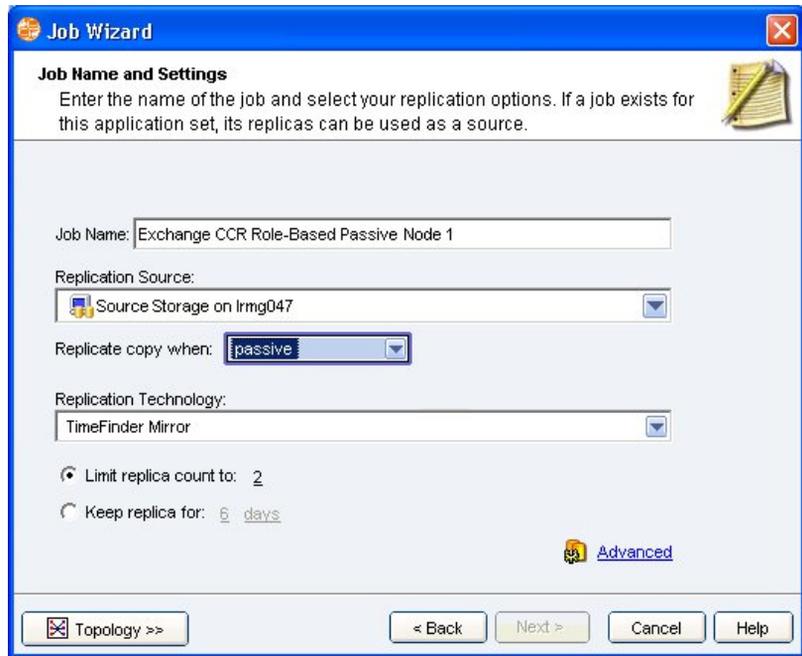


Figure 185 Role-based job replicating first physical node of Exchange CCR cluster

4. Create the second of the two linked jobs from the same Exchange application set and choose the following attributes:
 - Choose the other *physical node* in the **Source** field of the **Replication Options** panel.
 - Set **Replicate Copy When** to the storage group role to replicate (Either *Active* or *Passive*). Choose the same role you selected in step 3 above. Refer to [Figure 186 on page 531](#).

Note: Choose the same role here as was selected in the job for the other physical node. Then when these linked jobs run, only one replica will be created. The replica created will contain the copy of the storage group that represents the role (Either *Active* or *Passive*) selected.



Figure 186 Role-based job replicating second physical node of CCR cluster

- Link the jobs by choosing the first job (created above) in the **Start After Job** field of the **Starting the Job** panel. Refer to [Figure 187 on page 532](#).



Figure 187 Linking first and second job

Configuring replicas for a specific node (node-based)

Note: This scenario only requires one agent license.

Alternatively, Replication Manager can configure node-based replication by following these steps:

1. Make sure that you have registered the following hosts in the Replication Manager Console:
 - Both Exchange CCR physical nodes
 - Exchange Clustered Mailbox Server
2. Create an application set and select the storage groups that you want to replicate from the *Exchange Clustered Mailbox Server*. (Replication Manager maps all registered physical nodes as well so that a specific physical node can be selected as the source of the job later on.)

3. Create a job from the Exchange application set that you created above. In the job, choose the following attributes:
 - Choose the physical node to use in the **Source** field of the **Replication Options** panel.
 - Set **Replicate Copy When** to **active or passive**, that instructs Replication Manager to create a replica of the storage group copy on that physical host regardless of the role of the storage group that it is currently hosting.
 - Select the remaining job attributes as desired to replicate the storage groups on the physical node you chose as the source.

Configuring replicas for both nodes of the cluster regardless of role

Note: This scenario requires two licenses, one for each node.

Replication Manager can replicate both nodes of the cluster by following these steps:

1. Make sure that you have registered the following hosts in the Replication Manager Console:
 - Both Exchange CCR physical nodes
 - Exchange Clustered Mailbox Server
2. Create an application set and select the storage groups that you want to replicate from the *Exchange Clustered Mailbox Server*. (Replication Manager maps all registered physical nodes as well so that a specific physical node can be selected as the source of the job later on.)
3. Create the *first of two linked jobs* from the Exchange application set that you created above. [“Understanding link and copy jobs” on page 137](#) explains the concept of link and copy jobs.

In the first job, choose the following attributes:

- Choose one of the two *physical nodes* in the **Source** field of the **Replication Options** panel of the Job Wizard.
- Set **Replicate Copy When** in the **Job Name and Settings** panel to the storage group role to replicate (*Active or Passive*). That instructs Replication Manager to replicate the node regardless of the role.

4. Choose the Online - Copy consistency method for the first job to prevent Replication Manager from truncating the logs until the second job runs.
5. Create the second of the two linked jobs from the same Exchange application set and choose the following attributes:
 - Choose the other *physical node* in the **Source** field of the **Replication Options** panel.
 - Set **Replicate Copy When** to the storage group role (*Active or Passive*). That instructs Replication Manager to replicate the node regardless of the role.
6. Ensure that the Online - Full consistency method is selected for this second job. Online Full truncates the logs. These options are found in the **Advanced** panel of the Job Wizard. You can control the order in which the jobs run by linking them.



WARNING

Restoring replicas created by the first job requires planning. Because the second job truncates logs, the first job will always be missing logs for the time between job 1 and job 2. Those logs will be truncated and remain part of the second replica created only.

Understanding the topology of an Exchange CCR job

The job wizard panel offers a topology view that can be activated by clicking the **Topology** button. This section explains how an Exchange CCR job is represented within the Topology view.

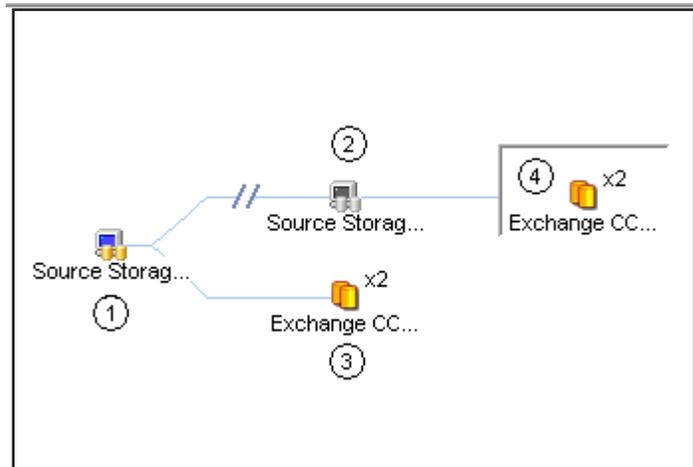


Figure 188 Topology view of a typical Exchange CCR job

The Topology view shows the entire Exchange CCR environment, including each of the following components:

- ◆ Storage for the active copy of the CCR storage groups
- ◆ Storage for the passive copy of the CCR storage groups
- ◆ Job to replicate the active copy of the CCR storage groups
- ◆ Job to replicate the passive copy of the CCR storage groups

Figure 188 on page 535 shows the topology view of an Exchange CCR environment that has jobs configured to replicate both the active and passive copies of the Exchange CCR storage groups. The numeric labels in the illustration correspond to the numbered components listed above.

Running a consistency check

Exchange 2007 jobs use the Consistency Check API to check the database and log files. A single mount host with the Exchange 2010 Management Tools can be used to run the consistency check for both Exchange 2010 and Exchange 2007.

If you select an Online Full replica that truncates transaction logs, Microsoft requires that you mount the replica as part of the job, and run a consistency check. If the mount fails for any reason or device errors occur, the logs are not truncated and the replica fails.

To mount the Exchange database for consistency checking only, you need to install the appropriate tools on the mount host, depending upon your version of Exchange. See the list below for details:

- ◆ For Exchange 2007 install the Microsoft Exchange Management Tools and reboot the system *before* installing the Replication Manager Exchange Agent. If you do not reboot the system, the COM+ component may not install correctly.

Note: If the job truncates logs, you should mount the replica and perform a consistency check. Consistency checks are optional but Microsoft requires that you mount the replica and perform a consistency check for every Exchange replica you create to prevent data loss if you ever need to restore a replica later on. If the consistency check fails on any Exchange database on the mount host, then the transaction logs are not truncated.

Consistency check advanced features

Replication Manager offers some advanced features that change how consistency checks are executed. Enabling these features can impact performance. For most users the default settings are sufficient. These advanced features include:

- ◆ **Single mount host consistency checks** — A single mount host with the Exchange 2010/Exchange 2013 Management Tools can be used to run the consistency check for both Exchange 2010/Exchange 2013 and Exchange 2007.
- ◆ **Minimize log checking** — Choosing this option from the mount options panel of the Job Wizard speeds up the log checking by instructing the consistency checking software to check only those logs that are required to recover the database. Selecting this option improves the performance of the consistency check.

If you clear that checkbox, then consistency check will be performed on all logs in the Exchange storage group when a mount occurs.

This command instructs Replication Manager to only check a subset of the Exchange logs that are included in the replica. If your backup window is small, you may find this option useful. However, the replica contains logs that have not been checked for consistency. If you attempt to restore the log volume, you may find that some log files are corrupt or the log sequence is not complete. Before restoring the log volume, you should mount the

replica and run **eseutil /k Enn** against the log path. For maximum protection, clear **Minimize log checking**. For maximum performance, select it.

You must also set the Working directory for Exchange 2007 replicas. The working directory is the directory to which the required log files will be copied in order to check them.

The **Minimize log checking** option is not available when the consistency method is Online - Differential.

- ◆ **Parallel consistency checks** — Consistency checks can run against the databases in parallel (all databases for a storage group at the same time) or sequentially. If the databases all reside on the same LUN, use the sequential option.
- ◆ **I/O throttling during consistency checks** — In some versions of Exchange, consistency checks can be paused to throttle the consistency checking operation. Specific functionality of consistency check throttling depends upon the version of Exchange as follows:
 - **For Exchange 2007:** Throttling options allow you to specify the number of I/Os after which to pause and the duration of the pause.

Managing Exchange errors

Exchange logs the following errors in the Application event log when they occur:

- ◆ **-1018** — The database tried and failed to verify information about a particular page in the database.
- ◆ **-1019** — This is similar to a -1018 error but indicates that the accessed page has returned an invalid page number (usually all zeros) rather than an invalid checksum.
- ◆ **-1022** — This is indicative of major hardware problems, particularly disk subsystem problems. If the database engine requests a page from disk but instead receives an error from the I/O subsystem, a -1022 error results.
- ◆ **447** — This is indicative of corruption in the logical database structure. This accompanies a message stating that the information store terminated abnormally.
- ◆ **448** — This error denotes an inconsistency or corruption in a table in the jet database. This accompanies a message stating that an information store data inconsistency has been detected in a table.

Replication Manager searches the application event log for these errors every time a replica is created. The first time it runs, Replication Manager searches the entire log. Subsequent runs search since the last successful run. If there are no existing replicas, then Replication Manager will search the entire log when creating the next replica.

You can configure Replication Manager to ignore any of these errors, all of these errors, or ignore them for one run of the job. These errors can cause Replication Manager to fail a replication task unless you specifically instruct Replication Manager to ignore them. For more information about ignoring Exchange errors, refer to the EMC Replication Manager online help.

Mounting Exchange 2007 replicas

Replication Manager can:

- ◆ Mount a replica to an alternate host using the same file path as on the production host.
- ◆ Mount a replica on an alternate host in a new location (determined by adding an alternate path to the beginning of path).
- ◆ Mount a replica on an alternate host in a new location (determined by path mapping). Refer to the following note.
- ◆ Mount a replica on the production host in a new location (determined by adding an alternate path to the beginning of the path).
- ◆ Mount a replica to the production host in a new location determined by path mapping. Refer to the following note.

Note: Specific information about how alternate paths and path mapping work can be found at [“Mounting using alternate path” on page 174](#) or [“Mounting using path mapping” on page 177](#).

When you mount an Exchange 2007 replica, there are certain considerations. Replication Manager can mount a replica to an alternate host that may or may not have Exchange Server installed.

It can perform:

- ◆ Mounts of new replicas as part of the replication job
- ◆ On-demand mounts of existing replicas

Note: Although Replication Manager can mount a replica to an alternate location on the production or mount host, alternate location mounts have a limited use. These mounts can only be used to check the database and logs for consistency. The databases cannot be brought online in Exchange System Manager.

Restoring Exchange 2007 replicas

When you restore a Microsoft Exchange full replica, you can choose to restore any of the following:

- ◆ All storage groups in the application set
- ◆ One or more storage groups from the application set
- ◆ One or more databases from the application set

For Exchange 2007, from a full replica, you can also restore the logs of one or more storage groups from the application set.

When restoring an Exchange 2007 storage group, first confirm that no databases have been added to the storage group after the replica was taken. If that occurs, Replication Manager tries to mount a database that no longer exists after the restore and that causes a failure.

For Exchange 2007, partial restore of a RecoverPoint replica is supported, if the RecoverPoint consistency groups were configured to contain no more than one Exchange storage group each. Restore is at the consistency group level.

When you are restoring just a database file, verify that the transaction log files needed for recovery are present. An unbroken sequence is required. To determine the minimum required range of logs, run the following command against each database after the restore:

```
ESEUTIL /mh <database name>
```

Look for the **Log Required** information in the **ESEUTIL** output.

A Microsoft Exchange replica includes a checkpoint file (.chk) that records the location in the transaction log files of the last complete transaction that Exchange wrote to the database. If you choose to perform a full restore of databases and logs, the checkpoint file helps Exchange know where to start if you choose to roll the database forward. If there is no .chk file, Exchange starts replay of the logs with the oldest transaction log file.

Restoring a differential backup in Exchange 2007

A restore from a differential backup restores only the logs. To restore from a differential backup, you must have the last full backup and the most recent differential backups available.

While restoring from multiple restore sets (say one FULL and one Differential), we need a way to instruct Exchange Server 2007 VSS Writer to replay logs only after the last restore set. Thus, a storage group or database restore needs to occur first with the option to recover and mount the databases deselected. It should then be followed with a Differential restore with the option to recover and mount the databases selected.

Note: Because restores are done at the physical disk level, all other data on the disks is overwritten.

Restoring with VSS

On Windows 2003, 2008, and 2012, Replication Manager coordinates with VSS and Exchange 2007 to restore the volumes that contain the databases and logs that you have selected to recover. If you choose to restore just one database from a storage group, Replication Manager takes all databases in the storage group offline because no temporary restore directory is available in which JET can play forward the logs. When restoring a VSS replication, you must restore it to the same location on the server from where it originated. Also, you cannot restore a VSS replica to a Recovery Storage group. The Exchange 2007 VSS Writer is always involved in the restore and recovery of the Exchange databases.

Therefore, Replication Manager selects the **Recover and mount databases** checkbox by default for Exchange 2007.

Restoring with Recover and Mount Database option selected

When the **Recover and mount databases** checkbox is selected for an Exchange 2007 restore, the Exchange 2007 VSS Writer performs the following tasks:

- ◆ Verifies that all of the required logs are available
- ◆ Verifies that there are no gaps in the log sequence
- ◆ Deletes the checkpoint file (as needed)
- ◆ Recovers the databases

Restoring without Recover and Mount Database option

If the **Recover and mount databases** checkbox is not selected, the Exchange 2007 VSS Writer will not delete the checkpoint file. You will be responsible for the following tasks:

- ◆ Deleting the checkpoint file (Enn.chk) although, the Exchange VSS Writer will delete the checkpoint file when appropriate.
- ◆ Recovering the databases.

To recover Exchange databases after this type of restore:

1. Use Replication Manager to restore the replica.
2. Run the **ESEUTIL** command as follows for each database restored:

```
ESEUTIL /mh <edbfile>
```

3. Check the **Logs Required** information. Make sure those logs are in the log path.
4. Run the **ESEUTIL** command as follows for each log of the storage group:

```
ESEUTIL /ml <Enn>
```

where *<Enn>* is the log prefix of the storage group. This will make sure the log sequence has no gaps.

5. For Exchange 2007 restores, determine the type of restore you did. If you restored the entire storage group, you do not need to delete the checkpoint file. If you restored a database from the most current replica, you will need to delete the checkpoint file.
6. To recover Exchange 2007 databases, use the Exchange Management Console to mount all of the databases in the storage group.

Selecting full or partial restores

The first panel of the Restore Wizard allows you to choose all or part of a replica to restore. This section describes what will be restored if you make certain selections on that panel.

To restore the entire contents of the application set, select the top node of the Replica tree as shown in [Figure 189 on page 543](#).

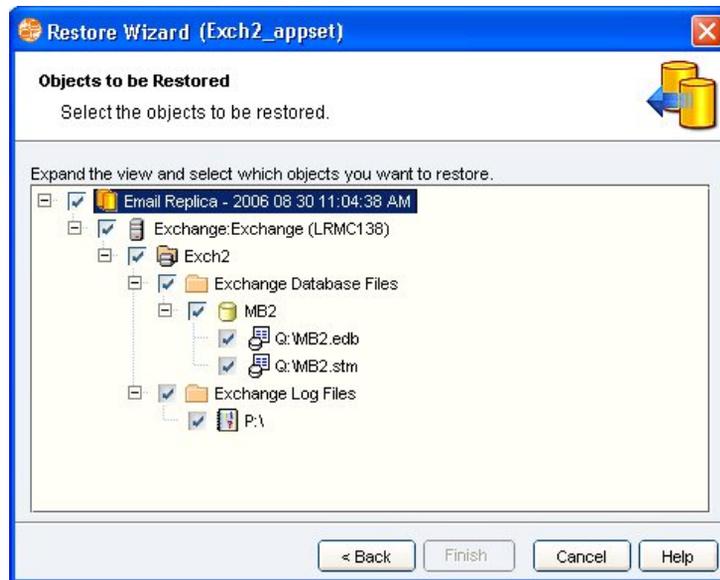


Figure 189 Selecting application host objects (example shows Exchange 2003)

To restore the entire contents of a storage group, select the top node of that storage group as shown in [Figure 190](#) on page 544.

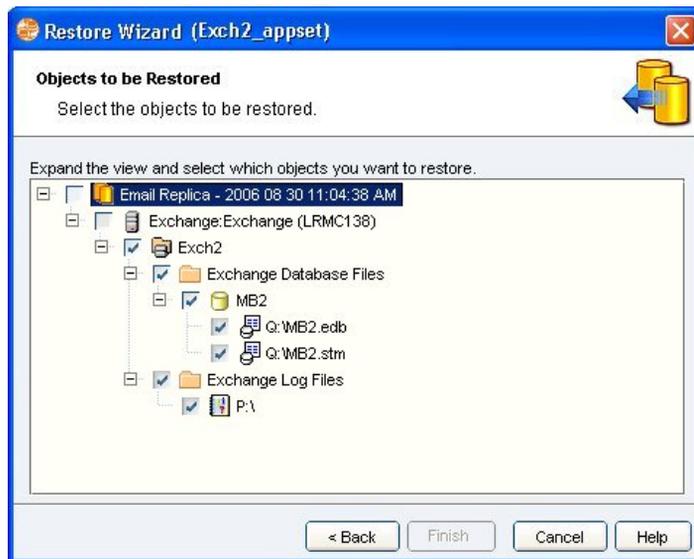


Figure 190 Objects to be restored (example shows Exchange 2003)

Note: If you attempt to restore a mailbox store (in any supported version of Exchange) that has been entirely deleted from the Exchange system since the replica was created, you must first re-create the specific mailbox store in Exchange before Replication Manager can restore the data from that mailbox store.

Restoring datafiles in Exchange 2007

To restore datafiles for one or more databases in an Exchange 2007 environment, select the individual databases.

Restoring log files in Exchange 2007

To restore Exchange 2007 log files, select them on the **Objects to be Restored** panel of the Restore Wizard.

Before restoring log files, restore just the datafiles from an older full replica, with the **Recover and Mount Databases** option cleared. Then restore just the log files from a newer full or differential replica with the **Recover and Mount Databases** option selected. This is important because while restoring from multiple restore sets, Replication Manager needs to instruct Exchange Server 2007 VSS Writer to replay logs only after the last restore set. Thus, a storage group or database

restore needs to have occurred first with the option **Recover and Mount Databases** cleared.

If you choose to restore logs, any newer logs will be overwritten. You might want to preserve the newer logs by copying them to other volumes before the restore.

Restore considerations

If the logs are selected for restore (from a full or differential replica), they will overwrite newer logs created since the replica was created. That means your database will represent the point in time when the replica was created.

Note: If you want to preserve logs as they exist on the production host, you should choose to restore only the datafiles to prevent Replication Manager from restoring older logs over newer logs.

If you want to restore at the storage-group level, Microsoft Exchange storage groups must be arranged on storage volumes so that each storage group uses separate physical volumes for the Exchange data and the logs. If you want to restore individual databases, each database must be stored on a separate physical volume.



CAUTION

If data other than that associated with the storage group resides on the same physical volumes, you may inadvertently restore data that you did not intend to restore or overwrite data that you did not intend to overwrite.

Celerra restrictions

Restoring a Celerra iSCSI replica is destructive to any newer replicas of the same volume. If you have a replica of a database and its logs and restore just the database file, Replication Manager will disable restore of the database file in any newer replicas. Also, this action prevents subsequent mounts of newer replicas.

When restoring logs from a differential replica, any newer differential replicas will be marked unrestoreable. In addition, the logs in any newer full or copy replicas will be marked unrestoreable.

Mounting in an Exchange 2007 CCR environment

Exchange CCR replicas of the active copy of the storage group and the passive copy of the storage group can be mounted to an alternate mount host for the purpose of running an Exchange consistency check only. Replicas created on the active node of the cluster and the passive node of the cluster can have different mount hosts. The

Restoring in an Exchange 2007 CCR environment

mount host used must have visibility to the array where the replica you are mounting was originally created.

In an Exchange CCR environment, replicas can be restored only to the same physical node on which the replica was originally created. In order to restore an Exchange CCR replica, the node must be hosting the active copy of the Exchange storage groups at the time of the restore. If the node you want to restore is not hosting the active copy of the storage groups, move the active copy to the other physical node by using the following PowerShell command:

Move-ClusteredMailboxServer

When Replication Manager performs a restore to an Exchange CCR environment, it first suspends replication of the storage groups that it is restoring; However, if Replication Manager restores just the database from the most current replica and existing logs are used to roll forward to the current point in time, you can resume replication using the following command:

Resume-StorageGroupCopy

When Replication Manager restores a storage group (database and logs), the active copy becomes older than the passive copy. When this situation occurs it is necessary to reseed the passive copy using the following command:

Update-StorageGroupCopy

This requires manual deletion of files on the passive copy or use of the **-DeleteExistingFiles** option to remove the logs, checkpoint, and database files at the target location.

Restoring in an Exchange 2007 SCR environment

If you are restoring to a storage group that is protected by Exchange Standby Cluster Replication (SCR), Replication Manager automatically suspends the SCR replications on each of the SCR targets prior to the restore operation. You can resume SCR protection after the restore using the following command:

```
Resume-StorageGroupCopy
```

Backing up Exchange 2007 replicas

Replication Manager can integrate with third-party backup software to create backups of Exchange replicas. The following procedure assumes that you have already created an Exchange replication job that successfully creates a replica and mounts that replica to the desired mount host.



CAUTION

In order to create a replica that you can import later from the backup, the database must be online when you perform the backup. Offline backups do not generate the necessary metadata needed to import the backup later.

To back up a replica:

1. In the content panel, right-click the replica that you want to back up and select **Mount**.
2. Select the replica you want to mount and click **Next**.
3. Select the checkbox next to **Mount components for host: xxxxxx**
4. Expand **General** mount options and look for the **Copy Metadata Files to** field. Select the checkbox next to that field and specify the directory path on the mount host to which you want to copy Exchange metadata files. Refer to [Figure 191 on page 548](#).

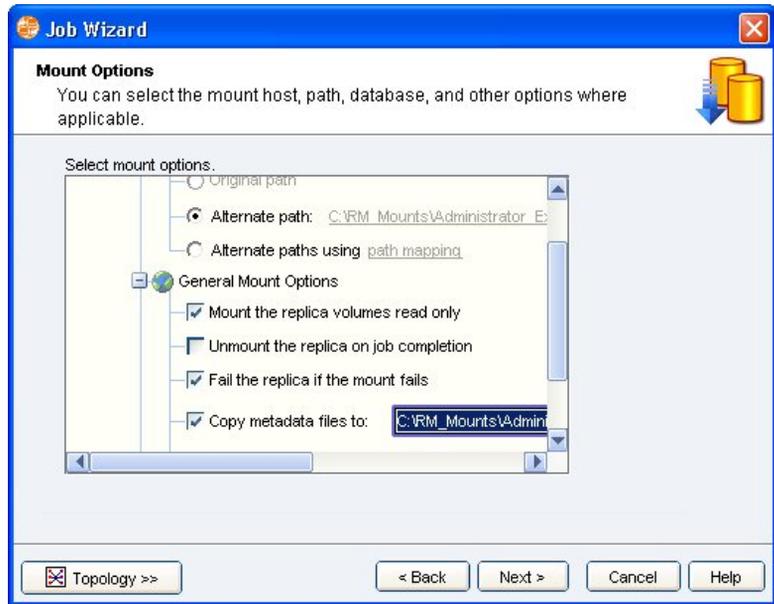


Figure 191 Mount and copy metadata files (Example shows Exchange 2003)

5. Mount and copy metadata files.
6. Mount the Exchange replica.
7. Using your third-party backup software, back up all volumes of the replica that reside on the mount host, along with all metadata files located in the directory path you specified in the **Copy Metadata files to** field.

Importing a backup as a replica

Replication Manager includes special procedures that create a replica from an existing backup of Exchange 2007 data. This is referred to as a *replica import* because you are importing an existing backup for use as a replica.

Note: Replicas created using copy on first write (COFW) technology, such as SnapView Snap, TimeFinder/Snap, and Celerra SnapSure, are not eligible for use during replica import from a backup. Use a replication technology that creates a full copy of the production volume, such as SnapView Clone or TimeFinder/Clone.

To import a backup as a replica:

1. Run a predefined job to create an extra, temporary Exchange replica. Replication Manager requires a temporary replica as a placeholder to successfully import the older data.

The temporary replica must have the same Exchange database names, number of devices, and paths as the replica that you want to restore from backup. In addition, the application set cannot be a composite that includes other data such as file systems or other non-Exchange databases.

Note: If the production databases you are importing no longer exist in the production environment, then it is necessary to recreate the database(s) with the same name(s) as those in the backup, then create a replica of those dummy databases to which you can import.

The temporary replica must reside on Symmetrix or CLARiiON or VNX storage.

Exchange must be running on Windows Server 2003, 2008, or 2012.

2. Next, determine which type of restore you plan to perform. Are you recovering just a database, or the entire storage group? Are you restoring to the point in time of the backed-up replica, or are you planning to roll forward to another point in time?
 - If you plan to restore to the point in time of the backup, delete any log files that exist in the log path of the mounted replica.
 - If you plan to roll forward to another point in time, save the log files on the mounted replica. You will need these files during the recovery phase.
3. In your third-party backup software, perform an in-place restore of all volumes of the replica to the mount host.

Note: If volumes are mounted as mount points, consider restoring files for the volume rather than the mount point, because the backup software may not be able to restore the mount point correctly.

4. In your third-party backup software, restore the metadata files to the mount host.
5. Right-click the temporary Exchange replica in the content panel and select **Import Replica**. The **Import Replica** dialog box appears. Refer to [Figure 192 on page 550](#).

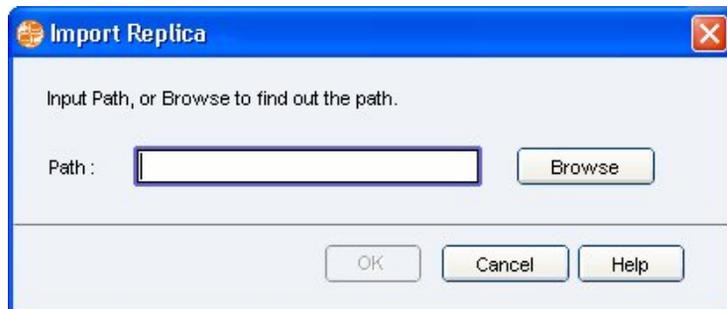


Figure 192 Import Replica

6. In the **Path** field, enter the path to the Exchange metadata files that you had specified when you mounted the replica and initiated the backup. Click **OK** to start the import process.

After Replication Manager successfully imports the replica, the timestamp of the replica changes to the date when the replica was backed up.

7. Run the following **ESEUTIL** command against the database and streaming files:

```
ESEUTIL /mh <database_name>
```

8. Examine the output of the **ESEUTIL** command to determine which log files you need to recover the Exchange databases. Look for the Logs Required line in the output, which should resemble the line in the following example:

```
Log Required: 66-68 (0x42-0x44)
```

To recover the database to the point in time of the replica, you will need log files Enn00042.log through Enn00044.log (Enn is the log prefix for the storage group). To roll forward, you will need the complete sequence of logs beginning at Enn00042.log through the highest generation currently on the production machine. You may have to restore log files from more than one backup to complete the range.

Note: You should also verify the database and log signatures. Refer to Microsoft Exchange documentation for more information.

9. Right-click the replica and select **Unmount**.

10. The next step depends on your individual restore needs, as outlined in step 1:
- If you are restoring the entire storage group (database and log files) to the point in time of the backup, right-click the replica in the content panel and select **Restore**.
 - If you want to restore one database to the point in time of the replica, but there are other databases on the LUN, you will need to use a recovery storage group. The procedure for using a recovery storage group is explained in white papers available from Microsoft and EMC.
 - If you want to restore a storage group or a database to a point in time beyond the backup, right-click the replica in the content panel and select **Restore**, and restore only the desired components. Then, copy or restore the necessary log files to the log path used by the storage group.

For example, if you determined in [Step 8](#) that the starting log file is E0100042.log and the ending log file is E0100100.log, you will need to restore every log file in the sequence from E0100042.log through E0100100.log and if the current log generation that you want to recover is E0100140.log, you will need to recover the sequence of logs from E0100100 through E0100140:

- a. After you have recovered all of the log files, run the following command to verify that all of the log files are present:

```
ESEUTIL /ml <Enn>
```

where *<Enn>* is the log prefix for the storage group.

- b. Use the **ESEUTIL** utility to recover the database files, or use Exchange System Manager to mount and recover the database files.

Exchange mailbox recovery procedures

If you want to perform a mailbox recovery from an existing replica:

1. Try to recover the mailbox using the built-in Undelete feature offered by Exchange. The default save time for undelete is 30 days, but that can be increased to 90 days by the administrator. If you are unable to retrieve the necessary Exchange information using that method, proceed to step 2.

2. For Exchange 2007, use the Recovery Storage Group to recover the mailbox.

Item level restore

Replication Manager can restore individual Exchange 2007 and Exchange 2010 mailboxes and messages when EMC® ItemPoint™ for Microsoft® Exchange™ Server is installed.

The following software must be installed and running on the mount host where you want to perform item level restores:

- ◆ Replication Manager Agent
- ◆ EMC ItemPoint for Microsoft Exchange Server
- ◆ Replication Manager Console

Note: You must be using the Replication Manager Console on the mount host to run item level restore. Before using item level restore, it is recommended that you disable expiration for the replica.

To restore one or more mailboxes or messages:

1. Right-click a specific Application Set or replica and choose **Restore** and then **Item Level**.
2. Select the replica from which you want to restore individual Exchange items.

If you select a RecoverPoint replica, you can choose to perform item level restores from a point in time or from an existing replica.
3. Select a database. ItemPoint can only launch one database at a time.
4. If you selected a replica that is not mounted, the Mount Path panel is displayed. Specify the mount path where the replica will be mounted.

The replica will be mounted in read only mode. It is best to select a path that the job does not use. This way, if the job runs, the replica will not be unmounted.

5. Select the location that EMC ItemPoint for Microsoft Exchange Server will use to recover the databases so it can access the mail items.

EMC ItemPoint for Microsoft Exchange Server is launched.

6. Use ItemPoint to restore mailbox items.

For more information about using EMC ItemPoint for Microsoft Exchange Server, refer to the EMC ItemPoint for Microsoft Exchange Server documentation.

Note: The ItemPoint documentation may describe operating systems and Exchange versions not supported by Replication Manager. The *EMC Replication Manager Support Matrix* is the authoritative source of information about supported operating systems and Exchange versions. To access the *EMC Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

7. Enable replica expiration (if expiration was disabled before the restore).

After performing item level restore

After you perform an item level restore, to avoid errors, you should close the ItemPoint console before using unmounting replicas in the Replication Manager GUI. Otherwise, if ItemPoint remains running on your system, you must make sure the Exchange store is closed before unmounting any replicas in the Replication Manager GUI.

To close an Exchange store in ItemPoint, in the source pane, select a source EDB, PST, or CAS file, and do one of the following:

- ◆ From the **File** menu, click **Close Store**.
- ◆ Right-click the source that you have just restored, and on the context menu click **Close Store**.

Using pre- and post-replication Exchange scripts

The Replication Manager performs certain default actions on the production Exchange server before and after splitting the mirror. You can add customized actions with your own user-supplied scripts.

To use a pre-replication and post-replication script:

1. Name the script and its location (by specifying the full path and filename) while configuring the job in the console.
2. Ensure that the user account associated with the application set can execute all pre- and post-replication scripts that you plan to run on replicas created by that application set.

A script must be in one of the following executable formats (.bat or .exe) or in a Power Shell cmdlet in Exchange 2007. Scripts are not required. However, if you need to perform some preparatory steps on the machine before you create the replica or if you want to perform some clean up afterward, pre- and post-replication scripts can help you to do that.

Suppress output from the script. The following techniques work for selected popular environments:

- ◆ **Windows** — Add @echo off to the first line of the .bat script.
- ◆ **UNIX sh** — Redirect output from the script as shown:

```
Prog.sh > /dev/null 2>&1
```

- ◆ **UNIX csh** — Redirect output from the script as shown:

```
Prog.csh >& /dev/null
```

Note: Suppressing script output prevents invalid characters from being added to replica history. Invalid characters in the history cause Replication Manager to stop responding when mounting or viewing a replica. If you need output from the script, set up and use a log file instead of directing output to standard out.

Considerations for Exchange in a cluster

Replication Manager can be used to replicate and restore Exchange data that resides on a Microsoft Cluster Service (MSCS) failover cluster. This section describes special procedures that you must follow in order to ensure success when manipulating Exchange data in a clustered environment.

Additional Exchange 2007 cluster considerations

Exchange 2007 supports two cluster configurations as follows:

- ◆ Single Copy Cluster (SCC)
- ◆ Cluster Continuous Replication (CCR)

Replication Manager supports both of these cluster technologies. The following considerations apply:

- ◆ If you run a job to replicate the passive copy of the storage group using the **Online -- Full** replication option, Exchange truncates the transaction logs for the active and passive copies of the storage group. Do not create separate jobs to truncate logs on both nodes.
- ◆ Extended Maintenance Mode QFEs are not required in Exchange CCR environments.
- ◆ The following are not supported in an Exchange CCR configuration:
 - RecoverPoint with Replication Manager
 - CLARiiON/VNX secondary devices in a MirrorView configuration
 - Symmetrix R2 devices in a SRDF configuration

Restoring Exchange replicas to a MSCS cluster

Replication Manager can restore Exchange replicas to an MSCS failover cluster, however, certain special procedures must be followed to prevent the cluster from failing over during the restore procedure.

To enable instant restore capabilities

In the following environment:

- ◆ Windows Server 2003
- ◆ CLARiiON storage arrays
- ◆ VNX storage arrays
- ◆ Celerra storage arrays

- ◆ VNXe storage arrays

Make sure that you set the protected restore flag on the CLARiiON, VNX, or VNXe array in order to get instant restore capability. Instant restore should be enabled on Symmetrix arrays.

Restoring Exchange 2007 data in an MSCS environment

To restore Exchange 2007 data in an SCC/MSCS environment that does not have the Extended Maintenance Mode hotfixes:

1. Remove any dependencies that exist on the physical disk resource(s) to which you intend to restore the data.
2. Delete the physical disk resources from the Exchange resource group.
3. Restore the selected Exchange storage group(s) or database(s) using the Replication Manager Console.

Note: Clear the **Recover and mount databases** checkbox. Refer to [“Restoring without Recover and Mount Database option”](#) on page 542 for more information on how to recover the databases manually.

4. Add the Physical Disk Resources back into the Exchange Resource Group using the Cluster Administrator.
5. Re-add any dependencies that you removed on the physical disk resource(s) to which you restored the Exchange data.

Troubleshooting Exchange issues

This section covers some troubleshooting issues that you might encounter with Exchange and its interaction with Replication Manager.

Considerations in DAG environments with REE enabled

In DAG environments with REE enabled, The MirrorView plugin should be enabled and the RecoverPoint plugin should be disabled on all servers in the DAG.

Use **Get-REEPluginInfo** to determine the status of the MirrorView and REERecoverPointPlugin. The REERecoverPointPlugin must be disabled. If the plugin is enabled, use the following command to disable it:

```
Disable-REEPlugin -Identity REERecoverPointPlugin
```

Exchange Interface Service for Exchange 2007 and 2010

The Exchange 2007 Management Tools and .NET Framework 2.0 must be installed before installing the Replication Manager Exchange 2007 agent. For Exchange 2010, the Exchange Management Tools and its prereqs need to be installed.

Note: You must reboot the system after the Exchange 2007 Management Tools are installed to ensure that the Replication Manager Exchange Interface component will register properly.

The Replication Manager Exchange Interface Service should be automatically registered when Replication Manager is installed. To manually install the service and set the user account:

1. Navigate to the following location, assuming Replication Manager is installed in the default location:

```
C:\Program Files\EMC\rm\client\bin
```

2. Run the following command to install the component:

```
RM_ExchangeInterface.exe /service /user <domain name\username> /password <password>
```

Example:

```
RM_ExchangeInterface.exe /service /user "test103\administrator" /password pwd123
```

Error messages are written to the Install.log file in the same directory. Additional parameters you may use are:

`/debug` — To write debugging information to the Install.log file.

`/nopriv` — Use when having problems setting the user account. The steps below will also set the user account for the service.

3. Run **DCOMCNFG**.
4. Expand **Component Services > Computers > My Computer > DCOM Config**.
5. Right-click **Replication Manager Exchange Interface** and select **Properties**.
6. Click the **Identity** tab.
7. Add the domain user account and password and click **OK**. See the sections on security for Exchange 2007/2010 earlier in this chapter.

To manually remove the service and DCOM Config:

1. Navigate to the following location, assuming Replication Manager is installed in the default location:

```
C:\Program Files\EMC\rm\client\bin
```

2. Run the following command to install the component:

```
RM_ExchangeInterface.exe /unregserver
```

This will stop and remove the service and COM configuration settings.

Modifying the Exchange Interface Service user account or password

If your user account and/or password information changes after the Replication Manager Exchange Interface Service is registered, you can use the following procedure to change the user account and/or password information:

1. Start component services.
2. Expand **Computers > My Computer > DCOM Configuration**.
3. Right-click **Replication Manager Exchange Interface** and select **Properties** from the context menu.
4. Select the **Identity** tab.
5. Modify the domain user account and/or password and click **OK**.

Activating diagnostic logging for the Exchange Replication Writer

This will also modify the account information for the Replication Manager Exchange Interface Service.

Activate the Exchange replication writer by following these steps:

1. Open the Exchange Management Shell.
2. Enter the following command:

```
Set-EventLogLevel "MSExchange Repl\Exchange VSS  
Writer" -Level Expert
```

This causes more diagnostic messages to be logged to the Application Event log.

Deactivating diagnostic logging for the Exchange Replication Writer

If you want to discontinue extra diagnostic information that you activated using the preceding procedure:

1. Open the Exchange Management Shell.
2. Enter the following command:

```
Set-EventLogLevel "MSExchange Repl\Exchange VSS  
Writer" -Level Lowest
```

This sets the minimum level of messages.

Activating diagnostic logging for Exchange 2007/2010 VSS Writer

If you need extra diagnostic information to troubleshoot issues with the VSS Writer, you can turn on diagnostic logging by following these steps:

1. Open the Exchange Management Shell.
2. Enter the following command:

```
set-eventloglevel "msexchangeis\9002 System\Exchange  
Writer" -level Expert
```

This causes more diagnostic messages to be logged to the Application Event log.

Deactivating diagnostic logging for Exchange 2007/2010 VSS Writer

If you want to discontinue extra diagnostic information that you activated using the preceding procedure:

1. Open the Exchange Management Shell.
2. Enter the following command:

```
set-eventloglevel "msexchangeis\9002 System\Exchange Writer" -level Lowest
```

This sets the minimum level of messages to be logged to the Application Event log.

Logging for the Replication Manager Exchange Interface Service

The Replication Manager Exchange Interface Service logs messages to the application event log. The source IDs are RM_ExchangeInterface and ECCRMAgent. Diagnostic logging is enabled by selecting the Replication Manager Debug logging options for the Replication Manager host.

Resolving failures when you restore databases without transaction logs

If a restore fails when you restore one or more databases without restoring the transaction logs, subsequent restores may fail with the following error:

```
027126 ERROR: Exchange Information Store has failed at pre restore. The error is VSS_E_WRITERERROR_RETRYABLE. The code is: 0x800423f3. Check the application event log for more information.
```

If this error occurs, delete any *Enn*restore.env files (where *nn* is the logfile prefix) in the transaction log directories and try your restore operation again.

Exchange restore fails with VDS errors due to “devices in use” error

This problem occurs when the Replication Manager agent (ircdd) crashes during restore.

This issue occurs when a user is attempting to restore a large number of storage groups in one replication. This causes an error because Replication Manager typically waits for 20 minutes for VSS to complete its post-processing operations after the restore. When there is a large number of storage groups, VSS processing may require longer to complete.

There are two possible resolutions to this problem:

Changing the VSS asynchronous processing timeout

- ◆ Restore fewer storage groups at one time by changing the restore configuration.
- ◆ Set a registry value to override the default VSS timeout as described below.

To change the VSS asynchronous processing timeout:

1. Create a DWORD value called `CC_VSS_ASYNC_TIME_OUT` in `HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\EMC\EMC ControlCenter Replication Manager\Client\<RM_version>`
2. Set a decimal value representing what the wait time should be in seconds. For example, set the value to 2700 for 45 minutes.

Storage group restore error: “volume cannot be locked for exclusive use”

This problem can occur in Exchange 2007 CCR clusters where the Microsoft Exchange Replication Service and a SVCHost service hold the log volume in use. Verify that there are no other applications active on the volume, such as a command window, Exchange Management Console, Windows Explorer, or an application instance that is not involved with the restore.

To solve this issue:



WARNING

Incorrectly modifying the Registry may cause serious system-wide problems that may require you to reinstall your system. Modify the Registry at your own risk.

1. Navigate to `HKEY_LOCAL_MACHINESYSTEM/CurrentControlSet/Services/lanmanworkstation/parameters`.
2. Create a new DWORD value called `KeepConn`.
3. Set the value to 1 (this is the number of seconds to keep the connection open). The Windows default value is 600 (10 minutes) if no value is present.
4. Restart the Workstation service.

Exchange replication error: "waiting 1200 seconds for BackupComplete to complete"

Running a replication while Exchange 2007 maintenance activities are taking place may result in the error "waiting 1200 seconds for BackupComplete to complete. BackupComplete was cancelled."

To prevent this failure, do not schedule replications during the Exchange maintenance window.

SharePoint Procedures

This appendix covers the specifics of support for SharePoint, including the following sections:

- ◆ Overview of support for SharePoint 566
- ◆ Configuring the SharePoint environment 567
- ◆ SharePoint application sets and jobs 576
- ◆ Mounting SharePoint replicas 581
- ◆ Restoring from SharePoint replicas 584
- ◆ Using SharePoint with RecoverPoint 587
- ◆ Troubleshooting SharePoint issues 591

Overview of support for SharePoint

The following is a summary of Replication Manager's support of SharePoint.

Supported versions

Replication Manager supports Microsoft SharePoint Server 2010 and Microsoft Office SharePoint Server 2007 SP1 and SP2.

The *EMC Replication Manager Support Matrix* provides the most up-to-date lists of application versions, hardware, operating systems, and service packs supported by Replication Manager. To access the *EMC Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.

General

- ◆ Replication Manager creates full replicas of an entire farm (content databases, configuration database, administration database, search databases, search index files). Partial replicas are not supported at the farm level.
- ◆ You can mount a replica of the entire farm, or databases and search indexes on a specific host in the farm.
- ◆ Because a SharePoint configuration is comprised of SQL Server databases, the options in Replication Manager for job creation, replica mount, restore, and recovery are similar or identical for SharePoint and SQL Server.
- ◆ You can extract and restore an individual content database using the `rmsqlrestore` utility, which is described in "Using the `rmsqlrestore` utility for SharePoint" on page 587.

SharePoint 2010

- ◆ Restore is supported at the content database level.

SharePoint 2007

- ◆ Full farm restore and content database restores are supported.

Configuring the SharePoint environment

This section describes SharePoint prerequisites and where to install Replication Manager software in the SharePoint environment.

SharePoint prerequisites

This section lists prerequisites for Replication Manager support of SharePoint:

General SharePoint configuration

- ◆ SharePoint search indexes and search databases must be located on separate LUNs from content databases or logs; this is a SharePoint best practice.
- ◆ At least one SharePoint host must have the Windows SharePoint Services VSS writer (the WSS-VSS writer) enabled on it. This host is referred to as the SharePoint Writer host. Each application set is associated with one writer host when the application set is created.

Typically this is installed on a web front end server, but any host in the farm with SharePoint installed may be used. Consider using the host where a Search Service Application administration component is enabled.

To install the WSS-VSS writer, run the following SharePoint **stsadm** command:

```
stsadm.exe -o registerwsswriter
```

stsadm is in the following location on the drive where SharePoint Products and Technologies is installed:

SharePoint 2010:

```
%COMMONPROGRAMFILES%\Microsoft Shared\Web Server Extensions\14\bin
```

SharePoint 2007:

```
%COMMONPROGRAMFILES%\Microsoft Shared\Web Server Extensions\12\bin
```

You must have local Administrator rights to use **stsadm**.

SharePoint Writer hosts running Windows Server 2008 R2 require Microsoft hotfix KB 2253693.

- ◆ You need to know the credentials for the SharePoint farm account (also known as the server farm account). You are prompted to enter them during application set creation.

This farm account must have local administrative rights on the SharePoint Writer host. Add the account, or the domain group of which it is a member, to the local Administrators group on the SharePoint Writer host.

SharePoint 2007 only

- ◆ Farms with multiple Shared Services Providers (SSPs) enabled are not supported.
- ◆ For full farm restore:
 - Microsoft Office SharePoint Server 2007 SP2 is required at the time of the replication.
 - Hosts running Replication Manager must be at version 5.2.3 or greater.
 - Replicas must be from application sets and jobs that were created in version 5.2.3 or greater. Replicas made from application sets and jobs created in version 5.2.2 cannot be used for full farm restore.
 - Replication Manager 5.2.3 or greater is required at the time of the replication.
- ◆ Verify that Windows SharePoint Services Search and Office SharePoint Services Search are running. These services are needed by Replication Manager to create SharePoint replicas and perform restores and other tasks. If a crawl is scheduled to run infrequently, these services might not be running after a reboot of the host.

SQL Server configuration

- ◆ One common set of credentials for all SQL Server instances used by SharePoint servers in the farm is required. You are prompted to enter SQL Server credentials during application set creation.
- ◆ Certain SQL Server permissions and rights are required to configure application sets and run jobs in a SharePoint environment. The user account must be configured to use either SQL Server authentication or Windows authentication. The Windows user account can either be a member of the local Administrators group or a non-Administrator account. To use a non-Administrator account, refer to the restrictions outlined in the SQL Server chapter of this guide, under [“Required permissions and rights” on page 435](#).

Storage-related

SharePoint data must reside on:

- CLARiiON or VNX storage accessed via FC, iSCSI, Hyper-V pass-through, or VMware RDM
- Symmetrix storage accessed via FC, Hyper-V pass-through, or VMware RDM

Local storage (for example, the C: drive) is not supported.

Replication Manager does not support SharePoint on VNXe.

SharePoint replicas created on CLARiiON or VNX arrays use local replication only. SharePoint replicas on Symmetrix arrays can use local and remote (SRDF/S) technologies.

DNS

If SharePoint hosts are registered in Replication Manager with IP addresses, reverse DNS must be configured.

SharePoint 2010 restrictions

Note the following restrictions on support for SharePoint 2010:

- ◆ If SharePoint Foundation Search is enabled, only one copy can be enabled in the farm.
- ◆ The following are not supported:
 - Full farm restore
 - Configurations using SQL Server database mirroring
 - Remote search
 - Search Service Applications (SSAs) restore

Replication Manager deployment in a SharePoint farm

[Figure 193 on page 570](#) illustrates a typical SharePoint 2010 farm and indicates where Replication Manager Agent is required to be installed. [Figure 194 on page 571](#) shows the same information for a SharePoint 2007 farm.

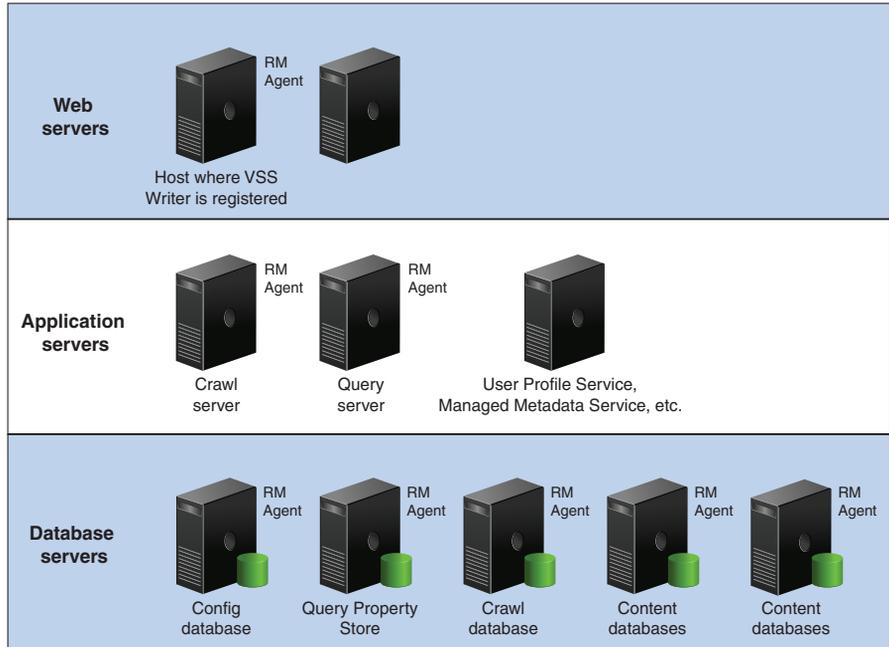


Figure 193 Replication Manager Agent on SharePoint 2010 servers

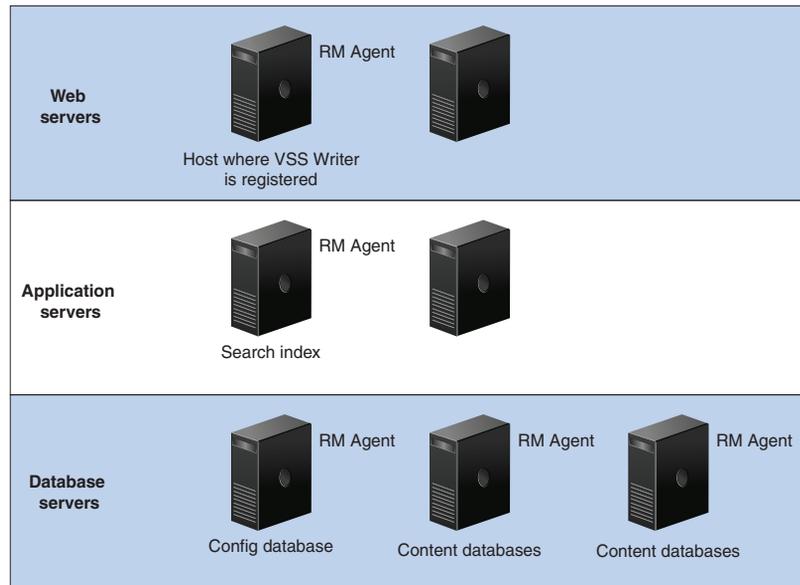


Figure 194 Replication Manager Agent on SharePoint 2007 servers

Where to install Replication Manager Agent on a SharePoint 2010 farm

To determine where to install Replication Manager Agent on a SharePoint 2010 farm, refer to [Figure 193 on page 570](#) and follow these steps:

1. Open the SharePoint Central Administration interface and go to System Settings > Manage servers in this farm.
2. Install Replication Manager Agent on each server running the following services:
 - Microsoft SharePoint Foundation Database
This is the name as seen in Central Administration. It refers to SQL Server instances used by the farm.
 - SharePoint Foundation Help Search
 - SharePoint Server Search
3. Install Replication Manager Agent for SharePoint on at least one host where the SharePoint Services VSS Writer is registered.

Where to install Replication Manager Agent on a SharePoint 2007 farm

Note that the host on which the VSS Writer is registered is not necessarily attached to storage; nevertheless, the Replication Manager Agent and its prerequisites (such as Solutions Enabler) must be present on that host.

To determine where to install Replication Manager Agent on a SharePoint 2007 farm, refer to [Figure 194 on page 571](#) and follow these steps:

1. Open the SharePoint Central Administration interface and go to Operations > Servers in Farm.
2. The following SharePoint servers require Replication Manager Agent. Note their server names:
 - Configuration database server
 - Each server running Windows SharePoint Services Database
 - For a cluster setup, install the Replication Manager Agent on all physical servers that make up the SQL Server cluster
3. Install Replication Manager Agent on each server running Office SharePoint Server Search. To find these servers, run the following command on any server in the farm:

```
stsadm -o enumssp -all
```

The output from this command contains the names of all hosts running Office SharePoint Server Search or Windows SharePoint Services Search. Look for lines beginning with IndexServer Server=. In the example below, host024 is the host name.

```
<IndexServer Server="host024" Path="H:\Program Files\Microsoft Office Servers\12.0\Data\Office Server\Applications" />
```

4. Install Replication Manager Agent on each server on which SPSearch index files are located. To find these servers, run the following command on any server in the farm:

```
stsadm -o spsearch -action list
```

Look for the following lines in the output from this command. In the example below, HOSTNAME123 is the host name:

```
*HOSTNAME123:
  Status: Online
  Default index location: D:\Data\Applications
```

5. Install Replication Manager Agent for SharePoint on at least one host where the SharePoint Services VSS Writer is registered.

Note that the host on which the VSS Writer is registered is not necessarily attached to storage; nevertheless, the Replication Manager Agent and its prerequisites (such as Solutions Enabler) must be present on that host.

License requirements

You need one license for each host in the farm where the Replication Manager Agent is installed.

Note that the SharePoint writer can be enabled on a standalone host, or on a host that is performing other SharePoint functions in the farm. If the SharePoint writer is enabled on one of the farm hosts, you only need one license to cover that host.

Host registration in a cluster environment

For a SharePoint configuration in a cluster environment, register the virtual server network names of the farm's SQL Server instances as Replication Manager hosts:

1. Run MSCS Administration tool to discover which resource group is associated with SQL Server.
2. View the Properties to discover the network name and corresponding IP address.
3. Register the host corresponding to that network name as a Replication Manager host.
4. Select that host when you configure the application set for SharePoint replicas.

Thin SharePoint replicas

The SharePoint 2010 Search Service Application (SSA) can be configured with redundant components, to achieve faster and more reliable searches.

When a SharePoint job starts, Replication Manager detects any SSA redundancies and determines the minimum components required to be replicated. This minimizes the number of LUNs needed for replication, allows replication to run even if hosts or SSA components are unavailable, and results in quicker detection of conditions that could cause a job to fail.

The SSA admin component, online crawl components, at least one query component per index partition, and all databases are replicated.

Thin SharePoint replicas are created by default, provided any redundant component in the SharePoint farm is on an alternate host and on its own LUN.

All clients in the farm must be at Replication Manager Agent version 5.3.1 or greater. If any client is below version 5.3.1, the replication runs but without optimization. Additionally, you should upgrade all clients to version 5.3.1 before updating existing application sets.

In the example in [Figure 195 on page 574](#), if all SharePoint servers are online, Replication Manager determines that there are redundant components for the index partition on Q2. Child jobs on Q2 are skipped, and the LUN for the query component mirror on Q2 will not be part of the replica.

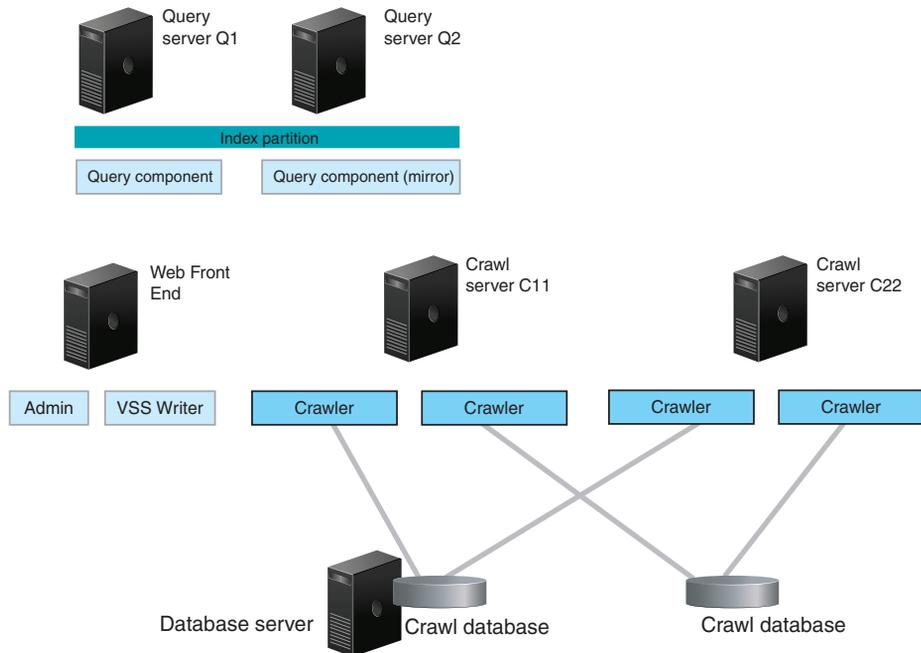


Figure 195 Thin SharePoint replicas

If a query server is down, on the other hand, the job will not fail because the redundant query component is available from the other query server.

In the same example, Replication Manager determines that there are redundant crawler components for the crawl store database. If both crawl servers C11 and C22 are up, the components remain in the replica because they are working on different content. If C22 alone is down (that is, the OSearch service on C22 or the host itself), it is pruned from the replica.

In the Replication Manager Console, job progress and replica history show which components were pruned.

SharePoint application sets and jobs

The following sections describe things you should consider when creating application sets and jobs for SharePoint replicas.

SharePoint application sets

Replication Manager supports creation of replicas at the farm level.

Each application set is associated with a host where the SharePoint services VSS writer was registered (as described in [“General SharePoint configuration” on page 567](#)). If you need to change the host where VSS writer is registered, you will need to re-create the application sets and jobs that were associated with the original VSS writer host.

During application set creation, all SQL Server instances must be running. If an instance is down, application set creation may fail or be incomplete. After creation, expand and examine the application set to make sure all content databases are present.

You are prompted to enter SQL Server and SharePoint farm account credentials during application set creation.

A SharePoint application set can contain only one SharePoint farm and it cannot contain objects from other applications or from file systems.

Figure 196 on page 577 shows the SharePoint application set creation.

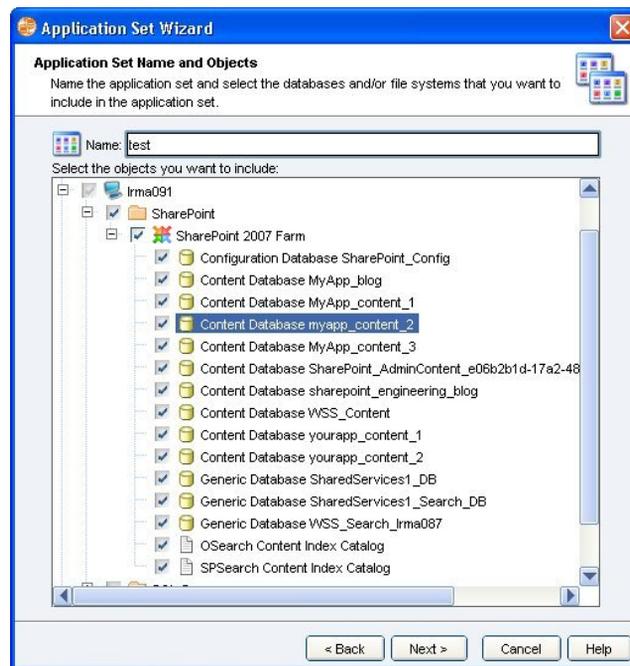


Figure 196 Creating a SharePoint application set

Farm changes that require an application set update

Certain changes to the storage layout of a SharePoint farm require you to update the corresponding application set. This section describes those changes and how to update the application set.

In general, an application set needs to be updated when there is creation, deletion, renaming, or movement of databases or search indexes.

If you add a host to or remove a host from a farm, you need to create a new application set to reflect the change. Do not update an existing application set in those cases.

Note that the SharePoint VSS Writer host is not considered part of the farm unless the host has databases or search components.

The specific changes for supported SharePoint versions are listed below.

SharePoint 2010

Changes to a SharePoint 2010 farm that require an application set update are:

- ◆ Modifications to the Search Service Application (SSA)
- ◆ Any change to content databases, such as creating, removing, or moving to other LUNs
- ◆ Enabling or disabling service applications
- ◆ Modifications to SharePoint Foundation Search

SharePoint 2007

Changes to a SharePoint 2007 farm that require an application set update are:

- ◆ Any change to content databases, such as creating, removing, or moving to other LUNs
- ◆ Modifications to other SharePoint databases
- ◆ Modifications to Office Search (OSearch) or Help Search (SPSearch)

How to update a SharePoint application set

Use the Update button (on the application set Objects tab) to query the SharePoint farm and make the appropriate changes to the application set.

Note: The Update button does not detect when a host is added to or removed from a farm. Create a new application set after adding or removing a host.

To update a SharePoint application set:

1. Right-click the SharePoint application set and select **Properties**.
2. Click the **Objects** tab.
3. Click the SharePoint farm object.
4. Click **Update**.
5. Click **OK** and confirm that you want to update the application set.

SharePoint jobs

SharePoint job creation prompts for the same configuration information as job creation for SQL Server 2005/2008 jobs.

Figure 197 on page 579 shows the advanced settings panel of the job wizard for SharePoint.

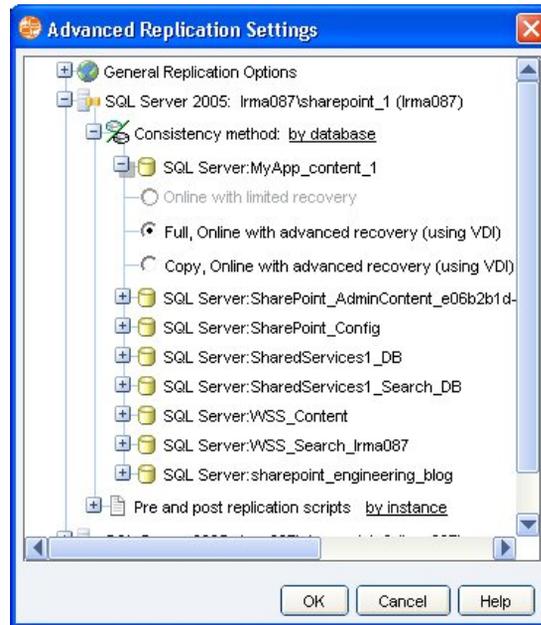


Figure 197 SharePoint job wizard, advanced replication settings

For information on consistency methods, recovery options, and mount options when creating a SharePoint job, refer to “Understanding SQL Server application sets and jobs” on page 439 in the SQL Server chapter of this guide. The SQL Server 2005/2008 information is applicable to SharePoint when a particular SQL Server version is specified.

Search activity is paused during replication

In order to ensure that the SharePoint farm is in a quiesced state, Replication Manager pauses all search activity on the crawl servers before starting a replication. Depending on the content that is being crawled, this step might take several seconds or even minutes.

If Replication Manager fails to pause the crawl, the replication also fails. When the replication is complete, Replication Manager resumes the crawl.

Overriding thin SharePoint replicas

Thin SharePoint replicas allow jobs to complete even if hosts or components are unavailable at the time of replication. However if the following conditions are true, by default Replication Manager will fail the job:

- ◆ the admin component is unavailable
- ◆ no query components are available for a given index partition
- ◆ no crawler component is available
- ◆ layout of the SSA changed from when the application set was created

To override the default behavior, set the following environment variable on the Replication Manager Server :

`EMC_ERM_SP2010_SMARTSSA_SUPPRESS_ERRORS`

When set, the variable replicates the SQL Server databases specified in the application set but none of the SSA components.

When setting an environment variable on the Replication Manager Server, first stop the Replication Manager Server Service, set the variable to '1', then restart the Replication Manager Server Service.

Mounting SharePoint replicas

This section describes Replication Manager mount functionality for SharePoint replicas. General information on mounting replicas is described in [Chapter 5, "Mount, Restore, and Recovery,"](#).

SharePoint mount capabilities

Replication Manager can perform the following mounts of SharePoint replicas:

- ◆ Mount to a production host or alternate host.
- ◆ Mount the entire farm, or all databases on a specific SharePoint server in the farm.
- ◆ From a replica made on a multiple server farm to a single mount host or to different mount hosts.
- ◆ Mount to an alternate mount host in the same location as the production host.
- ◆ Mount to an alternate mount host to a new location (determined by adding an alternate mount path to the pathname).
- ◆ Mounting back to the same cluster is not supported. You need a mount host to mount the replica for backup.

SharePoint mount prerequisites

This section describes prerequisites for mounting a SharePoint replica.

Production host requirements

When mounting a replica to the production host, select a different instance of SQL Server to prevent overwriting of the production database.

Alternate host requirements

To mount to an alternate host, the requirements are:

- ◆ The identical operating system, file system, HBA drivers, and application versions must be installed on the mount host as are installed on the production host.
- ◆ Replication Manager Agent software for SharePoint is installed. It must be the identical version that is installed on production host.
- ◆ The mount host must be registered with the Replication Manager Server.

SQL Server operations as part of SharePoint mount

If you intend to perform any SQL Server operations as part of the mount, alternate mount hosts must have the same version of SQL Server as the production hosts.

SharePoint mount options

Replica mount options for SharePoint are the same as those for SQL Server 2005/2008.

For information on mount options for SharePoint replicas, refer to the SQL Server chapter in this guide, beginning with “[Mount recovery modes with limited recovery](#)” on page 466. The SQL Server 2005/2008 information is applicable to SharePoint when a particular SQL Server version is specified in that chapter.

Mount options for multi-host SharePoint configurations

A multi-host SharePoint replica shares behaviors with those of a replica from a federated application set. Note the behavior of the following mount options in the context of a multi-host SharePoint replica:

- ◆ **Unmount the replica on job completion** unmounts the part of the replica that is mounted on the specified mount host. Replication Manager unmounts parts of the replica mounted on other mount hosts only when you choose this checkbox for the other hosts.

Note: If you have SAP BRbackup compliant jobs, this option is forced (unchecked).

- ◆ If you enable the **Fail the replica if the mount fails** option, if any one mount fails, the entire replication fails and anything that had been mounted will be unmounted.

Note: If you have SAP BRbackup compliant jobs, this option is forced (checked).

If you clear **Fail the replica if the mount fails**, then these are the behaviors:

- If all mounts succeed, then the replication status will say Mounted.
- If some mounts succeed, but some others fail, the status will be Mounted, but there will also be a warning flag attached to the replica.

- If all mounts fail, the status will say Mount Failed, and a warning flag will be attached to the replica.
- ◆ If you choose to include post mount scripts and/or backup scripts, remember that these scripts must reside on the selected target mount or backup host (respectively).
- ◆ If you enable **Fail the replica if the post mount script fails**, and the postmount script fails, the entire replication fails and anything that had been mounted will be unmounted.
- ◆ Clearing the checkbox **Fail the replica if the post mount script fails** results in the following behaviors:
 - If all scripts return success, then the replication status will say Mounted.
 - If some scripts fail, the status will be Mounted, but there will also be a warning flag attached to the replica.
- ◆ Thin SharePoint replicas of the SSA may omit components and hosts from a replication. Any scripts on omitted hosts will not run and these items will not be available for mount.

Partial mounts of multi-host SharePoint replicas

When you mount a multi-host SharePoint replica, you can mount data from selected hosts, leaving data from other hosts unmounted. Furthermore, you can remount a partially-mounted replica and include data from additional hosts by selecting the replica and running the Mount command again.

Restoring from SharePoint replicas

This section describes restore for SharePoint replicas. General information on restoring from replicas is described in “Mount, Restore, and Recovery” on page 161.

Restore capabilities for SharePoint replicas

When you restore from a SharePoint replica, you can choose to restore one or more content databases (SharePoint 2010 and SharePoint 2007) or the entire farm (SharePoint 2007 only).

Figure 198 on page 584 shows the selection of three content databases for restore.

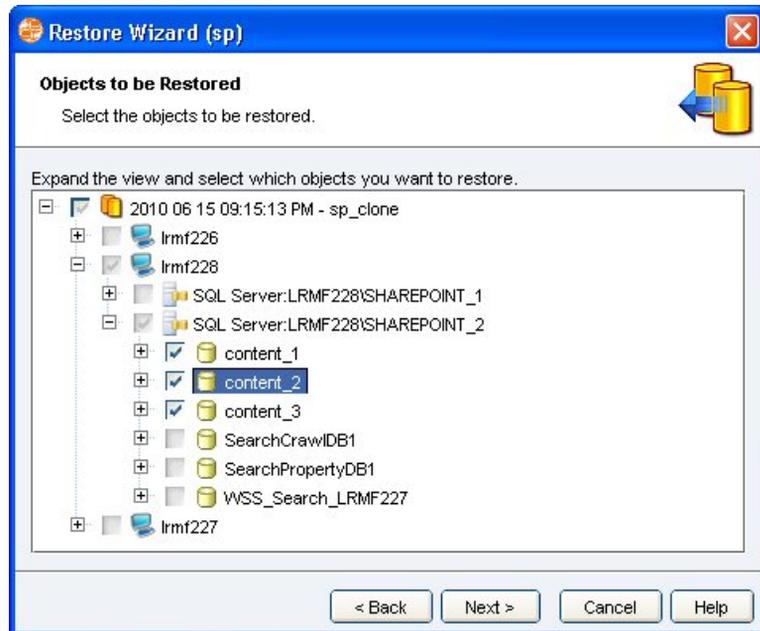


Figure 198 SharePoint restore wizard, objects to be restored

SharePoint restore prerequisites

Note the following before starting a restore of a SharePoint replica:

General

- ◆ The SharePoint Services VSS Writer must be running.
- ◆ EMC recommends that you do not schedule a search crawl to run during the period of the restore.
- ◆ SQL Server instances involved with the restore must be started and operational.

SharePoint 2010 only

- ◆ Stop the Windows SharePoint Services Timer on all SharePoint hosts and wait for any stored procedures to finish running (this can take several minutes). Do not restart the service until after the restore is complete.

SharePoint 2007 full farm restore

- ◆ Microsoft Office SharePoint Server 2007 SP2 must have been in use when the application sets and replicas were created.
- ◆ Replicas must be created from application sets and jobs that were created in Replication Manager 5.2.3 or greater.

Considerations for partial restore

Partial restore is available for content databases only. Changes made to the SharePoint configuration will persist when you restore a content database from a replica. For example, if you make a replica, then create a new site collection, the new site collection will still be present in the SharePoint configuration after you restore from the replica, even though the replica was made before you created the new site collection.

Replaying transaction logs after restore from SharePoint replica

To apply SQL Server transaction logs that were backed up, perform the following steps:

1. If a full farm restore is required, restore the entire farm. This will set the database mode to Recovery for all databases in the farm.
For restoring one or more content databases only, start with step 2.
2. In the Restore wizard, select the content databases to be restored, in the next panel select No Recovery
3. Refer to the Microsoft MSDN Library article entitled "How to: Restore a Transaction Log Backup (SQL Server Management Studio)" for the procedure to restore the SQL Server transaction logs into the content database.

SharePoint restore steps

4. Verify the contents of the SharePoint web application.

To restore from a SharePoint replica:

1. Expand Application Sets.
2. Select the application set that contains the replica you want to restore.
3. Right-click the replica you want to restore and select **Restore**.
4. Select the objects to restore. Note that individual selection of the following is disabled: indexes, configuration database, admin database.

When the parts of a replica that you select for restore share volumes with parts that you did not select, Replication Manager automatically selects those additional parts of the replica.

5. Select restore options. At this point you are dealing with SQL Server databases, so the options are the same as those available when you restore a SQL Server 2005/2008 replica:

- Back up the transaction logs before restoring
- Recover the database (No Recovery, Recovery, Standby)
- Replace (disables important checks during restore)

For full farm restore, only the **Replace** option is available.

For more detail on the options, refer to the online help.

After full farm restore

After a full farm restore is complete and before returning the farm to production:

1. Follow the procedure in Microsoft Knowledgebase article KB939308. This procedure will prevent issues caused when contents of the file system cache on the front-end servers are newer than the contents of the configuration database (which were just restored).
2. On each host running Windows SharePoint Services Timer, restart the World Wide Web Publishing Service.
3. After any restore, follow the steps in your recovery plan. This should include verifying:
 - Shared Services Provider
 - Expected content sources

- Search settings
- Crawl rules

Using the rsmqlrestore utility for SharePoint

Use the rsmqlrestore utility to restore individual databases that reside on the same LUN. The utility is described in the SQL Server chapter under “Using the rsmqlrestore utility” on page 473.

Using SharePoint with RecoverPoint

Beginning with Replication Manager 5.4.1, SharePoint is supported on RecoverPoint software. Specifically, Replication Manager 5.4.1 supports SharePoint 2010 farms protected by RecoverPoint.

Note: Thin SharePoint replicas are not supported.

Interoperability considerations

When using SharePoint with RecoverPoint, note:

- ◆ You must upgrade the Replication Manager server, console, and all Replication Manager clients to Replication Manager 5.4.1 for SharePoint support on RecoverPoint.
- ◆ All SharePoint components must be protected by RecoverPoint and mapped to RecoverPoint consistency groups due to RecoverPoint restores at consistency group granularity:
- ◆ Content databases should be in their own consistency groups, separate from all other databases and components.
- ◆ Spread content databases across consistency groups according to desired restore granularity.
- ◆ SSA databases and components should be in their own consistency groups, separate from any other databases.
- ◆ Service applications should be kept separate from config, content, and SSA.
- ◆ Physical and virtual hosts (VMware virtual machine using RDM disks and iSCSI) are supported.
- ◆ All RecoverPoint technologies that Replication Manager supports are supported with SharePoint.
- ◆ Creating CLARiiON SnapView clone and snap replicas of RecoverPoint targets is supported.

Creating SharePoint application sets

When creating the application set:

- ◆ Configure the farm and protect the farm with RecoverPoint first before creating application set.
- ◆ Enable RecoverPoint and enter the RPA management host name when registering each farm host.

RecoverPoint jobs mount options

When using SharePoint with RecoverPoint jobs, note:

- ◆ The following RecoverPoint bookmarks are supported replication options:
 - RecoverPoint CDP
 - RecoverPoint CRR
 - RecoverPoint CLR
- ◆ The mount options for a CLR job are shown in [Figure 199 on page 588](#):

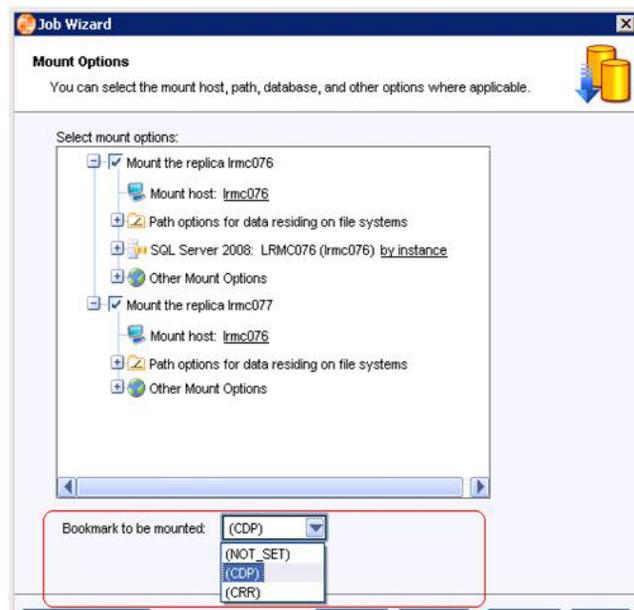


Figure 199 Mount options for CLR jobs

- There is one mount type for all hosts, either CDP mount or CRR mount, no mixed mount type.
- The drop-down list for selecting mount type determines which side of the bookmark to mount - CDP or CRR
- ◆ The mount options for a CDP/CRR job are as follows:
 - CDP job mounts CDP side bookmark
 - CRR job mounts CRR side bookmark

Note: For CDP and CRR jobs, there is no need to select which side to mount.

Running RecoverPoint jobs

When you run a RecoverPoint job:

- ◆ A CLR job creates two replicas, one For CDP, one for CRR.
- ◆ A CDP or CRR job creates one replica.

Any Point in Time (APIT) mount and restore

Consider the following when doing Any Point in Time mount and restore.

APIT mount

When doing an APIT mount:

- ◆ One protection window per application set using the longest consistency group protection window.

Note: The earliest time (date and time) that can be selected for APiT Mount is the start of the protection of the CG with the longest protection window.

- ◆ At least one SPiT replica created by running RecoverPoint job must exist before APiT mount can be enabled.

Note: A CDP replica created by running a CDP or CLR job must exist before APiT mount on the CDP side can be enabled. Same for CRR APiT mount.

APIT restore

When doing an APIT restore:

- ◆ Same as APiT mount, one protection window per application set using the longest CG protection window.
- ◆ At least one SPiT replica created by running RecoverPoint job must exist before APiT restore can be enabled.
- ◆ Full farm restore is disabled since APiT is crash consistent.

GUI intelligence for restoring SharePoint RecoverPoint replicas

Replication Manager performs an affected entities check based on the RecoverPoint consistency group. The GUI automatically:

- ◆ Selects the other replica objects in the same RecoverPoint consistency group when you select a replica object for restore.
- ◆ Displays an error message when you select some but not all replica objects in the same RecoverPoint consistency group for restore. The error message informs you that not all objects in the same consistency group are selected and restore will not be allowed.

For example, under User Databases, when WSS_Content is selected for restore, the GUI automatically selects WSS_Content_21249 because the two content databases are in the same consistency group. However, if you manually uncheck WSS_Content_21249 but leave WSS-Content checked and proceed, the GUI displays an error message informing WSS_Content_21249 is not selected and the restore will not be allowed. This is because the two databases reside on the same consistency group.

In addition, if the content database is configured with other non-restorable SharePoint components in the same consistency group, the GUI displays an error message informing that the content database can not be restored because it is in the same consistency group with non-restorable SharePoint components.

Troubleshooting SharePoint issues

This section covers troubleshooting of issues with SharePoint and its interaction with Replication Manager.

Unable to perform operations for SharePoint

Symptoms: Failure of various Replication Manager operations for SharePoint, such as creating application sets, or creating or running jobs.

Cause: A heavily loaded SharePoint host with low system memory (1GB) can cause the SharePoint VSS writer to stop functioning, and prevent Replication Manager operations for SharePoint from completing successfully.

Remedy: Increase SharePoint system memory.

Restore fails with VSS_UNKNOWN_ERROR

Symptom: Restoring a SharePoint replica fails and returns the error VSS_UNKNOWN_ERROR.

Cause: The replica was mounted with the No Recovery option, which makes the database unusable until it is restarted.

Remedy: Bring the database online manually before restoring or performing subsequent replications.

Replication or restore failure after SharePoint server reboot

Symptom: Replication or restore failure.

Cause: After a reboot, the Windows SharePoint Search service or Office SharePoint Server Search service (or both) are restarted by the Windows SharePoint Services Timer service. If an operation is run during the short period of time after the reboot and before the services are started, the operation fails.

Remedy: Wait for the services to start, or start them manually, and try the operation again.

Replica fails after restore of a content database

Symptom: After a content database restore, the next replication fails to find that content database.

Cause: SharePoint did not reconnect the content database to a web application.

Remedy:

1. Run SharePoint Central Administrator.
2. Navigate to **Applications Management > Content Databases**.
3. Select the web application that the content database belongs to.
4. Select **Add a content database**.
5. Enter the appropriate information and click **OK**.

Full farm restore fails with “replica does not contain detailed information” error

Symptom: Full farm restore fails with the error “The replica does not contain detailed information about SharePoint Shared Services Providers that is required to perform a full farm restore.”

Cause: The application set or replica was not created under Microsoft Office SharePoint Services 2007 SP2.

Remedy: Create application sets and replicas under SharePoint 2007 SP2 and with Replication Manager 5.2.3 or greater.

Unable to detect SharePoint 2010 farm when configuring an application set

Symptom: SharePoint VSS Writer has been enabled on a host and the service is running. When configuring an application set on that host, the SharePoint Farm is not visible.

Cause: Multiple SharePoint Foundation Help Search services have been enabled.

Remedy: Disable the extra instance of the SharePoint foundation Help search services in the farm, leaving at most one. Create the application set and update any existing application sets.

Crawl did not resume after replication

Symptom: After a replication job, one or more SharePoint search crawls that were running before the replication remain in the paused state.

Cause: Heavily loaded crawl servers may be very slow to respond to the resume crawl command issued by Replication Manager at the completion of a replication. Other failures during replication can prevent Replication Manager from resuming crawls automatically.

Remedy: Insure that all search hosts are online and functional. Manually resume search crawling and search background activity by running one of the following commands from a SharePoint host that has Replication Manager Agent installed.

SharePoint 2007:

```
"%ProgramFiles%\EMC\rm\client\bin\rm_sputil.exe"  
-resumecrawl
```

SharePoint 2010:

```
"%ProgramFiles%\EMC\rm\client\bin\rm_sputil14.exe"  
-resumecrawl
```


UNIX File System Procedures

Replication Manager protects UNIX file systems by creating and managing application sets that contain those file systems. This appendix covers the specifics of UNIX file system support. It includes the following sections:

- ◆ Creating UNIX file system replicas 596
- ◆ UNIX file system mount and restore..... 597
- ◆ UNIX raw partition concepts 599
- ◆ Linux logical volume manager support 600
- ◆ Using pre- and post-replication application scripts..... 601

Creating UNIX file system replicas

The Replication Manager UNIX File System Agent replicates databases and software applications that store their data in supported UNIX file systems. The agent supports other applications and databases, besides those specifically supported by Replication Manager, because you can implement your own database or application shutdown and startup scripts.

Databases or applications that are good candidates for replication and recovery include those that are:

- ◆ Large and reside on supported storage (for example, Symmetrix)
- ◆ Able to list all the filenames for the data
- ◆ Able to freeze the data, either by taking the application offline or by some other means that ensures the data remains internally consistent

Without the use of any scripts, the file system remains mounted for replication. Replication Manager flushes the file system I/O buffer immediately before the mirrors are split to ensure that all changes have been synchronized to disk before the replica is created.

UNIX file system mount and restore

The UNIX File System Agent allows you to initiate and control mounts and restores of UNIX file systems. Replication Manager can perform mounts:

- ◆ On an alternate mount host to the same location (path) as the production host
- ◆ On an alternate mount host to a new path (determined by adding an alternate root to the pathname)
- ◆ On an alternate mount host to a new path (determined by pathmapping)
- ◆ To the original production host in a new location (determined by adding an alternate root to the pathname)
- ◆ To the original production host in a new path (determined by pathmapping)

Note: The file-system version and volume manager (if applicable) on the mount host must be identical to the file system (and volume manager) version on the production host. Replication Manager software does not verify the configuration. *It is the responsibility of the user to verify that the file system versions are compatible.*

The Replication Manager UNIX File System Agent can restore specific data. The restore feature is subject to various limitations as described in [“Restore limitations for data in file systems” on page 599](#). You can restore data managed under logical volumes at the volume group level. Refer to [“Restoring UNIX logical volumes” on page 597](#), and [“Issues when multiple applications share volume groups or hypervolumes” on page 121](#).

Restoring UNIX logical volumes

When you are replicating a file system that is located on a logical volume, you can replicate, mount, and restore by volume group.

Replicas based on volume groups have the following characteristics:

- ◆ All of the devices in the volume group are replicated.
- ◆ All of the devices in the volume group are imported.
- ◆ On restore, all of the devices in the volume group are restored.

**CAUTION**

If you are restoring a Veritas volume group that has the same name as a logical volume that resides in a different volume group, the restore fails. Ensure that volume group and logical volume names are unique in your system.

For example, if there are two file systems on two different devices and both devices reside in a single Veritas volume group, if you ask to restore one of the file systems, both file systems will be restored. (First, all of the logical volumes are deported from the production data server; after the restore, all of the volumes are imported to the production data server.)

UNIX raw partition concepts

You can build a UNIX file system directly on raw partitions (slices). A Solaris example of a slice is file systems built on the special device file `/dev/dsk/c0t0d1s3`. In these situations, Replication Manager replicates the whole device containing the slice.

Replication Manager also restores the entire device; the intended file system will be restored along with other data, such as other file systems in other slices on the same device. However, Replication Manager will fail a restore if there are file systems mounted that it does not know about.

Mounts with multiple filesystems on different partitions of the same LUN are not supported.

Multiple partitions on UNIX and Linux LUNs are not supported.

Restore limitations for data in file systems

Lack of file-by-file granularity has a specific limitation: Restores reconstruct the entire file system, overwriting all other files already present.

Other restore limitations of the file system on a raw partition methodology are:

- ◆ You cannot request individual files for restore because the whole file system gets restored.
- ◆ You must manually re-create the mount point at the start of the restore session.
- ◆ The restore unmounts the file systems from the production host to ensure the files cannot be accessed and to invalidate any in-memory image of the file data on the host during the restore.
- ◆ After the restore, all restored file systems are mounted regardless of their state prior to the restore.

Limitations with MPIO devices

Replication Manager does not support using a filesystem mounted on a Linux partition while using MPIO devices.

Linux logical volume manager support

Replication Manager supports Linux environments built on logical volumes; however, there are certain restrictions associated with the setup of the environment of which users must be aware:

- ◆ The length of the full path to the logical volume cannot currently exceed 31 characters with certain versions of Solutions Enabler.
- ◆ Physical volumes should not be created on disk partitions, instead users should use the entire disk when creating physical volumes.
- ◆ Create all physical volumes that you intend to include in the same volume group on the same command line:

For example, do not create physical volumes for the same volume group in this manner:

```
pvcreate /dev/sdc
pvcreate /dev/sdd
vgcreate test_vg /dev/sdc /dev/sdd
```

Instead, use the following two commands:

```
pvcreate /dev/sdc /dev/sdd
vgcreate test_vg /dev/sdc /dev/sdd
```

- ◆ In this environment, the LUNs that hold the replica should not be visible to the production host, therefore production mount is not supported.
- ◆ For LVM1, mounting multiple replicas to the same mount host is not supported unless you mount the first replica (at least once) on the mount host before the second replica is created.

For LVM2, mounting any two replicas of the same source device simultaneously on the same mount host is not supported.

- ◆ When mounting replicas, if you are not using the default partition of 1, you must set the environment variable, `EMC_ERM_RHEL_PARTITION_SLICE_FIX`, to a non-zero integer.

Using pre- and post-replication application scripts

The Replication Manager UNIX File System Agent takes certain default actions on the host before and after splitting the mirrors. You can add custom before and after actions by specifying them in pre- and post-replication scripts. (Refer to [“Creating UNIX file system replicas” on page 596](#) for more background information.)

Your scripts can quiesce data to ensure consistency before a split by:

- ◆ Shutting down and starting up your application
- ◆ Putting your database or application into (and out of) an Online Backup mode, if such a mode is available

With or without the use of scripts, the agent always flushes the file-system I/O buffer immediately before creating a replica to ensure that all changes have been synchronized to disk. Unlike the Oracle agent, there are no default actions omitted when you use the scripts. (Also, unlike the Oracle agent, there is no checking for a state after which the replication fails.)

To use your own pre- and post-replication scripts for the Replication Manager UNIX File System Agent:

1. Name the script and its location (by specifying the full pathname) while configuring the job in the console.

Note: If a script acts on multiple file systems, you can place the script in any file system and provide a fully qualified pathname. Replication Manager accesses the script from that specified file system each time it runs that script.

2. Ensure that the script is owned by the root user (or the application user) and is executable by the root user (or application user) and group.
3. Do not assume the script inherits any UNIX environment variables. Explicitly set PATH and any other environment variables you might need within the script.

General guidelines for scripts

General guidelines for scripts are as follows:

- ◆ Each script should be a shell script, whether Bourne shell (sh), C shell (csh), or Korn shell (ksh). All scripts require their respective headers. For example, for sh, the script needs the header `#!/bin/sh` to work correctly.
- ◆ Do not send more than 255 characters of output to either stdout or stderr. If you do, these will be truncated. If more output or error data is needed, set up and use your own log file.
- ◆ Return a 0 (zero) status to continue with the replication; return a nonzero status to fail the job.

Replication Manager can protect Windows-based NTFS file systems by creating and managing application sets that contain the file system. This appendix covers the specifics of NTFS support. It includes the following sections:

- ◆ Creating NTFS replicas..... 604
- ◆ Windows host system layer concepts 606
- ◆ Using pre- and post replication application scripts 608

Creating NTFS replicas

The Replication Manager Windows NTFS Agent can replicate software applications that store their data in supported Windows-based NTFS file systems. The NTFS Agent allows you to implement your own database or application shutdown and startup scripts and allows you to support applications and databases besides those specifically supported by Replication Manager.

Replication Manager is best suited to replicate large datasets, collectively and rapidly. It replicates the entire physical disk on which the NTFS file system resides.

Databases or applications that are good candidates for replication and recovery include those that are:

- ◆ Large and reside on supported storage (for example, Symmetrix)
- ◆ Able to list all the filenames for the data
- ◆ Able to freeze the data, either by taking the application offline or by some other means that ensures the data remains internally consistent

NTFS replication operation

Without the use of any scripts, the file system remains mounted for replication. The system flushes the file-system I/O buffer immediately before splitting the mirrors to ensure that all changes have been synchronized to disk.

NTFS mount and restore functions

Using the NTFS agent, Replication Manager can initiate and control mounts and restores of NTFS. The product can perform:

- ◆ Mounts on an alternate mount host to the same location as the production host
- ◆ Mounts on an alternate mount host to a new location (determined by adding an alternate drive letter and optionally path information to the beginning of the pathname)
- ◆ Mounts on an alternate mount host to a new location (determined by pathmapping)
- ◆ Mounts to the original production host in a new location (determined by an alternate drive letter and optionally path information to the beginning of the pathname)

- ◆ Mounts to the original production host in a new location (determined by pathmapping)
- ◆ Restores from replicas to the production database server

The NTFS Agent can specify data from the replica to restore, subject to various limitations as described in [“Restoring file systems” on page 607](#).

Windows host system layer concepts

Because Replication Manager creates replicas at the device level, replication and restore granularity depends on device layout. Therefore, it is important to carefully lay out your devices so that they provide adequate granularity for your applications during replication and restore.

Windows host data layer concepts

For Windows, Replication Manager must replicate and restore each *file system as a device*.

The datafiles must reside on supported storage (for example, Symmetrix devices):

- ◆ Replication Manager creates a replica of the entire file system or logical volume when datafiles reside within file systems or logical volumes.
- ◆ Replication Manager can create replicas that contain only one partition per physical disk.
- ◆ Replication Manager supports partitions created on Master Boot Record (MBR) disks or GUID Partition Table (GPT) disks.
- ◆ Two terabyte or larger volumes are supported with some limitations. Refer to the *Replication Manager Support Matrix* for more information. To access the *EMC Replication Manager Support Matrix*, go to <http://elabnavigator.EMC.com/>, select PDFs and Guides, then scroll down to Replication Manager.
- ◆ Windows mount points are supported in Replication Manager. However, nested mount points are not supported in the same application set. For example, you cannot have L:\ and a mount point on L:\ (L:\MP1) in the same application set.
- ◆ Replication Manager does not support raw devices on Windows.
- ◆ Replication Manager does not support native dynamic disks on Windows.

Note: When you mount a replica that contains a Windows mount point similar to that just described, it is important to ensure that the mount host has a drive mounted to the drive letter where the mount point has been added. If that drive does not exist on the mount system, the mount fails.

Microsoft recommends that you choose to mount all NTFS replicas as part of the process of creating them.

Replicating NTFS file systems

Replication Manager replicates, and therefore restores, the data specified in the application set at the physical disk level. However, if you build the database directly onto logical volumes, granularity of the restore function is at the coarser logical-volume-group level.

Restoring file systems

Restores overwrite the entire physical disk, overwriting all the files in the file system. Other restore characteristics include:

- ◆ You cannot request individual files for restore because the whole file system is restored.
- ◆ The production file systems must be mounted at the start of the restore.
- ◆ The restore unmounts the file systems from the production server to ensure the files cannot be accessed and to invalidate any in-memory image of the file data on the host during the restore.

Note: Individual files cannot be restored automatically, but you can mount a replica to an alternate host, and then copy individual files manually.

Using pre- and post replication application scripts

The NTFS agent takes certain default actions on the host before and after creating the replica. You can add custom before and after actions by specifying them in pre- and post-replication scripts.

Scripts must be in one of the following formats: .bat or .exe.

With the NTFS agent, your custom pre- and post-replication scripts enable you to perform actions to ensure data consistency for applications that store data in many interrelated files or in a few very large files. Refer to [“Creating NTFS replicas” on page 604](#) for more background information.

Your pre-replication script could shut down your application and your post-replication script could start up your application or put your database or application into (and out of) an Online Backup mode, if such a mode is available. To use the pre- and post-replication scripts for the NTFS agent, you need to do the following:

1. Specify the name of the script and its full pathname when you configure the job in the console.
2. Do not assume the script will inherit any Windows environment variables. Explicitly set PATH and any other environment variables you might need in the script.

This appendix provides definitions to special terms used throughout this book.

A

- admsnap** Server-based software that provides a command line interface to SnapView software running in a storage-system SP. With admsnap, Replication Manager can start and stop sessions and activate and deactivate snapshots on a secondary server system.
- agent software** Different Replication Manager interfaces used to create replicas of database management systems, email applications, and/or file systems. Each of these applications is managed by a different kind of replication software, called an agent.
- any-point-in-time replica** A crash-consistent replica created from a time not marked by a job-inserted event. See also *specific-point-in-time replica*.
- application set** Defined set of production data, specifying a database (or selected tablespaces) or file systems to be replicated. Identified by an application set name.
- ASM** Automated Storage Management, a disk volume manager used for storing Oracle files, ASM allows administrators to add and remove disks while the database is available. Data is automatically striped across all disks in a disk group.

associated BCV A BCV that has been established recently enough for the Symmetrix system to be recording invalid tracks and be capable of doing an incremental establish.

Automated Storage Management See *ASM*.

B

BCVs Business continuance volumes that create copies of data from a standard Symmetrix device (which is online for regular I/O operations from the host). Data is stored on BCV devices to mirror the primary data. Uses for the BCV copies can include backup, restore, decision support, and applications testing. Each BCV device has its own host address, and is configured as a stand-alone Symmetrix device.

C

Celerra iSCSI Network Servers A storage system connected to hosts using iSCSI that is part of the EMC network-attached storage (NAS) product line.

Celerra Replicator Replication Manager uses the Celerra Replicator option to create a read-only, point-in-time replica of a production Celerra file system.

Celerra SnapSure snapshots A point-in-time representation of the data stored on an iSCSI LUN. The snapshot is not necessarily a complete copy and therefore should not be relied on as a data backup replacement.

CLARiiON EMC's line of mid-range storage arrays.

CLARiiON Snapshots See *SnapView snapshots*.

clones See *TimeFinder/Clones*, *SnapView clones*, or *iSCSI host*.

Cluster Continuous Replication (CCR) Exchange CCR combines the replication and replay features in Exchange 2007 with failover features in Microsoft Cluster services. CCR is a solution that can be deployed with no single point of failure in a single data center or between two data centers. EMC Replication Manager integrates with Exchange CCR to offer further protection for Exchange data.

consistent split The ability of a Symmetrix or CLARiiON or VNX storage array to create a replica of data while maintaining the application data on devices in a crash-consistent format. In other words, the data is consistent as of the point in time when the replica was initiated.

Transactions that were in progress, but not yet committed, at the time the replica was initiated will not be part of the replica.

continuous data protection The method of data protection in which all changes to data are continuously captured and tracked, allowing for data recovery to any point in time.

copy on first write An algorithm that copies current contents of a source LUN before it is modified (written to). The copy-on-first-write operation is on a chunk: before the first modification of any disk blocks within a chunk, the software reads and stores the original data of the chunk in the reserved LUN pool. This policy applies only to the first modification of the data. Overwrite of any data that has already had a copy-on-first-write does not require any extra processing since the software saved the original data in the reserved LUN pool.

D

Database Administrator (DBA) The Database Administrator is the user who understands and manages relational database systems (or other similar applications) in your environment. (Database Administrators manage one or more database—or other applications—served by the replication software.) The database administrator should be the expert who is well versed in the operation of a database. The Database Administrator is responsible for restoring replicas of databases and performing recovery when necessary. See [“User roles” on page 36](#) for specific information about each role’s permissions.

disaster recovery (DR) Includes the setup, failover, and failback procedures required to actuate live data from a failed production system to a disaster recovery system, and then back to a new or restored production system.

disaster recovery system The components consisting of a Windows server, Celerra Network Server, and associated software, located at a secondary data center, that replicate the data from the production system at the primary data center.

E**EMC snapshots**

See *TimeFinder/Snaps*.

ERM Administrator

The ERM Administrator is the user who understands and manages both storage configurations and the replication server operating systems in your environment. The ERM Administrator controls storage resource allocation. The ERM Administrator is also the system administrator for the overall information replication and recovery system. Initial tasks include installing the replication server software, the agent software, and the replication console software. Users in this role understand file systems but not necessarily databases or other coordinated applications. See “[User roles](#)” on [page 36](#) for specific information about each role’s permissions.

establishing

The process of initially starting the live synchronization between the production devices and the mirror devices.

**establishing/
reestablishing
(synchronizing) the
mirrors**

The process of initially/subsequently starting the live synchronization between the production devices and the mirror devices.

F**failback**

The process by which live data on the disaster recovery system is failed over to a restored or new production system in a primary data center. Failback as used in this document does not imply a return to any preexisting state on the production server in the primary data center. It involves failing over from the disaster recovery server to the new production server once it is online.

failover

The process by which the replicated data set on any given server transitions to live data, due to failure of the system from which the replication was made. Initially, failover occurs from the primary data center to the secondary data center. In recovery, failover occurs from the secondary data center to the newly functional system at the primary data center.

federated database

A database management system (DBMS) that supports applications and users submitting requests for data that reference two or more RDBMSs or databases to satisfy a single goal. Federated databases make distributed requests to multiple databases.

flashback recovery area An Oracle technology that allows Oracle DBAs to recover databases quickly. Replication Manager integrates with Oracle flashback recovery.

fracture The process of breaking off a CLARiiON or VNX clone from its source. Once a clone is fractured, it can receive server I/O requests.

I

incremental establish Start of the live synchronization between the production devices and the associated *BCVs*.

instant restore General term for capabilities that enable near-instant access to production data in the event of data corruption to a production database. Includes both mounting replicas to an alternate host for restore purposes and restoring the replicas to the production computer.

iSCSI host A computer hosting an iSCSI Initiator.

iSCSI (Internet SCSI) A protocol for sending SCSI packets over TCP/IP networks.

iSCSI Initiator An iSCSI endpoint, identified by a unique iSCSI name, which begins an iSCSI session by issuing a command to the other endpoint (the target).

iSCSI target An iSCSI endpoint, identified by a unique iSCSI name, which executes commands issued by the iSCSI Initiator.

J

job A set of actions that create a replica of a given application set. Optionally, jobs perform other actions on that replica. A job can create, mount, run pre-and post-replication scripts, and run backup operations on a replica. Jobs can be scheduled or initiated on demand. Formerly known as an *Activity*.

L**Local Continuous Replication (LCR)**

Exchange LCR is a single-server solution that uses built-in technology to create and maintain a copy of a storage group on a second set of disks that are connected to the same server as the production storage group. LCR provides asynchronous log shipping, log replay, and a quick manual switch to a copy of the data. EMC Replication Manager integrates with Exchange LCR to offer further protection for Exchange data.

LUN

Acronym for a logical unit. This refers to a device or set of devices, usually in a CLARiiON or VNX storage array. LUNs can be the source or target for a Replication Manager replica.

M**mirror devices**

Storage devices used for replicas. (TimeFinder/Mirror devices are called BCVs.) Mirror devices can be synchronized with the production devices, so as to mirror all data and data changes, and can have the synchronization split, thereby preserving the data at the time of the split. See also *BCVs*.

MirrorView/A

EMC MirrorView/Asynchronous is remote replication software for EMC CLARiiON or VNX arrays that provides highly available data protection across a campus or metro area environment. Data is written to the source array first and then the remote array asynchronously.

MirrorView/S

EMC MirrorView/Synchronous is remote replication software for EMC CLARiiON or VNX arrays that provides highly available data protection across a campus or metro area environment. Data is written to the source array and the remote array simultaneously.

mount hosts

The systems that Replication Manager uses to mount a replica. This may be separate from the production system.

mount replicas

The operation of importing, mounting, and opening the replica on an alternate, surrogate server so that it is available as an independent copy of the original database. (Mounting a replica on the production server is supported for some platforms.) Mounting can include correcting file system, logical volume manager, and database inconsistencies. Alternatively, mounts can import a selected subset of the replica back to the production database. Usually a mounted replica can be opened in Read/Write or Read-Only mode.

multiobject application sets Application sets that contain more than one application (either multiple instances of the same application, or multiple instances of different applications). This functionality accommodates federated databases and other instances where multiple applications are related in some way and should be replicated as a single group.

O

Open Replicator EMC technology that creates remote point-in-time copies of Symmetrix sources for high-speed data mobility, remote vaulting, migration, and distribution. EMC Open Replicator software supports copying data between EMC Symmetrix DMX and qualified storage systems.

Operator Role for users who operate the information replication and recovery software on a day-to-day basis. The only permitted tasks are running jobs and mounting replicas to alternate hosts. See [“User roles” on page 36](#) for specific information about each role’s permissions.

OPS Oracle Parallel Server; a database server option that allows multiple instances to mount a single database from more than one node in a cluster environment. OPS was later replaced with RAC.

Oracle Database management system (DBMS) that is widely used for business applications.

Oracle ASM See [ASM](#).

P

Power DBA Role for users who control the application databases that are replicated using the Replication Manager product. See [“User roles” on page 36](#) for specific information about each role’s permissions.

Power user Role for users who operate the replication and recovery software, but need somewhat more control over the application than an Operator. See [“User roles” on page 36](#) for specific information about each role’s permissions.

primary Replication Manager Server Host that is running Replication Manager Server software, and that controls the replication.

production devices Storage used for the live data, usually made available to a corporate or public audience to facilitate business needs.

production server Production computer that hosts the information system that manages the production data: a database server, Web server, application server, or file server.

protected restore When selected, a process that prevents source writes from being copied to a CLARiiON or VNX clone or Symmetrix BCV during a reverse synchronization.

R

RAC Real Application Clusters; allows multiple Oracle instances on different nodes of a cluster to access a shared database on the cluster to facilitate load balancing.

RecoverPoint Appliance (RPA) The RecoverPoint component that supplies continuous data protection services for applications operating on production hosts and using production storage. The RPA is on a dedicated server that is external to the Replication Manager Server and to the production and recovery hosts.

Real Application Clusters See [RAC](#).

recovery The combined task of restoring a replica and rolling forward the database using redo logs.

reestablishing the mirrors Subsequent resynchronization between the standard production devices and the mirror devices in a storage array. Reestablish implies that the devices have been established previously and that only some tracks require updates.

replica A copy of an application set, resulting from a successful replication operation. Identified by a timestamp. Replicas can be selected for mounting and for restore.

replica history A set of messages that capture the steps that created a replica (including the output of user callout scripts), plus the operations performed on the replica since its creation.

Replication Manager The name of the product described in this document. This product now includes functionality previously sold under the Replication Manager/SE name.

Replication Manager SE	See <i>Replication Manager</i> .
replication operation	The process of copying an application set. Replication Manager creates a replica by synchronizing a set of mirrors to the devices that hold the application set, and then splitting those mirrors.
replication server	Computer that hosts the Replication Manager Server software.
restore mirrors	The act of resynchronizing mirror devices and production devices by having the mirrors overwrite the production devices.
restore replicas	The operation of copying data from a replica to the production database (on the production server), thereby reverting to an earlier state of the database. Includes correcting file system, logical volume manager, and database inconsistencies. This operation may be performed on an entire replica or on a selected subset of the replica (if that subset resides on distinct volumes).
rotations	Replication Manager users can specify a rotation by specifying a maximum number of replicas that will be part of a given rotational set. Replication Manager creates a new replica each time the job runs. However, if the number of existing replicas of that application set exceeds the maximum, Replication Manager first deletes the oldest replica of that application set to prevent the system from creating more than the maximum number of replicas.
S	
SAN Copy	Functionality that allows Replication Manager to move data from one CLARiiON or VNX storage array to another or from a Symmetrix storage array to a CLARiiON or VNX storage array over a high-speed SAN or WAN infrastructure.
secondary Replication Manager Server	Host that is running Replication Manager Server software with a read-only configuration. The Replication Manager database is automatically kept synchronized with the primary server host's Replication Manager database.
simulation	A preliminary execution of a job that helps to identify potential problems without actually creating a replica.
snapshot, CLARiiON, VNX, or Symmetrix	Replica created using CLARiiON Snapshots, SnapView snapshots, or TimeFinder/Snaps (VDEVs). Refer to <i>CLARiiON Snapshots</i> or <i>VDEVs</i> .

- snapshot, SQL Server** A backup functionality that quickly creates a point in time backup of the SQL Server database. Replication Manager uses this functionality to create replicas while the SQL Server database remains online.
- SnapView** Allows you to obtain a copy of a LUN by creating a clone or snapshot. The clone or snapshot can serve for backup, decision support scenarios, or as a base for temporary operations on the production data without damaging the original data on the source LUN.
- SnapView clones** A LUN that is an actual copy of a specified source LUN. The state of the clone determines if it is a byte-for-byte copy of its source. Replication Manager can make SnapView clones available to another host through mount commands or restore using the restore capabilities of the product.
- SnapView snapshots** Instantaneous point in time copies of CLARiiON or VNX LUNs created by CLARiiON SnapView software or VNX Snapshot software. SnapView snapshots are views of a point-in-time image of a source LUN(s). A snapshot occupies no disk space, but appears like a normal LUN to secondary servers. Replication Manager can make SnapView snapshots available to another host through mount commands or restore using the restore capabilities of the product.
- SP** Storage processor on a CLARiiON or VNX storage system. On a CLARiiON or VNX storage system, a circuit board with memory modules and control logic that manages the storage system I/O between the host's Fibre Channel adapter and the disk modules.
- SP A** Storage processor A. A generic term for the first storage processor in a CLARiiON or VNX storage system.
- SP B** Storage processor B. A generic term for the second storage processor in a CLARiiON or VNX storage system.
- specific-point-in-time replica** An application-consistent replica created from an event that was inserted by a job. A job-inserted event puts an Oracle database (for example) application in hot backup mode temporarily and saves additional files so the database can be rolled forward when the replica is mounted or restored.
- splitting the mirrors** The process of stopping the live synchronization between the production devices and the mirror devices.

- SRDF** An EMC technology that allows two or more Symmetrix systems to maintain a remote mirror of data in more than one location. The systems can be located within the same facility, in a campus, or hundreds of miles apart using fibre or dedicated high-speed circuits.
- The SRDF family of replication software offers various levels of high-availability configurations, such as SRDF/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A).
- standard volume** A standard volume, also referred to as STD, is a local volume on a Symmetrix system that typically holds the original production data. In some cases, Replication Manager can also use these volumes as the target of a replica.
- STD** See *standard volume*.
- Storage group** A collection of one or more LUNs on a CLARiiON or VNX array. Storage groups are created using Navisphere.
- Storage pool** User-defined set of storage from which Replication Manager selects the storage to be used for replicas for a given job.
- Symmetrix** The line of high-end storage arrays produced by EMC that provide centralized, sharable enterprise storage. These arrays can create an information infrastructure capable of managing large, complex ultra-dynamic storage environments by consolidating storage from multiple heterogeneous hosts onto a single system.
- Symmetrix Remote Data Facility** See *SRDF*.
- T**
- Thin SharePoint replicas** Efficient, highly available replication of complex SharePoint search configurations.
- TimeFinder** Symmetrix TimeFinder is a business continuance solution that allows you to use special Symmetrix devices called business continuance volumes (BCVs) to create mirrors of Symmetrix data.

TimeFinder/Clones Copies of a source device on multiple target devices. The source and target devices can be either standard devices or BCV devices as long as they are all of the same size and emulation type. Clone copies of striped or concatenated meta devices can also be created, but the source and target meta devices must be completely identical in stripe count, stripe size, and capacity. Once activated, the copy can be instantly accessed by a target's host, even before the data is fully copied to the target device.

TimeFinder/Snaps A host-accessible device containing track-level location information (pointers), that indicates where the copy session data is located in the physical storage. TimeFinder/Snap operations provide instant snapshot device copies, using virtual devices (VDEVs).

U

UDB Universal Database; IBM's database management system (DBMS) that is widely used for business applications.

V

VDEVs Virtual devices; TimeFinder/Snap operations provide instant snap device copies, using virtual devices. A virtual device is a host-accessible device containing track-level location information (pointers) that indicates where the copy session data is located in the physical storage. VDEVs consume minimal physical disk storage, as they store only the address pointers to the data stored on the source device or a pool of save devices.

VNX EMC VNX series (VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500) targets the high-performance, high-scalability requirements of midsize and large enterprises. VNX family software is available in comprehensive packs or modular suites.

VNXe EMC VNXe is a unified storage system that simultaneously provides block and file level access to storage. VNXe includes the EMC Unisphere, which is an application-enabled interface that enables you to easily manage your VNXe system.

SnapView snapshots	See <i>SnapView snapshots</i> .
VNXe SnapSure snapshots	See <i>Celerra SnapSure snapshots</i> .
VSS (Volume Shadow Copy Service)	A Windows service and architecture that coordinate various components to create consistent point-in-time copies of data called shadow copies.

A

- access control lists 120
- active tasks, in console 45
- Administrator's Guide* 40
- agent software 33, 40
- alternate path, mounts using 183
- application sets 609
 - acknowledging failed replicas in 310
 - composite 113
 - creating 116
 - federated 112
 - granting users access 120
 - in console 45
 - mapping two to the same data 117
 - RecoverPoint consistency groups 113
 - removing objects from 119
 - validation 117
- applications, supported 17
- archive logs
 - replicating in Oracle 315
- archive logs, location in Oracle 338
- ASM RAC mount to alternate cluster 374
- ASM. *See* Oracle ASM.
- associated BCV 610
- ATA disk support, CLARiiON 74
- audience, intended readers 19

B

- backing up to tape 514, 547
- backup and recovery 28
- backup scripts 186
- BCVs 610
- Block Change Tracking, Oracle 354

C

- callout scripts
 - alternate file locations 241
 - identifiers for mount, failover, restore 239
 - identifiers for replication 238
 - overview 235
 - privileges 240
- canceling a task 157
- CDS 123
- Celerra
 - IQNs 158
 - NFS 324
 - SnapSure snapshot restore granularity 232
- Celerra NFS 292
- Celerra Replicator
 - overview 80, 81
 - remote snapshot 80, 81
 - restore granularity 232
- Celerra storage failover 292
- CLARiiON
 - ATA disk support 74
 - clones 68, 70
 - definition 610
 - restore granularity 231
 - SAN Copy options 289
 - snaps 71, 610
- clones
 - CLARiiON 68, 70
 - defined 610
 - TimeFinder 58, 60
 - VNX 69, 71
- cluster
 - Exchange restore 556

- Microsoft Cluster restore 234
- Oracle Restore 382
- SQL Server restore 480
- UNIX cluster mount and restore 193
- cluster continuous replication (CCR) 525
- comments, submitting 23, 23
- components 40
- composite application sets 113, 260
- concurrent BCV replicas 123
- consistent split 611
 - CLARiiON and VNX 250
 - mounting replicas 256
 - Oracle 254
 - replicas, creating 253
 - restore 258
 - SQL Server 255
 - Symmetrix 250
 - with federated application sets 124
 - with federated databases 259
- console
 - main window 43
 - overview of 40
 - starting 41
- content panel 43
- control files
 - and consistent split restrictions 315
- control files, location in Oracle 338
- copy job 137
- copy replicas 28
- creating jobs 125
- Cross-Platform Data Sharing (CDS) 123
- customer support 23

D

- daily operations 155
- data layout planning 121
- database
 - administrator (DBA) 611
 - determining mode in Oracle 338
 - federated 112
 - renaming in Oracle 365
 - restoring 226
- datafiles
 - backing up as device 606
 - location in Oracle 338
- DB2DIR environment variable 425

- DB2INSTANCE environment variable 425
- DBA 611
- debug logs 54
- deleting replicas on demand 300
- disaster restart 28
- documentation, related 20

E

- emulation mode 265
- Engenuity Consistency Assist (ECA). *See consistent split.*
- environment variable
 - DB2DIR 425
 - DB2INSTANCE 425
 - ERM_TEMP_BASE 417
 - LD_LIBRARY_PATH 425
 - PATH 425
- evaluation phase 29
- Exchange
 - application set 497, 523
 - circular logging 521
 - creating replicas 496, 522
 - import replica from tape 517, 550
 - locating data and logs on separate volumes 520
 - pre- and post-replication scripts 555
 - production host setup 491
 - replica import from tape 516, 548
 - restore considerations 510, 545
 - restore options 543
 - restoring at storage group level 545
 - restoring with VSS 508, 541
 - rules for mounting replicas 184
 - selecting application host objects 543
 - selecting full or partial restores 509, 543

F

- failover, Celerra 292
- federated application sets
 - configuring 262
 - consistent split requirement 124
 - creating 262
 - creating jobs for 265
 - mount options 271
 - performance 274
 - restrictions when changing 119, 264

- terminology 260
- file-systems restore limitations 599
- flash recovery area 113
- full SAN Copy 74, 285

G

- getting started panel 43
- granularity, of replicas 121

H

- HACMP
 - and RAC 384
 - mount and restore 193
- help
 - from technical support 23, 23
 - menu 44
- hosts, in console 45
- HP Serviceguard 193
- Hyper-V 104
- hypervolumes, shared by multiple applications
 - 121

I

- icons 46
- importing from tape, Exchange 516, 548
- incremental SAN Copy 77, 286
- installing Replication Manager 40
- IQNs, Celerra 158
- IQNs, VNXe 158
- IRCCD 33
- IRD 33

J

- jobs
 - canceling 157
 - creating 125
 - deleting 149
 - for federated application sets 265
 - in console 45
 - managing 124
 - monitoring progress while running 157
 - running on demand 157
 - running simultaneously on same application
 - set 149

- scheduling 152
- simulating 142

K

- Keep LUNs Visible option 94

L

- latent error phase 29
- LD_LIBRARY_PATH environment variable 425
- limitations, file system restores 599
- link job 137
- linked copy job 138
- Linux
 - raw devices 405
- local continuous replication (LCR) 525
- log files 54
- logging into the console 41
- Logical Partitions (LPARs) 107
- logical volume manager (LVM)
 - Linux 600
 - NTFS 606
 - support for heterogeneous 113
 - UNIX 597
- LUNs
 - visibility after unmount 189

M

- managing
 - application sets 116
 - job schedules 152
 - jobs 124
- Microsoft Exchange. *See* Exchange
- Microsoft iSCSI Initiator 102
- Microsoft SQL Server. *See* SQL Server
- mirror sessions in SQL Server 2005/2008 448
- MirrorView
 - clones of secondary device 70, 71
 - in CLARiiON and VNX summary 74
 - in overview of remote replication types 290
 - snaps of secondary device 73
- monitoring running tasks 157
- mount as Real Application Cluster 334
- mount location, alternate 168
- mount scripts 186
- mounting RecoverPoint replicas

- image access options 216
- procedure 216
- mounting replicas
 - CLARiiON static mount 189
 - performance guidelines 222
 - to the production host 182
 - troubleshooting 245, 248
 - using a substitution table (path map) 178
 - using an alternate root path 175
 - using callout scripts 239
 - using path mapping 177
- mountpoints, support for in Windows 606
- multiple applications
 - restoring replicas that contain 233
- Multi-Threaded Server (MTS) *See Oracle*.

N

- network file system 292
- NFS 292
- NFS datastore 87
- normal logs 54
- NTFS
 - agent scripts 608
 - file system agent scripts 608
 - logical volumes 606
 - mount and restore functions 604
 - pre- and post-replication scripts 608
 - pre- and postscripts 608

O

- online help 54
- online replications without hot backup mode 254
- Open Replicator 67
- Oracle
 - archive logs, locating 338, 343
 - ASM RAC mount to alternate cluster 374
 - ASM RAC mount to cluster 334
 - configuring listener 338
 - consistent split 254
 - control files, locating 338
 - controlling where temp files are stored 355
 - database mode 338
 - datafiles, locating 338
 - do not perform database operations 362
 - failsafe restores 383
 - flash recovery area 113

- mount and restore functions 361, 379
- mount restrictions 182
- Multi-Threaded Server configuration 314
- pre- and post-replication scripts 390
- RAC and HACMP 384
- RAC to RAC 334, 374
- recovering replicas created on raw devices in
 - Linux 405
- RecoverPoint replica 387
- redo logs, locating 338
- replicating flashback recovery area 315
- rules for mounting replicas 184
- SFRAC 397
- sp-file 317, 371
- SYSDBA users, locating 338
- System Global Area (SGA) 407
- user privileges impact on available features 337

Oracle ASM

- configuration with RAC 338
- database layout restrictions 344
- disk groups 345
- rebalancing restrictions 344
- redundancy level restrictions 347

ORACLE_HOME

- specifying multiple for ASM support 318

P

- PATH environment variable 425
- path mapping 177
- path root, changing when mounting replicas 175
- performance guidelines, mount 222
- physical target access (mount option) 216
- platforms, supported 17
- PowerShell commands in callout scripts 240
- pre- and post-replication scripts
 - in Exchange 555
 - in NTFS 608
 - in Oracle 390
 - in SQL Server 483, 484
 - in UDB 424
 - in UNIX 601
- preventing automatic replica deletion 300
- privileges for callout scripts 240
- production host setup, Exchange 490
- production mount 182

R

- RAC mount to cluster 334
- RAC to RAC mount 334
- RAID 5 BCVs
 - mixing with TimeFinder/Mirror BCVs 265
 - restrictions when using 229
- raw devices
 - on Linux 405
 - on UNIX 599
- RDM 97
- Real Application Cluster (RAC) 334
- RecoverPoint
 - any point in time replica 82
 - application-consistent replica 82
 - CE management 217
 - consistency groups 113
 - crash-consistent replica 82
 - mount and restore 215
 - mount SQL Server replica 468
 - Oracle replica 387
 - replica management 298
 - replica rotation 304
 - restore considerations 218
 - specific point in time replica 82
 - SQL Server mount restrictions 469
 - SRM management 217
- recovery, of information from a replica 161
- redo logs
 - and consistent split restrictions 315
 - location in Oracle 338
- remote replication 283
- renaming databases in Oracle 365
- renaming SIDs in Oracle 365
- replicas
 - acknowledging failed 310
 - created using SAN Copy 284
 - created using SRDF 284
 - deleting 300
 - disabling expiration 301
 - flagged 310
 - granularity of 121
 - managing 277
 - modifying rotation 302
 - number supported 110
 - objects included 280
 - properties 278
 - reenabling expirations 301
 - restrictions on creating simultaneously 149
 - rotating 303
 - storage devices included 281
 - suspending automatic deletion 300
 - unacknowledging failed 310
 - viewing 278
 - viewing history 282
 - viewing log data 282
- Replication Manager
 - agent, overview of 40
 - architecture 32
 - components 40
 - console, overview of 40
 - installing 40
 - overview 26
 - server, overview of 40
- replications
 - using callout scripts during 239, 242, 243
- repurposing 27
- restarting a scheduled task 153
- restore 161
 - and roll forward phase 30
 - composite application sets 233
 - considerations with RecoverPoint 218
 - database considerations 226
 - from tape, Exchange data 516, 548
 - granularity of 231
 - Microsoft Cluster 234
 - protective measures 225
 - UNIX limitations 599
 - using callout scripts during 239
 - volume-type granularity by storage array 231
- retention period 124
- rm_hacmp_pvidupdt.pl command 196
- rm_serviceguard_vgidupdt.pl command 197
- rmsqlrestore 473
- rotations 124
- rotations, managing replicas in 150
- running jobs on demand 157

S

- SAN Copy
 - clone sync rate (CLARiiON) 289
 - full 74

- incremental 77
 - link utilization 289
 - replica 284
 - session throttle 289
- save pool
 - customized 64
- scheduled task, stopping and restarting 153
- scheduling a job 152
- screen components 43
- scripts
 - callout 235
 - general guidelines 390
 - in a federated environment 267
 - NTFS 608
 - PowerShell commands in 240
 - UNIX 601
- SCSI targets
 - uniqueness across controllers 93
- server software 33, 40
- SharePoint
 - application set, updating 119
 - application sets 576
 - changes to configuration 119
 - configuring the environment 569
 - full farm restore 568, 585, 586, 592
 - mounting replicas 581
 - overview of support for 566
 - prerequisites 567
 - restoring with rmsqlrestore 587
 - troubleshooting 591
- SID
 - renaming in Oracle 365
- simulation 142
- snaps
 - CLARiiON 71
 - VNX 72
- Snapshot Clones
 - restore granularity 231
- Snapshot snaps
 - restore granularity 231
- SnapView Clones
 - CLARiiON 68
 - restore granularity 231
 - VNX 69
- SnapView Snaps
 - CLARiiON 71
 - restore granularity 231
- VNX 72
- sp-file, Oracle 317, 371
- SQL Server
 - application sets 439
 - consistent split 255
 - copy replication 441
 - file restore 475
 - file system mount 467, 468
 - filegroup replication defined 439
 - filegroup restores 475
 - filestream datatype 448
 - full database replication defined 439
 - full replication 441
 - full restores 470
 - mount and restore functions 464
 - mount RecoverPoint replica 468
 - mount recovery modes
 - with advanced recovery 468
 - with limited recovery 466
 - mount restrictions 182
 - no recovery mode 468
 - online replication 443
 - online with advanced recovery (using VDI)
 - replications 443
 - online with limited recovery replication 441, 443
 - pre- and post-scripts 483, 484
 - production mount 466
 - RecoverPoint mount restrictions 469
 - recovery mode 468
 - replicating SQL Server 2005/2008 databases
 - with mirror sessions 448
 - replication types (SQL Server 2005/2008) 441
 - restore to cluster 480
 - restoring with rmsqlrestore 473, 474
 - rules for mounting replicas 184
 - snapshots 481, 482
 - standby mode 468
 - transparent data encryption 469
- SRDF remote BCV replica 61, 284
- SRDF/A 61, 63, 66, 86, 92
- SRDF/S 61, 63, 66, 86, 92
- Standby Continuous Replication (SCR) 525
- status bar 43
- stopping scheduled task 153
- storage
 - device name 159

- state 159
- viewing arrays and devices 158
- visible hosts 159
- storage allocation
 - CLARiiON snaps 71
 - incremental SAN Copy 77
 - RecoverPoint 82
 - Remote BCVs 61
 - TimeFinder/Snaps 63, 65, 66
 - VDEVs 63, 65, 66
 - VNX snaps 72
- storage arrays, supported 17
- storage pools in the console 45
- storage services
 - in console 45
 - viewing 158
- storage types supported 31
- support 23
- supported applications and platforms 17
- surgery phase 30
- Symmetrix clones, restoring from 229
- Symmetrix storage arrays, restore granularity 231
- SYSDBA users, finding a list in Oracle 338

T

- tape backup of Exchange data 514, 547
- tape restore of Exchange data 516, 548
- task panel 43
- tasks
 - canceling 157
 - monitoring progress while running 157
- technical support 23
- TimeFinder
 - clones 58, 60
 - mirror (remote) 61
 - mirror of R2 61
 - restore granularity of snaps 231
 - restoring from clones 229
 - snaps 63, 65, 66
 - snapshot restrictions 228
- timeouts, VSS and Exchange 113
- toolbar 43
- topology view 136
- tree panel 43

U

UDB

- application sets and replicas 414
- collecting information 412
- configuring environment 410
- controlling where temp files are stored 417
- database mount options 419
- list databases 412
- list instances 412
- LOGRETAIN settings 410
- MIRROR mode 423
- mount restrictions 182
- mounting and restoring 418, 419
- offline replications 414, 424
- online replication illustrated 416
- online replications 414
- pre- and postreplication scripts 424, 425
- Prepare Only mount 421
- restores 422
- rules for mounting replicas 184
- sample scripts 425, 426
- script guidelines 425
- show database configuration 412
- show tablespace locations 412
- Snapshot mode 421
- test username and password for an instance 412
- user-supplied scripts 424
- write-suspend mode 426

Universal Database. *See* UDB

UNIX

- cluster mount and restore 193
- pre- and post-replication scripts 601
- rules for mounting replicas 185
- scripts 601

unmounting replicas, using callout scripts while 240

user

- granting application set access 120
- in console 45

V

- validating application sets 117
- vCenter credentials 116
- VDEVs
 - restore granularity 231

- restoring from 228
- restrictions 228
- storage allocation 63, 65, 66
- Veritas volume group maintenance 122
- Virtual I/O (VIO) 107
- virtual target access (mount option) 216
- virtual target with roll (mount option) 216
- VMware 84, 104
- VMware Microsoft iSCSI Initiator
 - on Celerra 103
 - on CLARiiON or VNX 103
- VMware Raw Device Mapping 97
- VMware Raw Disk Mapping
 - on Celerra 99
 - on CLARiiON or VNX 99
 - on Symmetrix 100
- VMware vCenter 116
- VMware virtual disks
 - on Celerra 95
 - on CLARiiON or VNX 95
 - on Symmetrix 95
- VMware VirtualCenter 116
- VMware VMFS
 - and federated application sets 261
 - on Celerra 88
 - on CLARiiON or VNX 89
 - supported environments 84, 261
- VNX
 - clones 69, 71
 - restore granularity 231
 - SAN Copy options 289
 - snaps 72, 621
- VNX Replicator
 - overview 80
 - remote snapshot 80
- VNXe
 - IQNs 158
- volume groups, shared by multiple applications 121
- VSS
 - restoring Exchange 2003/2007 using 541
 - restoring Exchange 2007/2013 using 508
 - timeout 113

X

- XML generated for callout scripts 241