

EMC[®] SRDF/Cluster Enabler Plug-in

Version 4.1.4

Product Guide

P/N 300-014-898
REV 01

Copyright © 2012 EMC Corporation. All rights reserved. Published in the USA.

Published December, 2012

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 the technical documentation and advisories section on the EMC online support website.

CONTENTS

Preface

Chapter 1

About Cluster Enabler

Cluster Enabler overview	16
Cluster Enabler plug-in architecture.....	17
Cluster Enabler components	18
Cluster Enabler Manager interface	19
The Cluster Enabler Manager window	19
The Cluster Enabler Manager wizards	19
Cluster Enabler logging	20
Disk space requirements.....	21
Extracting logs	21
Changing the logging level	21
Changing the logging directory	22
Changing logging retention period	22
Changing the maximum log file size	22
Windows event log messages.....	23
Microsoft Windows Server support.....	24
Quorum model support	24
Multiple CE cluster management.....	25
Setting Up devices on Windows Server 2008 or 2012	25
Virtualization support	26
Hyper-V support	27
Cluster Shared Volumes	28
VMware support.....	31
Supported functionality	33
Delay Failback.....	33
Mount point support	34
Multiple storage array support	35
Delegating CE administration	35
Viewing cluster dependency.....	37
SRDF/CE support matrix	40

Chapter 2

About SRDF/Cluster Enabler

SRDF/Cluster Enabler plug-in overview	42
SRDF overview	43
SRDF/CE supported features	45
SRDF/Asynchronous compatibility	45
SRDF/CE swap support.....	46
Symmetrix Virtual Provisioning.....	47
Supported devices	47
SRDF/CE configuration with multiple remote adapters.....	47
Monitoring SRDF link status	47
SRDF composite groups.....	48
Concurrent SRDF	48
Restrictions and limitations.....	49
Failover/Failback behavior	49
Cascaded SRDF	51
Cascaded SRDF/CE requirements	51
Restrictions and limitations.....	52
Failover/Failback behavior	52
Configuring cascaded SRDF with CE Manager	53
Pre-SRDF/CE clustering considerations	54

Chapter 3	Clustering Concepts	
	Microsoft Failover Clusters	58
	Microsoft Failover Cluster concepts	59
	Microsoft Failover Cluster modes of operation	60
	CE geographic cluster system	61
	Cluster Enabler modes of operation.....	62
	Cluster behavior and failover operations	64
	Application software in a cluster environment	64
Chapter 4	Cluster Behavior	
	Cluster failover operation	66
	SRDF/CE failover and recovery behavior	67
	SRDF/CE unique behavior	68
	Complete site failure and recovery	69
	Response to complete site failure.....	70
	Failure behavior when using MNS with File Share Witness	73
Chapter 5	SRDF/CE Installation	
	Installation overview	76
	Before you begin	76
	Getting started with Symmetrix arrays	78
	Installing the SRDF/CE plug-in module	79
	Uninstalling the SRDF/CE plug-in module	79
	Uninstalling the plug-in from some cluster nodes.....	79
	Uninstalling the plug-in from all cluster nodes/deconfigure the cluster .	80
	Uninstalling the plug-in from all cluster nodes/destroy the cluster	80
Chapter 6	Using Cluster Enabler Manager	
	Getting started using the CE Manager.....	82
	The Cluster Enabler Manager window	82
	Cluster Enabler wizards	83
	Using the CE Configuration Wizard	84
	Adding nodes.....	86
	Managing a CE cluster.....	87
	Storage Discover Wizard.....	87
	Update Mirrored Pairs Wizard	88
	Change Quorum Model Wizard	88
	Managing a CE cluster group	91
	Create Group Wizard	91
	Modify Group Wizard	93
	Configure a CE Group	95
	Deconfigure a CE group	96
	Delete a CE group.....	96
	Storage component.....	97
	Adding and removing devices from a group	98
	Viewing information	98
	Displaying group information	99
	Displaying node information	102
	Displaying site information	104
	Restore and recovery operations	106
	SRDF/CE recovery procedures	106
	Configuring a custom resource	111

CONTENTS

Using CE Manager to create a custom resource CE Group..... 113
Using CE Manager to edit a custom resource CE Group 115

Appendix A

Base Component Installation and Upgrade

Installation overview 120
Before you begin 120
Installing the Cluster Enabler Base Component..... 121
 Installing the Base Component separate from the plug-ins
 (clean install) 121
 Installing the Base Component along with the plug-ins (clean install) . 121
 Upgrading the Base Component along with the plug-ins..... 121
 Upgrading only the Base Component 122
Uninstalling the Cluster Enabler Base Component 123
 Uninstalling the Base component from some cluster nodes..... 123
 Uninstalling the base component from all cluster nodes/deconfigure
 the cluster..... 123
 Uninstalling the base component from all cluster nodes/destroy the
 cluster..... 124
Configuring a CE cluster on Server Core 124
 Requirements and considerations 124
 R2 Server Core configuration 124
Upgrading Windows Server 2008 to Windows Server 2008 R2 125

Glossary

FIGURES

	Title	Page
1	Overview example of a typical CE cluster configuration.....	17
2	Cluster Enabler Manager window	19
3	CE Manager with virtual machine cluster group	28
4	Cluster Shared Volumes tree view	30
5	Lateral and peer nodes.....	33
6	Sample Dependency Report	39
7	Overview example of an SRDF/CE cluster configuration	43
8	Basic SRDF configuration	44
9	SRDF/CE with concurrent SRDF.....	49
10	Sample SRDF Cascaded configuration	51
11	Recommended cabling configuration	54
12	Typical two-node Microsoft Failover Cluster	58
13	A typical four-node Microsoft Failover Cluster.....	59
14	A geographically distributed two-node CE cluster	61
15	A geographically distributed four-node CE cluster.....	62
16	Two-node two-cluster CE configuration	63
17	SRDF/Cluster Enabler failover operation.....	66
18	Types of complete site failure.....	69
19	Lateral and peer nodes.....	71
20	MNS clusters with File Share Witness	73
21	Cluster Enabler Manager window	82
22	CE Manager Configuration Wizard	84
23	CE Manager expanded navigation tree.....	86
24	Windows Server 2008 supported quorum models	89
25	Create Group Wizard, Select Devices for the Group	92
26	Create Group Wizard, Select Group Policy.....	93
27	Modify Group Wizard, Select Devices	94
28	Configure CE Group option	96
29	Example of Symmetrix storage array view	97
30	CE Manager storage actions	98
31	CE Manager Groups component	99
32	CE Manager groups information	99
33	CE Manager VM groups information.....	101
34	CE Manager Nodes component.....	102
35	CE Manager node information	102
36	CE Manager Sites component.....	104
37	CE Manager Symmetrix site information	105
38	Recover CE Cluster Enter Node Name.....	110
39	Recover CE Cluster Choose Tasks	110
40	Recover CE Cluster Change Cluster Number	111
41	Microsoft Cluster Administrator, Generic Application Resource Type.....	112
42	Cluster properties	112
43	Cluster properties with Generic Application	113
44	Select Group Policy, custom resource.....	114
45	Microsoft Cluster Administrator, EMC_Group 4	115
46	Validate selection, custom resource	116
47	Summary of Group 4, custom resource.....	117

TABLES

	Title	Page
1	Cluster Enabler Manager Wizards	20
2	Windows event log messages.....	23
3	CSV feature support matrix.....	31
4	Cluster mount point example	35
5	SRDF/CE support matrix	40
6	Cluster Enabler Manager wizards	45
7	Storage component display information.....	97
8	Groups component displayed information.....	100
9	Groups event information.....	100
10	CSV Group information.....	101
11	Nodes component displayed information.....	103
12	Node event information.....	103
13	Site component displayed information.....	105

PREFACE

As part of an effort to improve its product lines, 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 representative if a product does not function properly or does not function as described in this document.

Note: This document was accurate at publication time. New versions of this document might be released on the EMC online support website. Check the EMC online support website to ensure that you are using the latest version of this document.

Purpose

This guide is part of the EMC Cluster Enabler for Microsoft Failover Clusters documentation set and is intended for use by system administrators during installation, system setup, and routine operations.

Audience

System administrators working with Cluster Enabler must be proficient in the use of the following products:

- ◆ Microsoft products:
- ◆ Windows Server 2008 or 2012 Enterprise and Datacenter Editions, as installed
- ◆ Windows Server 2008 R2 or Windows Server 2008 R2 Server Core, or Windows Server 2012 Server Core Enterprise and Datacenter Editions, as installed.
- ◆ Microsoft Failover Clusters

EMC Symmetrix storage arrays, as per your Cluster Enabler product version and the following applicable software:

- ◆ Solutions Enabler (SYMCLI/SYMAPI)
- ◆ EMC Symmetrix Remote Data Facility (SRDF)
- ◆ EMC ControlCenter Symmetrix Remote Data Facility (SRDF) Manager, if installed
- ◆ EMC PowerPath, if installed

Required documentation

The following documentation is part of the EMC Cluster Enabler for Microsoft Failover Clusters documentation set, and is required for SRDF/Cluster Enabler:

- ◆ *EMC Cluster Enabler Base Component Release Notes*
- ◆ *EMC SRDF/Cluster Enabler Plug-in Release Notes*
- ◆ *EMC SRDF/Cluster Enabler Plug-in Product Guide*

Related third-party documentation

The following Microsoft documentation available at microsoft.com contains information about or related to the products discussed in this guide:

- ◆ *Windows Server 2008 or 2012 Clustering Whitepapers*, containing various whitepapers and datasheets overviewing Windows Server 2008 or 2012 Clustering.

Related documentation

The following documentation from EMC Corporation contains information that may be helpful in a Cluster Enabler environment.

EMC Solutions Enabler:

- ◆ *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*
- ◆ *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide*
- ◆ *EMC Solutions Enabler Symmetrix TimeFinder Family CLI Product Guide*
- ◆ *EMC Solutions Enabler Installation Guide*

EMC ControlCenter:

- ◆ *EMC ControlCenter Planning and Installation Guide*
- ◆ *Symmetrix SRDF Host Component Product Guide*

EMC PowerPath:

- ◆ *EMC PowerPath Product Guide*

Fibre Channel:

- ◆ *Symmetrix Fibre Channel Product Guide*

For detailed interoperability information, please refer to E-Lab Interoperability Navigator, which can be reached at <http://elabnavigator.EMC.com>.

Conventions used in this document

EMC uses the following conventions for special notices:



CAUTION, used with the safety alert symbol, indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.

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

IMPORTANT

An important notice contains information essential to software or hardware operation.

Typographical conventions

EMC uses the following type style conventions in this document:

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, functions, and utilities URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, and notifications
Bold	Used in running (nonprocedural) text for names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, and man pages Used in procedures for: <ul style="list-style-type: none"> Names of interface elements, such as names of windows, dialog boxes, buttons, fields, and menus What the user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> Full titles of publications referenced in text Emphasis, for example, a new term Variables
Courier	Used for: <ul style="list-style-type: none"> System output, such as an error message or script URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
Courier bold	Used for specific user input, such as commands
<i>Courier italic</i>	Used in procedures for: <ul style="list-style-type: none"> Variables on the command line User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user
[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections — the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Where to get help

EMC support, product, and licensing information can be obtained on EMC Online Support, as described next.

Note: To open a service request through EMC Online Support, you must have a valid support agreement. Contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

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

EMC offers a variety of support options.

Support by Product — EMC offers consolidated, product-specific information on the Web at:

<https://support.EMC.com/products>

The Support by Product web pages offer quick links to Documentation, White Papers, Advisories (such as frequently used Knowledgebase articles), and Downloads, as well as more dynamic content, such as presentations, discussion, relevant Customer Support Forum entries, and a link to EMC Live Chat.

EMC Live Chat — Open a Chat or instant message session with an EMC Support Engineer.

eLicensing support

To activate your entitlements and obtain your Symmetrix license files, visit the Service Center on <https://support.EMC.com>, as directed on your License Authorization Code (LAC) letter emailed to you.

For help with missing or incorrect entitlements after activation (that is, expected functionality remains unavailable because it is not licensed), contact your EMC Account Representative or Authorized Reseller.

For help with any errors applying license files through Solutions Enabler, contact the EMC Customer Support Center.

If you are missing a LAC letter, or require further instructions on activating your licenses through the Online Support site, contact EMC's worldwide Licensing team at licensing@emc.com or call:

- ◆ North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.
- ◆ EMEA: +353 (0) 21 4879862 and follow the voice prompts.

CHAPTER 1

About Cluster Enabler

This chapter provides a high-level overview of clustering and explains how EMC Cluster Enabler provides disaster-recovery protection in geographically distributed Microsoft Failover Clusters.

IMPORTANT

EMC recommends reading this chapter in its entirety before installing and configuring Cluster Enabler for Microsoft Failover Clusters.

◆ Cluster Enabler overview	16
◆ Cluster Enabler Manager interface	19
◆ Cluster Enabler logging	20
◆ Microsoft Windows Server support.....	24
◆ Virtualization support	26
◆ Supported functionality	33

Cluster Enabler overview

Cluster Enabler (CE) for Microsoft Failover Clusters is a software extension of failover clusters functionality. Cluster Enabler allows Windows Server 2008 (including R2) and Windows Server 2012 Enterprise and Datacenter editions running Microsoft Failover Clusters to operate across multiple connected storage arrays in geographically distributed clusters. Each cluster node is connected through a storage network to the supported storage arrays. The method of automatic failover for mirrored pairs during a node failure depends on the storage environment.

CE software supports the following replication technologies:

- ◆ SRDF[®]/Cluster Enabler for Microsoft Failover Clusters (for Symmetrix[®] storage arrays)
- ◆ MirrorView[™]/Cluster Enabler for Microsoft Failover Clusters (supported up to version 4.1 for CLARiiON[®] storage arrays)
- ◆ RecoverPoint[™]/Cluster Enabler for Microsoft Failover Clusters (for storage arrays supported by RecoverPoint)

Note: Refer to EMC Online support for Cluster Enabler plug-in software module availability for your replication technology, or check with your EMC sales representative.

Once configured using the EMC Cluster Enabler Manager graphic user interface (GUI), Microsoft Failover Clusters are referred to as *CE clusters*.

Cluster Enabler expands the range of cluster storage and management capabilities while ensuring full business-continuance protection. An iSCSI or Fibre Channel connection from each cluster node is made to its own storage array. Two connected storage arrays provide automatic failover of mirrored volumes during a Microsoft Failover Cluster node failover.

This connection effectively extends the distance between cluster nodes (depending on network latency) and forms a geographically distributed cluster (stretch cluster) with disaster-tolerant capabilities.¹

[Figure 1](#) provides an example of a typical Cluster Enabler configuration. There are two hardware sites. Primary Site A has a storage array connected to Microsoft Cluster Servers, and Secondary Site B has a storage array connected to another set of Microsoft Cluster Servers. The Microsoft Cluster Servers are connected by means of a Local Area Network (LAN) connection, and the storage arrays are connected by way of the storage array's links.

1. The *EMC Networked Storage Topology Guide* provides additional information regarding distance restrictions for your specific configuration.

Cluster Enabler protects data from storage, system, and site failures, 24 hour a day, 7 days a week, and 365 days per year:

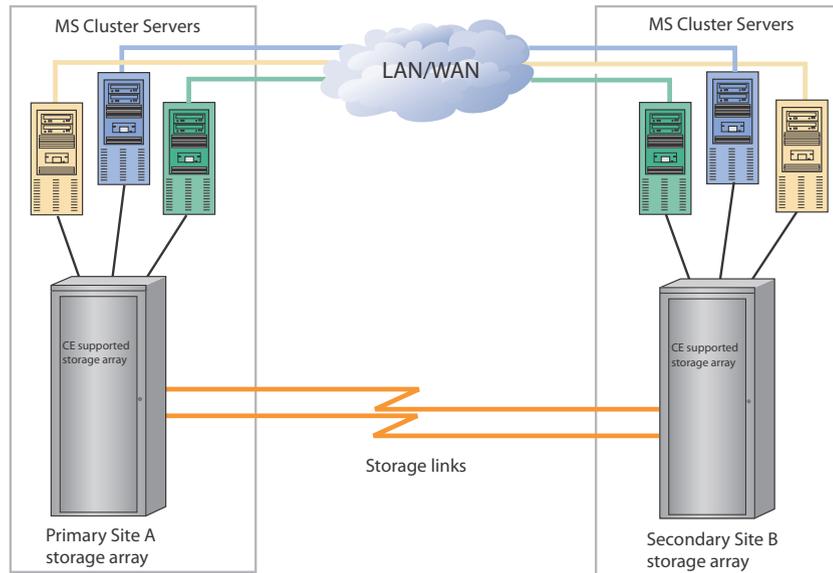


Figure 1 Overview example of a typical CE cluster configuration

Cluster Enabler plug-in architecture

EMC Cluster Enabler for Microsoft Failover Clusters provides a plug-in architecture consisting of a CE base module and separately available plug-in modules. Each CE plug-in module supports a different storage replication technology as follows:

- ◆ SRDF/Cluster Enabler for Microsoft Failover Clusters (for Symmetrix storage arrays)
- ◆ MirrorView/Cluster Enabler for Microsoft Failover Clusters (supported up to version 4.1 for CLARiiON storage arrays)
- ◆ RecoverPoint/Cluster Enabler for Microsoft Failover Clusters (for multiple RecoverPoint supported storage arrays)

Select your plug-in module based on your storage environment's requirements. The new Cluster Enabler architecture supports the coexistence of multiple plug-ins, which can be installed on the same cluster node.

Note: You cannot mix replication technologies and storage configurations within the same cluster group. For example, Symmetrix SRDF and CLARiiON MirrorView devices cannot be part of the same CE cluster group.

The base module must be installed prior to installing a plug-in module. ([“Appendix A”](#) provides detailed installation instructions for the CE Base Component.)

Note: Refer to EMC Online support for CE plug-in software module availability for your replication technology, or check with your EMC sales representative.

Cluster Enabler components

Cluster Enabler integrates Microsoft Failover Cluster software with replication technology software and supported storage hardware, allowing the seamless use of disks to function as a single SCSI disk. Cluster Enabler achieves this using several components:

- ◆ **CE Manager**—An MMC-based (Microsoft Management Console) user interface that allows you to configure operational parameters and perform cluster tasks.
- ◆ **CE Resource DLL**—A Dynamic Link Library (DLL) that is used by Microsoft Failover Cluster to perform group failover/failback operations for all storage group resources.
- ◆ **CE VM Resource DLL**—A Dynamic Link Library that is used by Microsoft Failover Cluster to perform failover/failback of Hyper-V child partitions residing on Cluster Shared Volumes (CSVs).
- ◆ **CE WMI provider**—A Windows Management Instrumentation component that interfaces with the underlying storage array and performs various operations, such as failover, group creation, and so on, on the storage array.
- ◆ **CE Service**—A Windows service dependent on Cluster Service, used for Quorum and CSV Device Failover, and to manage the Preferred Owners' list.
- ◆ **Quorum Filter Driver**—A component that performs arbitration or *ownership protocol* for the Microsoft Failover Cluster database quorum.

Cluster Enabler documentation

EMC Cluster Enabler product documentation consists of an integrated online help system and the following documents:

EMC Cluster Enabler Base Component:

- ◆ EMC Cluster Enabler Base Component Release Notes

EMC SRDF/Cluster Enabler Plug-in:

- ◆ EMC SRDF/Cluster Enabler Plug-in Release Notes
- ◆ EMC SRDF/Cluster Enabler Plug-in Product Guide

EMC MirrorView/Cluster Enabler Plug-in:

- ◆ EMC MirrorView/Cluster Enabler Plug-in Release Notes
- ◆ EMC MirrorView/Cluster Enabler Plug-in Product Guide

EMC RecoverPoint/Cluster Enabler Plug-in:

- ◆ EMC RecoverPoint/Cluster Enabler Plug-in Release Notes
- ◆ EMC RecoverPoint/Cluster Enabler Plug-in Product Guide

Note: Additional related documentation is provided with each replication technology plug-in module.

Cluster Enabler Manager interface

Cluster Enabler for Microsoft Failover Clusters provides a graphic user interface called Cluster Enabler Manager. The CE Manager provides several wizards to streamline cluster tasks and reduce the complexity of typical cluster management.

The CE Manager allows you to configure your Microsoft Failover Clusters for disaster-recovery protection. The CE Manager allows you to set up and configure disk-based resources to automatically move geographically dispersed resource groups back and forth.

The Cluster Enabler Manager window

The CE Manager window, shown in [Figure 2](#), contains a menu bar, two views, and a navigation tree. After cluster configuration, the navigation tree can be expanded to show four separate components: Groups, Storage, Sites, and Nodes.

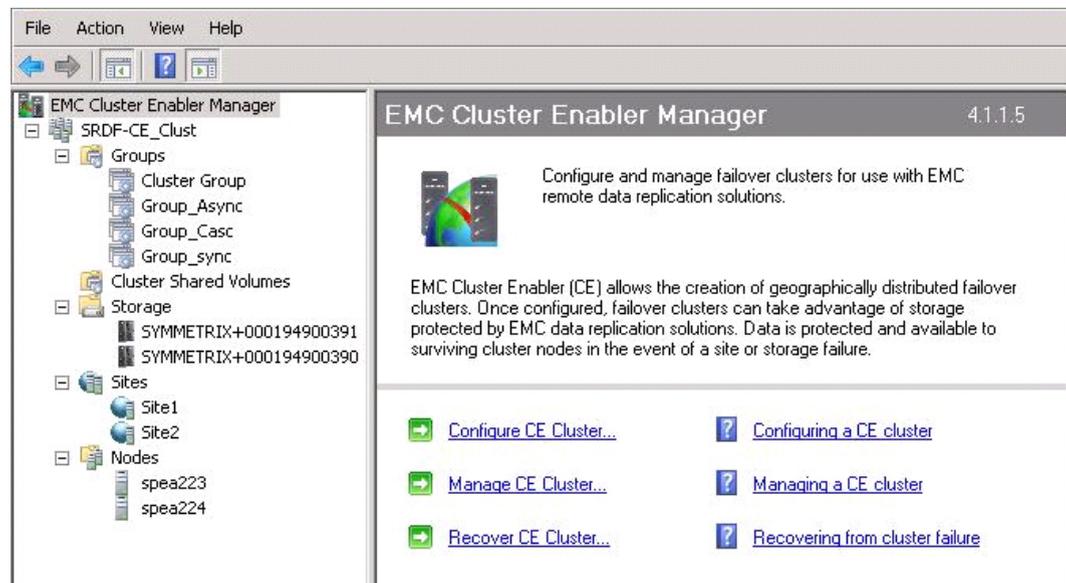


Figure 2 Cluster Enabler Manager window

The Cluster Enabler Manager wizards

The CE Manager provides several wizards to assist you in completing various cluster tasks. Wizards are a series of dialog boxes that step you through the completion of a complex task. The first step towards managing disaster recovery for distributed failover clusters is to run the Configuration Wizard to configure a CE cluster.

[Table 1](#) lists the various wizards that are included in the CE Manager.

Table 1 Cluster Enabler Manager Wizards

Wizard	Functionality
Configuration Wizard	Configures a CE cluster. The configuration process is the first step towards managing disaster recovery for distributed failover clusters. The Configuration Wizard steps you through the process of configuring your failover cluster for management with CE.
Create Group Wizard	Creates a CE Group, add devices, and select a group policy.
Modify Group Wizard	Steps you through the process of adding or removing devices in a CE group.
Recover CE Cluster Wizard	Recover a supported cluster.
Change Quorum Wizard	Changes a cluster's quorum model type.
Update Mirror Pairs Wizard	Discovers storage, updates the storage configuration, validates the storage groups, and sets up the storage group definitions in the cluster properties database to update the mirrored pairs in a cluster.
Storage Discovery Wizard	Discovers and sets up the attached storage. Performs a storage discovery after any changes to the storage configuration.

Cluster Enabler logging

Cluster Enabler provides detailed logging features and implements a simplified extraction process for extracting log file data. If there is a problem with Cluster Enabler, detailed logs provide EMC Customer Support with the technical information necessary to help diagnose the problem and help Cluster Enabler engineers with debugging.

Cluster Enabler incorporates various logging capabilities to create application detail logs. The amount of detail that these logs contain is controlled by the logging level. You can adjust the logging level to suit your needs. Refer to [“Changing the logging level” on page 21](#) for more information. Under normal operating conditions, error, warning, and information entries are written to the application detail log. When verbose logging is enabled, these logs contain enough information to help developers diagnose various application failures.

By default, logs are stored in the `C:\Program Files\EMC\Cluster-Enabler\Logs` directory. The latest log file is named `ce_event_trace_current.txt`.

The stored logs are saved as text files and can be viewed using any text editor. Note that the current log file is an active file and therefore may not contain a complete set of log entries, as some may still be in process. Some text editors may not be able to access the current log file. To obtain a complete copy of the current log file, you can use the `CE_EventTraceDump.exe` program. [“Extracting logs” on page 21](#) provides more detail and some extraction examples.

By default, when the log file exceeds 100 MB in size, it will be closed out and renamed from `ce_event_trace_current.txt` to `ce_event_trace_yyyymmddhhmmss.txt`, where `yyymmddhhmmss` is the current date and time. The maximum file size of the log is controlled by a registry key entry and can be changed. [“Changing the maximum log file size” on page 22](#) provides more information.

To help manage logging disk space, older log files are automatically purged. By default, 7 logs are saved. You can control how many logs are saved. [“Changing logging retention period” on page 22](#) provides additional information.

Disk space requirements

The amount of disk space required depends on the logging level and the amount of cluster activity taking place. As a general guide, you might expect 50 KB per day for a normal logging level. If the logging level is set to verbose, and cluster activity is greater than normal, you might expect 200 MB or more per day.

Extracting logs

To extract a current log file, enter type `CE_EventTraceDump.exe` from the command line of the working directory. This extracts the dump file to the designated log directory and names it `ce_event_trace_yyyymmddhhmmss.txt`, where `yyymmddhhmmss` is the current date and time. You can use the `-o file name` option to change the name of the output file.

Examples Each of the following examples assume that the current working directory is `C:\Program Files\EMC\Cluster-Enabler`, and that the default log directory is `C:\Program Files\EMC\Cluster-Enabler\Logs`.

Example 1

To extract the dump file to the log directory and name it `ce_event_trace_yyyymmddhhmmss.txt`, enter:

```
CE_EventTraceDump.exe
```

Example 2

To extract a named dump file to a particular location, use the following format:

```
CE_EventTraceDump.exe -o C:\filename.txt
```

Note: Use the `-h` option to display usage information.

Changing the logging level

The logging level is controlled by a registry key. To change the logging level, follow these steps:

1. Open a command prompt and enter:


```
regedit
```
2. Edit the registry key value for:


```
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CE\Configuration\
EventTraceLevel
```
3. Using the Microsoft Service Management API, restart the `ce_eventtrace` service.

By default, the level is set to 4. At this level, error, warning, and informational messages appear in the log file. To create verbose logs, you can change the value to 5. At this level, error, warning, informational, and verbose messages are sent to the log file. Be aware that changing this level to 5 dramatically increases the amount of data that is sent to the log file.

Changing the logging directory

The logging directory is controlled by a registry key. To change the logging directory, follow these steps:

1. Open a command prompt and enter:

```
net stop ce_eventtrace
```

2. Then enter:

```
regedit
```

3. Edit the registry key value for:

```
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CE\Configuration\EventTraceDirectory
```

Your edited path must have a trailing backslash (\) and must exist before you make this change.

4. Then enter:

```
net start ce_eventtrace
```

Changing logging retention period

The log retention period is controlled by a registry key. To change the log retention period, follow these steps:

1. Open a command prompt and enter:

```
regedit
```

2. Edit the registry key value for:

```
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CE\Configuration\EventTraceLogRetention
```

The `DWORD` value represents the number of logs to keep. The `ce_eventtrace` service does not need to be restarted. The new value takes effect almost immediately.

Changing the maximum log file size

The maximum log file size is controlled by a registry key. To change the maximum log file size, follow these steps:

1. Open a command prompt and enter:

```
regedit
```

2. Edit the registry key value for:

```
HKEY_LOCAL_MACHINE\SOFTWARE\EMC\CE\Configuration\EventTraceFileMaxSize
```

The `DWORD` value represents the file size in MBs. The `ce_eventtrace` service does not need to be restarted. The new value will take effect almost immediately.

Windows event log messages

The Windows event log displays descriptive event messages for some of the more common events encountered when using Cluster Enabler. [Table 2](#) lists the event log messages by Event ID, Event type, Description, and the Action that should be taken when the specific event has been encountered.

Note: Event types are described as an error, warning, or information.

Table 2 Windows event log messages (page 1 of 2)

Event ID	Event type	Description	Action
1	Informational	Generic ID used to report informational messages.	Action varies based on description text.
2	Warning	Generic ID used to report warning messages.	Action varies based on description text.
3	Error	Generic ID used to report error messages.	Action varies based on description text.
4	Informational	Informational message generated when a group comes online successfully.	No action necessary.
5	Error	Error message generated when a group fails to come online.	The description text indicates the name of the group that failed to come online. Look at the previous event log messages and application logs to find the root cause of the failure.
6	Error	An unexpected application error occurred.	<ol style="list-style-type: none"> 1. Attempt the action again. 2. Turn on verbose logging (logging level 5) and attempt again. 3. If failure occurs again, save the Windows event log and the CE application log, and contact EMC support.
7	Error	The link between the storage arrays is down for storage group (<i>GroupName</i>).	Use storage array CLI interfaces to determine the root cause of the problem.
8	Informational	The link between the storage arrays is replicating data to the remote storage array.	No action necessary.
9	Error	Communication or data access to the WMI (Windows Management Instrumentation component) service failed.	<ol style="list-style-type: none"> 1. Read the event log messages and application logs to find the root cause of the problem. 2. If failure occurs again, save the Windows event log and the CE application log, and contact EMC support.
10	Error	A failure occurred while reading or writing storage group information.	<ol style="list-style-type: none"> 1. Attempt the action again. 2. Turn on verbose logging (logging level 5) and attempt again. 3. If failure occurs again, save the Windows event log and the CE application log, and contact EMC support.

Table 2 Windows event log messages (page 2 of 2)

Event ID	Event type	Description	Action
11	Error	A failure occurred while reading or writing storage group information to the cluster registry.	<ol style="list-style-type: none"> 1. Attempt the action again. 2. Turn on verbose logging (logging level 5) and attempt again. 3. If failure occurs again, save the Windows event log and the CE application log, and contact EMC support.
12	Error	A failure occurred while deleting a mirror group.	Read the event log messages and application logs to find the root cause of the problem.
13	Error	A failure occurred while creating a mirror group.	Read the event log messages and application logs to find the root cause of the problem.

Microsoft Windows Server support

EMC Cluster Enabler for Microsoft Failover Clusters is supported on Microsoft Windows Server 2008 and 2012 systems, including Windows Server 2008 R2 and Core editions.

There are two Windows processor architectures that are supported:

- x86 (Windows Server 2008)
- x64 (AMD64 and Intel EM64T)

Note: Microsoft does not support mixed architecture clusters. All nodes must be the same Windows architecture.

[“Appendix A”](#) provides installation instructions and lists the prerequisites and requirements for supported Microsoft Windows Servers.

Quorum model support

Quorum model support and changing the cluster model type of a cluster depends on your chosen CE plug-in module and storage replication technology. Cluster Enabler provides a wizard for changing the quorum model of a cluster. The Change Quorum Wizard steps you through the process of changing a cluster’s quorum model type. [“Supported model type descriptions”](#) provides detailed descriptions of each model.

Supported model type descriptions

The following supported model types and quorum options exist for Windows Servers 2008 and 2012:

- ◆ No Majority: Disk Only
- ◆ Node Majority
- ◆ Node and Disk Majority
- ◆ Node and File Share Majority

No Majority: Disk Only — This quorum model can sustain failures of all nodes except one (if the disk is online). In the event of a quorum disk failure using this quorum model, the entire cluster would shut down if the quorum disk is lost.

Node Majority — This cluster model is recommended for clusters with an odd number of nodes. The cluster can sustain failures of half the nodes (rounding up) minus one.

Note: If no majority exists, the cluster service is disabled.

Node and Disk Majority — This cluster model is recommended for clusters with an even number of nodes. It can sustain failures of half the nodes (rounding up) if the witness disk remains online. For example, a six node cluster in which the witness disk is online could sustain three node failures. It can sustain failures of half the nodes (rounding up) minus one if the witness disk goes offline or fails. For example, a six node cluster with a failed witness disk could sustain two (3-1=2) node failures.

Node and File Share Majority — This cluster model is recommended for clusters with special configurations. It works in a similar way to Node and Disk Majority, but instead of a witness disk, this cluster uses a witness file share. Note that if you use Node and File Share Majority, at least one of the available cluster nodes must contain a current copy of the cluster configuration before you can start the cluster. Otherwise, you must force the starting of the cluster through a particular node.

Using the Change Quorum Wizard

Once your Microsoft cluster has been configured as a CE cluster, you must use this wizard for all quorum model changes. If your configured CE clusters are No Majority: Disk Only model type, you can use this wizard to change the selected quorum disk. You can also use this wizard to change the file share for configured CE clusters of Node and File Share Majority model type.

Note: To change the quorum model to Node and File Share Majority in Windows Server 2008, you must first update the FileShare permissions to add the Cluster Name and allow Change and Read permissions for the file share. Your Windows documentation provides instructions on changing permissions for FileShare.

Multiple CE cluster management

The Cluster Enabler CE Manager lets you manage multiple CE clusters simultaneously, as long as all of the clusters are either Windows Server 2008 or 2012 clusters and are in the same domain. To manage the cluster, CE Manager runs under a domain administrator account. This account is part of local administrator group of every node of the cluster it manages.

Note: Mixing both Windows Server 2008 and 2012 clusters in one CE Manager session is not supported.

Setting Up devices on Windows Server 2008 or 2012

For Windows Server 2008 or 2012, all disks must first be added to Failover Cluster Management before they can be configured for Cluster Enabler. By default, Failover Cluster assigns all disks to a group called *Available Storage*. You must ensure that Failover Cluster can bring these disks online before using them in Cluster Enabler.

Follow these steps to correctly set up devices on the Windows Server:

1. Choose the appropriate instructions from the following three scenarios, as listed for disks shown in Available Storage:
 - a. If there are no disks in Available Storage, ensure that all disks to be added are write-enabled on the same site (for example, site A).
 - b. If there are already disks in Available Storage, and you want to add more disks, ensure that all disks to be added are write-enabled on the same site where Available Storage is online.
 - c. If some existing disks in Available Storage are not online, move them to the site where the Available Storage is online. If this does not solve the problem, then you need to do the following:
 - Remove those disks from Available Storage.
 - Move all groups and devices to the same node in Failover Cluster. Manually move the corresponding devices to ensure that devices are write-enabled on the node to which you are moving the group.
 - Evict all remaining peer nodes.
2. Ensure that you have access to the disks where they are write-enabled. If not, you must reboot and reformat them.
3. Right-click **Storage** in Failover Cluster Management, and select **Add a Disk**. All available disks will display. You can select disks to add to the cluster. All added disks will be in the group Available Storage. Verify that all disks are online in Available Storage.
4. The devices should now be available for use in Cluster Enabler.

Virtualization support

CE version 4.0 and higher supports the following virtualization tools and features:

- ◆ Windows Server 2008 (x64) Hyper-V
- ◆ Windows Server 2008 R2 (x64) Hyper-V including R2 Server
- ◆ Windows Server 2012 Hyper-V Server
- ◆ VMWare ESX Servers

Windows Server 2008 R2 Cluster Shared Volumes, Windows Server 2012, and Windows Server 2008 (x64) Hyper-V server virtualization is supported for Symmetrix and CLARiiON arrays. Once configured as a CE group using the CE Configuration Wizard, groups with Hyper-V resources display as regular device groups. Windows Server 2008 R2 and Windows Server 2012 Cluster Shared Volumes (CSV) are supported. CSV is a Failover Clustering feature that allows all nodes in a cluster concurrent access to disk on every CSV-enabled shared disk. Once converted using the CE Configuration wizard, CSV disks display under Cluster Shared Volumes in the left pane navigation tree of the CE Manager. Using Cluster Enabler, you can view the properties or change the failover policy of a CSV disk.

Note: The virtual machine and the CSV disks must first be configured in Microsoft Failover Cluster Manager.

Hyper-V support

CE supports Windows Server 2008 (including R2) and Windows Server 2012 Hyper-V server virtualization. Hyper-V is installed and managed as a role under Windows Server 2008 and requires an x64-based processor. SRDF/CE support for Hyper-V is limited to configurations employing *Host Clustering*. Host clustering allows you to host and failover virtual machines between nodes or sites, thereby making them highly available. Once configured using the CE Configuration Wizard, groups with Hyper-V resources display as regular device groups.

The following descriptions explain the difference between *host* clustering and *guest* clustering:

Host Clustering — With *host* clustering, the physical host is the cluster node. If the host stops running, all of its guests (virtual machines) are restarted on another physical host. Host clustering protects against the failure of a physical host (hardware failure of a computer).

Guest Clustering — With *guest* clustering, a guest (Virtual Machine) is a cluster node, and therefore the guest runs applications that are monitored in some way by the Cluster service, either because they are designed to work with clustering (cluster-aware) or because they are configured in the cluster as a Generic Service, Generic Application, or Generic Script resource. With guest clustering, if either the guest operating system or the clustered application fails, the guest can fail over to another guest, either on the same host or on a different host. Guest clustering protects against failure of a cluster-aware application on a guest, as well as failure of an individual instance of a guest.

Note: In Windows Server 2008 R2, guest clustering is only supported using iSCSI disks.

The following listed Microsoft documentation should be consulted for Hyper-V configuration instructions:

- ◆ The *Hyper-V Getting Started Guide* is available at:

<http://technet.microsoft.com>

- ◆ The *Virtualization with Hyper-V: FAQ* is available at:

<http://www.microsoft.com/windowsserver2008>

The following steps are provided as a guide for getting started with Hyper-V and CE version 4.0 and higher for a non-CSV disk:

1. Follow the instructions provided in Microsoft's *Hyper-V Getting Started Guide* to install Hyper-V using the Server Manager.
2. Follow the instructions provided in Microsoft's *Hyper-V Getting Started Guide* to create and set up a virtual machine (guest machine) using the Hyper-V Manager.
3. Install an operating system on the virtual machine.
4. Install the application that you want to be highly available on the operating system of the virtual machine.
5. Using Microsoft Failover Cluster Manager, configure a failover cluster for the virtual machine resources that you just created. Consult your Microsoft Failover Cluster documentation for instructions.

Note: Turn off the virtual machine before adding it to the cluster.

6. Bring the virtual machines online in Failover Cluster Management.
7. Open the CE Manager and configure a CE cluster using the CE Configuration Wizard.
8. On the Current Nodes wizard page, add a second node to the cluster.
9. Once added, follow the steps in the wizard accepting the default settings.
10. Once the CE cluster is configured, note that the CE resource is part of each virtual machine service group. The physical device where the virtual machine was created is dependent on the CE resource. The CE group with the Hyper-V resource displays as a regular device group. [Figure 3](#) shows an example of the CE Manager GUI with a Hyper-V resource.

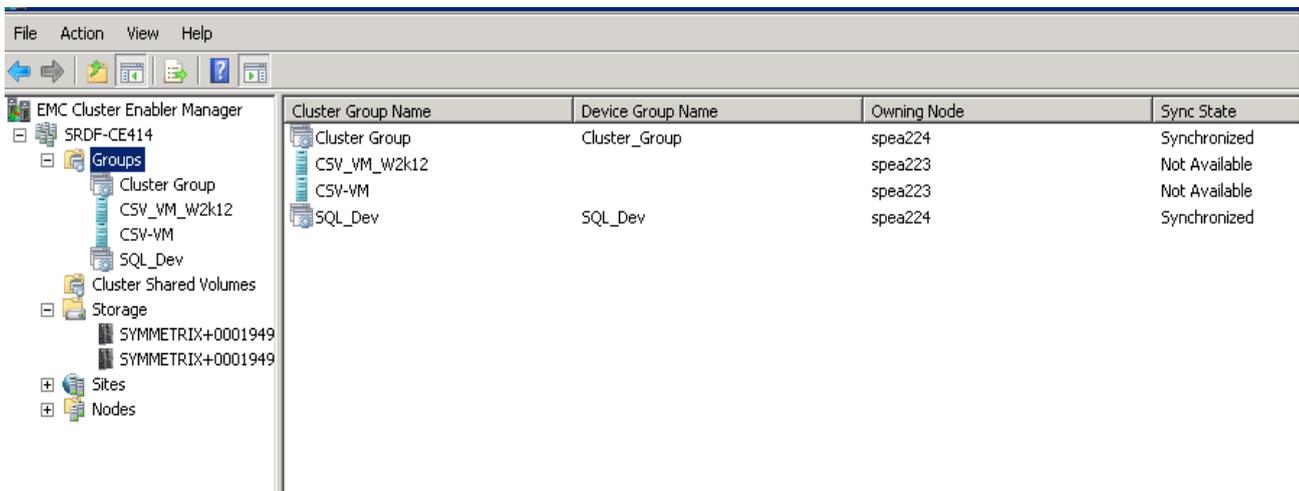


Figure 3 CE Manager with virtual machine cluster group

Cluster Shared Volumes

Cluster Enabler version 4.1.4 supports Windows Server 2008 R2 (x64) and Windows Server 2012 (x64) Cluster Shared Volumes (CSV). CSV is a failover clustering feature that allows all nodes in a cluster concurrent access to data on every CSV-enabled shared disk. Once converted using the CE Configuration wizard, CSV disks display under Cluster Shared Volumes in the navigation tree of the CE Manager. Using Cluster Enabler, you can view the properties or change the failover policies for a CSV disk.

For Windows Server 2008 R2, VMs can exist only on the primary (R1) site. CSV VMs cannot exist on the secondary (R2) site as the R2 devices are read/write disabled (except in the case of failover without swap).

This is different from Failover Cluster behavior without CE configured, where VMs would be allowed on the secondary but be in redirected access mode. The reason for this is that in geocustering, site to site network transfers would have higher network latencies and more expensive bandwidth requirements. So CE restricts VMs to remain on the site on which they have direct access to the disk, and move them only when the CSV disk fails over to the secondary site.

For Windows Server 2012, CSV VMs can run on any node irrespective of where its CSV disk is online. This means that the VM can failover to a node where its CSV disk is marked as write-disabled.

Note: The virtual machine and the CSV disks must first be configured in Microsoft Failover Cluster Manager. CE Manager does not allow custom resource configuration for a specific VM, instead CE Manager configuration wizard can be run to configure for all the VMs in the cluster.

Note: For Windows Server 2008 R2, EMC Cluster Enabler does not support CSV VM live and quick migration between sites.

Converting CSV disks for CE

Before you can manage CSV disks with CE Manager, you must convert the CSV disks using the CE Configuration wizard. Follow the steps in the CE Configuration wizard to configure CSV as you would a CE cluster. All VMs should be brought to the CSV primary site before configuration or they will begin failing over automatically.

Note: If I/O is attempted on a cluster node containing an R2 CSV disk, the node (and only that node) transitions to redirected access. The node returns to direct access only when the mirror is promoted/swapped to a primary mirror.

Note: For MirrorView/CE, when a CSV is converted using CE Configuration wizard, a corresponding consistency group is not created on the CLARiiON, instead the mirror remains an individual mirror.

Note: For Windows Server 2008 R2, a new CE VM resource will be added as part of the virtual machine. This resource is responsible for keeping the virtual machine on the read/write enabled site on which the CSV disk is.

During the various wizards steps, you will notice that the virtual machine and CSV group cluster disks will be validated and converted for management with CE. During the conversion process, the Configuration wizard sets failover policies for each CSV disk and the VM group FailoverThreshold is updated. After completing the wizard, Open CE Manager to view the cluster disk resources listed under the “Cluster Shared Volumes” folder.

Note: When CSV are configured for CE, note that there are no disk resources listed under the new virtual machine. Disk resources are listed under Cluster Shared Volumes.

Managing CSV disks with CE

Once converted, CSV disks can be managed using the CE Manager. The CSV Folder view displays the set of VMs residing on each CSV disk. The CSV disk and the VM details are populated in a tree view. The parent node contains all of the CSV-related data (CSV Path, Owning Node, Device group name, Sync State, etc.). VM details are grouped under the appropriate CSV Parent Node, on which the VM resides. This representation allows you to

see the set of VMs hosted on each CSV disk, and whether, the CSV disk is configured using CE or not. The tree view can be expanded by selecting Expand All or collapsed by selecting Collapse All, which is useful if there is a large number of CSV disks to manage.

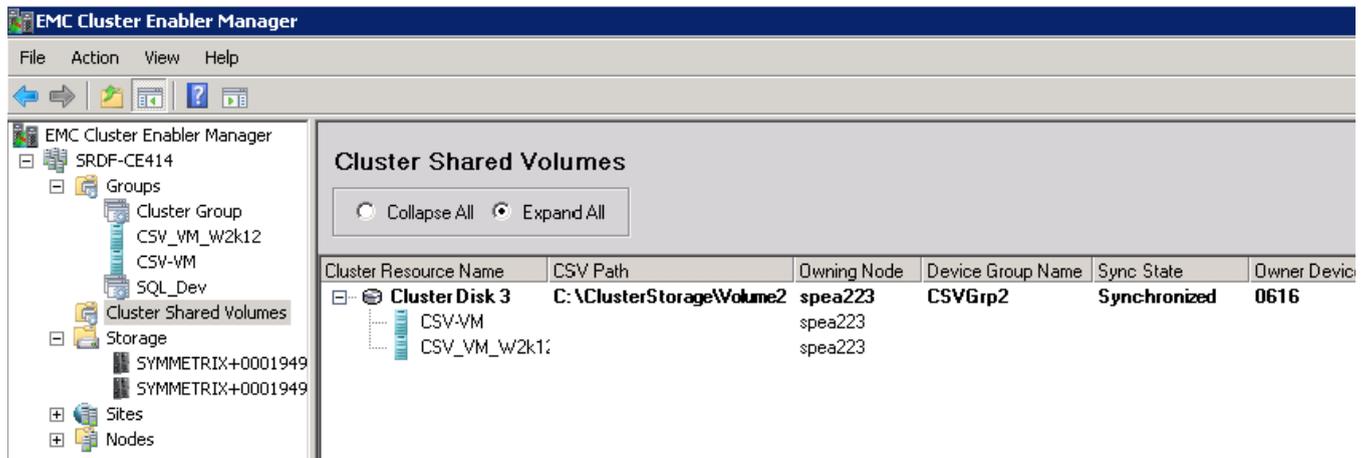


Figure 4 Cluster Shared Volumes tree view

You can change the failover policy for a CSV disk or deconfigure it to remove it from CE Manager control. Right-click on a **cluster disk** to access the CSV action menu. Selecting **Deconfigure CSV From CE** deconfigures the disk from CE Manager control. A dialog box pop-up will appear asking you to confirm the action. Click **Yes** to deconfigure or **No** to abort the action.

Note: If the CSV disk is deconfigured, CE failover support to the remote nodes will no longer be operational. Also all Virtual Machines dependent on that CSV will no longer be managed by CE. To make the CSV disk failover operational again, you will need to reconfigure the CSV and virtual machines using the CE Configuration Wizard in the CE Manager.

Selecting the **Properties** option displays the current properties of a CSV disk. Selecting the **Policies** tab allows you to change the failover behavior for the CSV disk. Selecting the **Refresh** option from the menu updates the CSV view. You can select either **Restrict Group Movement** or **Automatic Failover**. Once selected, click **OK**.

The *Restrict Group Movement* selection restricts the CSV disk from failing over to a peer node. In a replication link failure, this setting will only attempt to move disk laterally. If the replications link is up, this setting has no impact.

The *Automatic Failover* selection allows the CSV disk to automatically failover to any remote site node in the event of a replication link failure.

Selecting a **VM Group** from CE manager displays the dependent CSV properties in the right panel.

Note: CSV disk supports SRDF/Synchronous mode only. Therefore, the SRDF/Asynchronous option is greyed out in the advanced tab setting of the selected CSV disk's properties. The failover behavior for CSV disk currently supports the Restrict Group Movement policy only and is the default selection.

Table 3 depicts CSV features supported for Windows Server 2008 and 2012.

Table 3 CSV feature support matrix

Operating system	CSV disk feature		Virtual machine feature	
	Planned failover	Unplanned failover	Planned failover	Unplanned failover
Windows Server 2008 R2	Yes	No	Partial ^a	No
Windows Server 2012	Yes	Yes	Yes	Yes ^b

a. Has to failover along with CSV disk.

b. Only disk consistency is met.

VMware support

Cluster Enabler version 4.1.4 supports the configuration of a four-node Windows Server 2008 cluster (including R2) or a four-node Windows Server 2012 cluster in VMware ESX Server environments. This section provides instructions for configuring CE in VMware environments.

CE supports two different system configurations, for either:

- ◆ A virtual machine cluster, where the virtual machines reside on two separate physical ESX servers, or
- ◆ A physical-virtual machine cluster, where one node is a physical host, and the other node is a virtual machine on a node in a VMware ESX cluster group.

You must adhere to the following instructions when configuring CE in VMware environments:

1. Ensure that the following applicable software and versions are installed:
 - ESX Server version 5.0 and later
 - Windows Server 2008
 - Windows Server 2008 R2
 - Windows Server 2012
2. You must set the timeout in the `boot.ini` file on all virtual machines to 180 seconds. If the `boot.ini` file currently includes only one entry, the timeout is not effective. You must populate the `boot.ini` with two separate entries. The same entry can appear twice and can be copied and pasted from the original entry. See below for an example of the `boot.ini` file.

```
[boot loader]
Timeout=180
default=multi(0)disk(0)rdisk(0)partition(2)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(2)\WINDOWS="Microsoft Windows XP Professional"
/noexecute=optin /fastdetect
multi(0)disk(0)rdisk(0)partition(2)\WINDOWS="Microsoft Windows XP Professional"
/noexecute=optin /fastdetect
```

Note: No changes are necessary for physical hosts.

3. Configure a dedicated network interface card (NIC) for a heartbeat and associate the virtual NIC to a separate VLAN or network.
4. All data devices are presented to the virtual machines as raw device mapping (RDMs) disks in physical compatibility mode on a separate dedicated shared SCSI adapter.

Note: All gatekeeper devices are presented to the virtual machines as RDMs in physical compatibility mode on a separate dedicated SCSI adapter. The virtual SCSI adapter for the gatekeepers should not be shared with the adapter used for accessing the devices. Gatekeepers presented to the virtual machine should not be presented to any other virtual machine configured in the VMware ESX Server cluster group.

5. You must adhere to all other VMware instructions for the configuration of Failover Clusters. For additional information, refer to the *Setup for Microsoft Cluster Service* technical papers available from VMware at:

<http://www.vmware.com>

Supported functionality

In addition to the Wizards noted in [Table 1, “Cluster Enabler Manager Wizards,”](#) on [page 20](#), the CE Manager provides various features that manage and monitor information for clusters, groups, storage devices, sites, and nodes. [Chapter 6](#) provides information on how to use the Cluster Enabler Manager GUI to complete cluster management tasks.

The following sections explain some of the base functionality available with Cluster Enabler.

Delay Failback

Delay Failback capability is implemented as part of Cluster Enabler’s default functionality. Delay Failback automatically modifies the Preferred Owner list for each failover cluster group so that a failover will occur to a lateral node first, and if the lateral node is unavailable, to a peer node. Lateral nodes are defined as nodes connected to the same storage array. Peer nodes are defined as nodes connected to different storage arrays, located across the link from each other, as shown in [Figure 5](#).

Cluster Enabler manipulates the Microsoft Failover Cluster Preferred Owners list whenever a group is brought online. CE then examines the group Preferred Owners list and determines which node is the lateral node. It can then modify the Preferred Owner list so that the current node and its lateral partner are the first two in the list.

Therefore, no matter which side a group is moved to, the Preferred Owner list is modified to allow a group to fail over to a lateral node, and not fail back or fail over across the link as a first option. Microsoft Failover Clusters only moves a group across the link as a last resort. This prevents the failover clusters from arbitrarily performing what amounts to a failback/failover across the link in an automatic fashion. This feature delays the actual failback of a group from the primary node, and is therefore termed *delay failback*.

Note: The Delay Failback feature overrides all previous configurations in all quorum-based solutions.

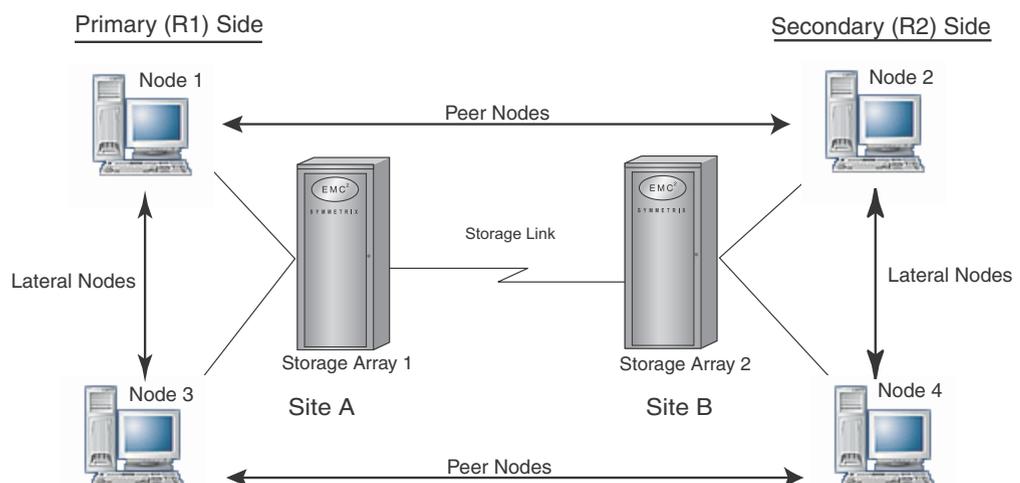


Figure 5 Lateral and peer nodes

Delay Failback runs simultaneously on all nodes. Therefore, when a group comes online on any node, the Preferred Owner list is updated, regardless of whether it is a primary or secondary device. For example, the sequence for Group *x* on Node 1 will be the following:

1. Delay Failback first determines if it *knows* the other nodes in the four-node cluster. This information is gathered by CE during normal operations. If not, the default feature is bypassed because it is unable to differentiate between a lateral node, peer node, and so on.
2. If Delay Failback knows the other nodes, then it determines if Group *x* has come online on Node 1.
3. If Group *x* has come online on Node 1, the Microsoft Failover Cluster Preferred Owner list is modified so that Node 1 is the first Preferred Owner, followed by the lateral node and then the peer nodes.

Enabling and disabling Delay Failback

The Delay Failback feature is enabled by default. This feature can be enabled or disabled by setting a cluster private property using the command-line interface. In the steps shown below, the command prompt is shown as `C:\>`.

1. To verify the current Delay Failback setting, issue the following command:

```
C:\> cluster /priv
```

2. Check the output for the text **DelayFailBackEnabled**. A value of 0 means the feature is disabled. A value of 1 means the feature is enabled. If the **DelayFailBackEnabled** property has not been adjusted, it will not be visible in the `cluster /priv` command output and is assumed to be the default value of 1 (enabled).
3. To disable the Delay Failback setting, issue the following command:

```
C:\> cluster /priv DelayFailbackEnabled=0:DWORD
```

To enable the Delay Failback setting, issue the following command:

```
C:\> cluster /priv DelayFailbackEnabled=1:DWORD
```

Mount point support

Cluster Enabler supports mount points. By using mount points, you can overcome the limitation on drive letters, which makes it possible for a cluster to support more than 26 volumes.

For mount points to work correctly, all related disks must belong to the same cluster group. If related disks are spread across multiple cluster groups, volumes cannot be brought online because cluster groups can be online on different nodes. To avoid this scenario, Cluster Enabler first groups all related disks by identifying the mount points on a given disk and any disks upon which the given disk is mounted. Cluster Enabler then creates a parent/child relationship between the disks.

When a user chooses a disk to create a group (or adds a disk to an existing group), Cluster Enabler finds all related disks by traversing its parent/child relationships and adding every related disk to the group. It then adds appropriate dependencies between the disks so that the resources can be brought online in an orderly fashion.

[Table 4 on page 35](#) illustrates a cluster example consisting of drive letters and mount points for six volumes. Using this configuration, you can see various parent/child relationships among the disks.

For example, the user chooses `E:\MNT1`. Therefore:

- ◆ `E:\MNT1` is a mount point with `E:\` as its parent.
- ◆ `E:\` is a child of `F:\`. Thus, disk `F:\` is included in the group.
- ◆ `F:\` has additional children `F:\MNT2` and `F:\MNT2\MNT3`. Thus, the group includes these disks too.

The result of these parent/child relationships is that the group will include volumes `OBCE`, `OBCF`, `OBD0`, `OBD1`, and `OBD2`. Each disk is dependent on its parent to come online. In this example, `OBCF` is dependent on `OBCE`, and `OBD0` is dependent on `OBCE`, and so forth.

Of course, each group is also dependent on the Cluster Enabler resource.

Table 4 Cluster mount point example

Drive letter and mount point	Symmetrix volume ID
<code>F:\</code>	<code>OBCE</code>
<code>F:\MNT1</code> , <code>E:\</code>	<code>OBCF</code>
<code>F:\MNT2</code>	<code>OBD0</code>
<code>F:\MNT2\MNT3</code>	<code>OBD1</code>
<code>D:\</code>	<code>OBCD</code>
<code>E:\MNT1</code>	<code>OBD2</code>

When you delete a device, Cluster Enabler finds all related disks and deletes them too. For example, if the current mount points are `F:\` and `F:\MNT2` and `F:\MNT2\MNT3`, and if the device that corresponds to `F:\MNT2` is deleted from the group, all three devices corresponding to `F:\`, `F:\MNT2`, and `F:\MNT2\MNT3` are deleted.

However, if you were to first delete mount point `F:\MNT2` from the operating system and then delete its corresponding device from the group, Cluster Enabler would delete only the devices that correspond to `F:\MNT2` and `F:\MNT2\MNT3`. The device corresponding to `F:\` would be left in the group because, after the mount point deletion, it is no longer related to `F:\MNT2`.

Multiple storage array support

Cluster Enabler for Microsoft Failover Clusters supports the use of multiple storage arrays per cluster. This feature provides greater flexibility to you and your storage provisioning.

Delegating CE administration

The CE Manager lets you manage multiple CE clusters simultaneously, as long as all of the clusters are either Windows Server 2008 or 2012 clusters and are in the same domain. To manage the cluster, CE Manager and Cluster Administrator are used with a domain account, which is part of local administrator group on every cluster node. This effectively grants full control of every cluster node to the domain account used to manage the cluster.

CE provides a utility that allows the most common CE and cluster management tasks to be delegated to a non-local administrator. To support this option, a command-line utility, called `cesec.exe` is used on each cluster node after the cluster has been fully configured.

Using the cesec.exe utility

The `cesec.exe` command-line utility allows the local administrator to delegate the most common CE cluster tasks to non-local administrators by adding a domain group (recommended) or a domain user.

A CE cluster must have already been configured by a local administrator using the CE Configuration Wizard. The `cesec.exe` utility is located in the CE install directory (typically `C:\Program Files\EMC\Cluster-Enabler`) and must be run on all nodes in the cluster by a local administrator. On Windows Server 2008, the utility must be run from an elevated command prompt.

Note: Due to a Microsoft limitation, the Windows Server 2008 Failover Cluster Manager cannot be used by a non-local administrator account, even if that account has been granted full control of the cluster. Use the `cluster.exe` command-line utility instead.

System security changes

Running the `cesec.exe` command-line utility allows you to change the following security administration privileges:

- ◆ Allows a non-local administrator to manage the cluster.
- ◆ Allows a user to make remote DCOM connections.
- ◆ Opens the Windows Firewall for the Windows Management Instrumentation (WMI) rule group on Windows Server 2008.
- ◆ Allows remote write access to the following WMI namespaces: `Root/CIMV2`, `Root/EMC`, and `Root/MSCluster`.
- ◆ Allows a user to query the Service Control Manager and to control the following CE-related services: Cluster Service (`clussvc`), CE Event Trace Service (`ce_eventtrace`), and CE Service (`cesvc`).
- ◆ Allows remote access to the CE portion of the registry (`HKLM\SOFTWARE\EMC\CE`).
- ◆ Allows the user to export CE log files by granting write access to the CE log directory (typically `C:\Program Files\EMC\Cluster-Enabler\Logs`).

Restrictions

Certain CE configuration operations are not allowed. The following CE configuration changes are blocked:

- ◆ CE install/uninstall
- ◆ Using the Configuration Wizard to convert MS clusters to CE clusters
- ◆ Adding and deleting nodes for an existing cluster
- ◆ De-configuring a CE cluster

Command syntax

The following is the `cesec.exe` command syntax:

```
cesec.exe -ce <action> <principal>

action set | remove | list

principal domain\user | domain\group | user@domain.com |
group@domain.com
```

Options:

`-ce <action> <principal>`

Modifies the CE security settings for a principal.

`-ce list`

Lists the security settings relevant to CE.

Usage examples

The following examples assume that the domain users who manage CE have been added to a domain group called `DomainName\CE Admins`.

To allow the domain group to manage CE, enter type:

```
cesec.exe -ce set DomainName\CE Admins
```

To remove the domain group, enter type:

```
cesec.exe -ce remove DomainName\CE Admins
```

To list your current security settings, enter type:

```
cesec.exe -ce list
```

Viewing cluster dependency

Cluster Enabler provides a cluster dependency viewer that allows you to view or print cluster configuration data (Dependency Report) showing all CE cluster groups and device dependencies for the cluster. This tool can be used to graphically display complex storage site configurations for a CE cluster. Point-to-point, cascaded, and concurrent configurations are supported. The expanded view displays all devices involved in each site and the replication mode between sites.

Groups in each site are sorted alphabetically, and devices are color coded by site. The dependency viewer also allows you to sort CE groups by site. Interconnection between devices is labeled by the mode of replication (that is Sync or Async). Remote adapter (RA) numbers are displayed for each leg of all configurations. The CSV group is displayed with a CSV group name, instead of the GUID. The CSV-Virtual Machine groups are grouped under the particular CSV group on which they reside.

Follow these steps to view and print a Dependency Report for a cluster:

1. Open the CE Manager, select the cluster in the Navigation Tree, and select **Action**, and **View Dependency** from the menu bar. The View Dependency option can also be launched using the right-click menu from a selected cluster.

Note: There may be a wait period while acquiring site information before displaying the actual diagram.

2. Select each cluster group, and double-click the disk objects to expand the view for each site. Devices are color coded by site as noted in the right-side key display.

Note: Select the **Expand All/Collapse All** option to expand or collapse group details.

3. From the Dependency Report top menu bar, click **Sort Group by Site** to change the Site view.
4. To preview a diagram print layout, select the **Print Preview** icon from the Dependency Report top menu bar. To print the diagram select the **Print** icon.

Figure 6 shows a sample Dependency Report.

Note: CE Dependency Viewer displays both the CSV disk group name and the path for a configured CSV disk. If not configured, only the path is shown. For example, if configured, the group name displays as:

CSVGrp1 [C:\ClusterStorage\Volume1].

Otherwise, only the CSV Path is displayed as:

[C:\ClusterStorage\Volume1].

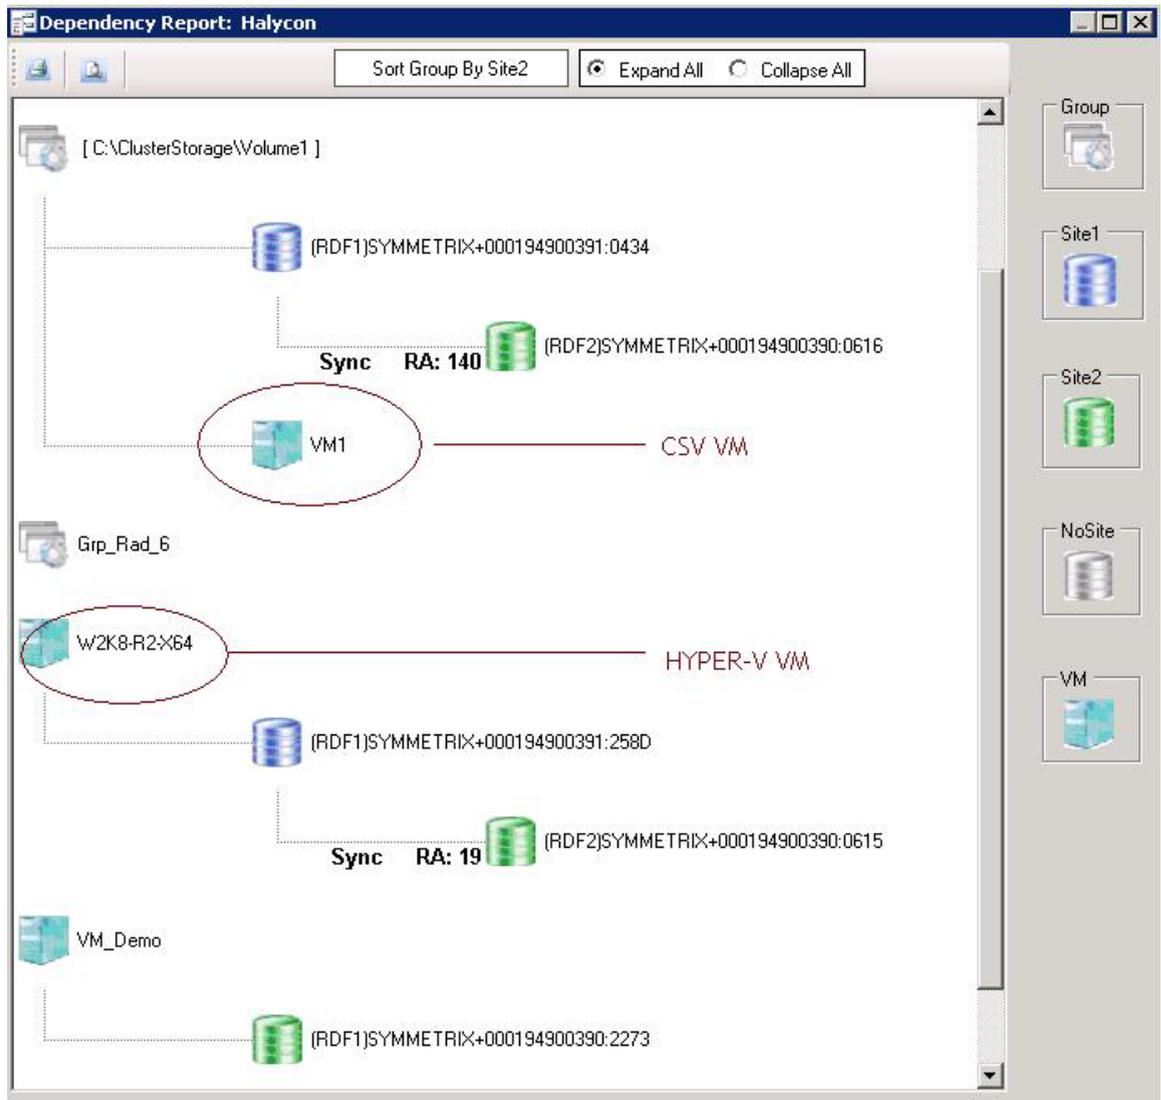


Figure 6 Sample Dependency Report

Note: The following notations apply to [Figure 6](#):
 CSV VM – VM resides on the CSV disk.
 HYPER-V VM – VM resides on regular RDF disk.

SRDF/CE support matrix

Table 5 shows the SRDF/CE support matrix for Microsoft Windows Server operating systems by CE version and the required minimum Solution Enabler version.

Table 5 SRDF/CE support matrix

Microsoft Windows Server OS (Enterprise or Datacenter Editions only)	Operating System Support Matrix for SRDF/CE							
	CE 4.1.4	CE 4.1.3	CE 4.1.2	CE 4.1.1	CE 4.1.0	CE 4.0.1	CE 4.0.0	CE 3.1
W2K12 x64	✓	⊗	⊗	⊗	⊗	⊗	⊗	⊗
W2K12 x64 CORE	✓	⊗	⊗	⊗	⊗	⊗	⊗	⊗
W2K8 R2 x64 SP1	✓	✓	✓	✓	✓	✓	✓	⊗
W2K8 R2 x64 CORE SP1	✓	✓	✓	✓	✓	✓	✓	⊗
W2K8 R2 IA64 SP1	⊗	✓	✓	✓	✓	✓	✓	⊗
W2K8 x64 SP2	✓	✓	✓	✓	✓	✓	✓	✓
W2K8 IA64 SP2	⊗	✓	✓	✓	✓	✓	✓	✓
W2K8 x86 SP2	✓	✓	✓	✓	✓	✓	✓	✓
W2K3 x64 SP2	⊗	✓	✓	✓	✓	✓	✓	✓
W2K3 xIA64 SP2	⊗	✓	✓	✓	✓	✓	✓	✓
W2K3 x86 SP2	⊗	✓	✓	✓	✓	✓	✓	✓
Minimum Solutions Enabler (SE) version required	V7.3.0.1			V7.0.1				V6.5.2

CHAPTER 2

About SRDF/Cluster Enabler

This chapter provides an introduction to the SRDF/Cluster Enabler plug-in module and explains how EMC Cluster Enabler provides disaster recovery protection in geographically distributed Microsoft Failover Clusters using the Symmetrix Remote Data Facility (SRDF).

- ◆ SRDF/Cluster Enabler plug-in overview 42
- ◆ SRDF overview 43
- ◆ SRDF/CE supported features 45
- ◆ Concurrent SRDF 48
- ◆ Cascaded SRDF 51
- ◆ Pre-SRDF/CE clustering considerations 54

SRDF/Cluster Enabler plug-in overview

SRDF/Cluster Enabler (SRDF/CE) is a software plug-in module to EMC Cluster Enabler for Microsoft Failover Clusters software. The Cluster Enabler (CE) plug-in architecture consists of a CE base module component and separately available plug-in modules, which support your chosen storage replication technology. The CE base component must be installed prior to installing a plug-in module.

Note: Refer to the EMC Online Support for CE plug-in software module availability for your replication technology or check with your EMC sales representative.

The SRDF/CE plug-in module provides a software extension of failover clusters functionality that allows Windows Server 2008 (including R2) and Windows Server 2012 Enterprise and Datacenter editions running Microsoft Failover Clusters to operate across multiple connected Symmetrix arrays in geographically distributed clusters.

Each cluster node is connected through a storage network to the supported Symmetrix array. Once configured using the EMC Cluster Enabler Manager graphic user interface (GUI), Microsoft Failover Clusters are referred to as CE clusters.

Important: Mixed replication technologies/storage configurations are not supported. For example, Symmetrix® SRDF and CLARiiON MirrorView devices cannot be part of the same CE cluster group.

Cluster Enabler expands the range of cluster storage and management capabilities while ensuring full business continuance protection. An iSCSI or Fibre Channel connection from each cluster node is made to its own Symmetrix array. Two connected Symmetrix arrays provide automatic failover of mirrored volumes during a Microsoft failover cluster node failover.

This connection effectively extends the distance between cluster nodes (depending on network latency) and forms a geographically distributed cluster with disaster-tolerant capabilities. ²

[Figure 7 on page 43](#) provides a graphical example of using Cluster Enabler in an SRDF Symmetrix array environment.

Cluster Enabler protects data from the following types of failures, 24 hour a day, 7 days a week, and 365 days per year:

- ◆ Storage failures
- ◆ System failures
- ◆ Site failures

2. The *EMC Networked Storage Topology Guide* provides additional information regarding distance restrictions for your specific configuration.

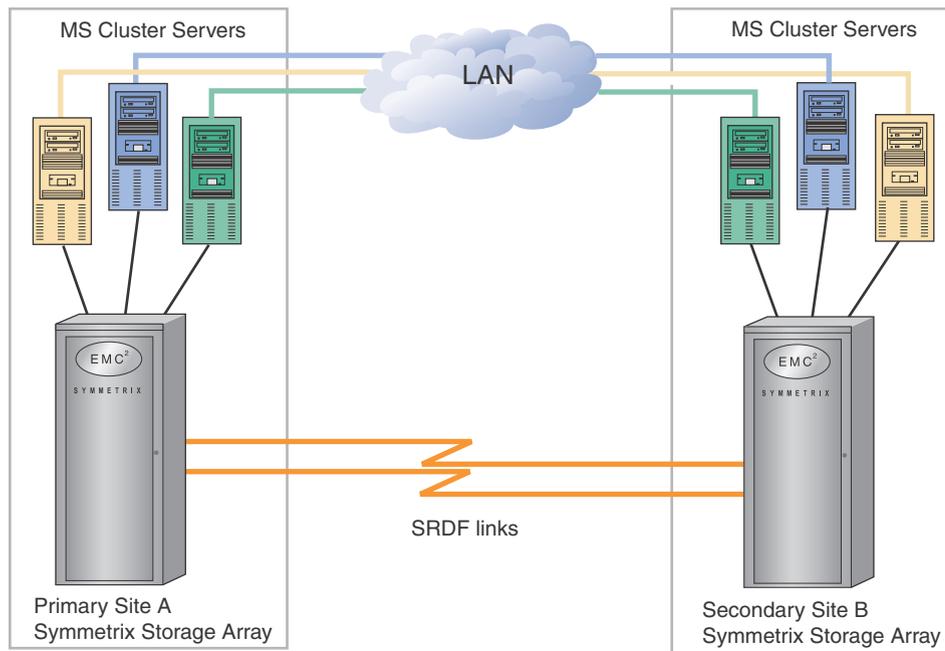


Figure 7 Overview example of an SRDF/CE cluster configuration

SRDF overview

The Symmetrix Remote Data Facility (SRDF) is a Symmetrix-based business continuance and disaster recovery solution sold as a separate license by EMC Corporation. In basic terms, SRDF is a configuration of multiple Symmetrix arrays whose purpose is to maintain multiple, real-time copies of logical volume data in more than one location.

SRDF duplicates production (*source*) site data to a recovery (*target*) site transparently to users, applications, databases, and host processors. If the primary site is not able to continue processing, data at the secondary site is current up to the last I/O transaction.

SRDF can be used in several key areas including, but not limited to:

- ◆ Disaster Recovery
- ◆ Remote Backup
- ◆ Data Center Migration
- ◆ SDMS—Symmetrix Data Migration Service
- ◆ Data Center Decision Solutions

When primary (*source*) systems are down, SRDF enables fast switch over to the recovery (*target*) copy of the data, allowing critical information to become available in minutes. Business operations and related applications may resume full functionality with minimal interruption.

Protecting against data loss allows the operations and applications to resume at the secondary site. SRDF can be used:

- ◆ By itself, and data processing can be resumed by powering up a standby system and manually restarting.

- ◆ In combination with more sophisticated software to automatically resume operations.

Figure 8 on page 44 illustrates a basic SRDF configuration.

SRDF/CE combines Microsoft Failover Clusters and SRDF to provide a more sophisticated solution. SRDF/CE provides an automated configuration wizard to be used in conjunction with the Microsoft Cluster Administrator to administer the SRDF-enabled cluster.

Note: For greater detail on SRDF, consult the SRDF documentation set, “[Related documentation](#)” on page 12.

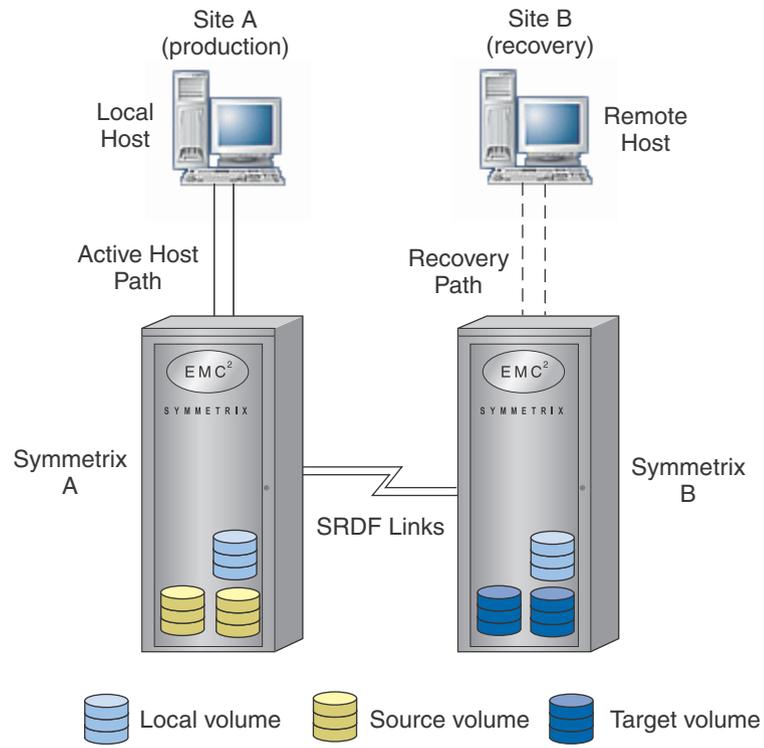


Figure 8 Basic SRDF configuration

SRDF/CE supports both SRDF/Synchronous (SRDF/S) and SRDF/Asynchronous (SRDF/A) modes of transfer.

SRDF/CE supported features

SRDF/Cluster Enabler for Microsoft Failover Clusters provides a graphic user interface called Cluster Enabler (CE) Manager. The CE Manager provides several wizard processes to streamline cluster tasks and reduce the complexity of typical cluster management. [Table 6](#) lists the various wizard processes that are included in the CE Manager.

Table 6 Cluster Enabler Manager wizards

Wizard	Functionality
Configuration Wizard	Configures failover clusters for management with Cluster Enabler
Create Group Wizard	Creates a CE group, adds devices and selects a group policy
Modify Group Wizard	Modifies a CE group to add or remove devices
Recover CE Cluster Wizard	Recovers a CE shared quorum cluster
Change Quorum Wizard	Changes the quorum model of a cluster
Update Mirror Pairs Wizard	Updates the mirrored pairs in a cluster
Storage Discovery Wizard	Discovers and sets up the attached storage for Cluster Enabler

Note: [“Using Cluster Enabler Manager” on page 81](#) provides detailed descriptions and functionality for each wizard.

In addition to the wizard processes noted in [Table 6](#), the CE Manager provides various features that manage and monitor cluster, group, storage device, site, and node information. [“Using Cluster Enabler Manager” on page 81](#) provides information on additional GUI features.

The following listed SRDF/CE plug-in module features are specific to the SRDF and Symmetrix replication technology:

- ◆ [“SRDF/Asynchronous compatibility” on page 45](#)
- ◆ [“SRDF/CE swap support” on page 46](#)
- ◆ [“Symmetrix Virtual Provisioning” on page 47](#)
- ◆ [“Supported devices” on page 47](#)
- ◆ [“SRDF/CE configuration with multiple remote adapters” on page 47](#)
- ◆ [“Monitoring SRDF link status” on page 47](#)
- ◆ [“SRDF composite groups” on page 48](#)
- ◆ [“Concurrent SRDF” on page 48](#)
- ◆ [“Cascaded SRDF” on page 51](#)

SRDF/Asynchronous compatibility

SRDF/CE is compatible with EMC SRDF/Asynchronous (SRDF/A). SRDF/A is a high-performance, extended-distance asynchronous replication that uses a delta set architecture for reduced bandwidth requirements and no host performance impact.

Asynchronous mode provides a point-in-time image on the target (R2) device that is only slightly behind the source (R1) device. SRDF/A session data is transferred to the remote Symmetrix system in delta sets, eliminating the redundancy of same-track changes being

transferred over the link, thereby reducing the required bandwidth. SRDF/A only needs enough bandwidth to support the average production workload versus peak workloads, provided there is enough Symmetrix cache to support the peak workloads.

SRDF/A is intended for users who require no host application impact while maintaining a consistent, restartable image of their data on the R2 side at all times.

Note: SRDF/CE always enables consistency on SRDF/A groups. SRDF/A consistency ensures that applications have a consistent copy on the remote side when they failover.

SRDF/CE supports Enginuity™ releases as outlined in the *E-Lab Interoperability Navigator*. At the 5x70 Enginuity level and later, you can specify a single SYMCLI group whose device members have been previously defined as SRDF/A enabled. Once configured, SRDF/CE automatically fails over this group to the target side as necessary.

Note: SRDF/CE does not support clusters where the target (R2) side is larger than the source (R1) side. When the system fails over to the R2 side, it can never fail back since the R2 cannot resynchronize all its data back to the R1 side.

Note: SRDF/A is not supported for quorum group in shared quorum models. Other groups in the cluster may use synchronous or asynchronous modes as desired.

SRDF/CE swap support

An R1/R2 personality swap (or R1/R2 swap) refers to swapping the RDF personality of the RDF device designations of a specified device group, so that source R1 devices become target R2 devices and target R2 devices become source R1 devices.

R1/R2 RDF swaps are available with Enginuity Version 5567 or later. There are two types of R1/R2 swaps: FastSwap and Dynamic Swap. A FastSwap occurs immediately after failover if the group is fully synchronized. A Dynamic Swap takes longer because after failover, the tracks are checked to determine if they are synchronized, and then the swap occurs. If you enable an R1/R2 swap for a group, SRDF/CE automatically checks during a failover to determine whether FastSwap is available. If FastSwap is available, SRDF/CE will use it. If FastSwap is not supported, SRDF/CE will automatically use Dynamic Swap.

R1/R2 swap benefits

This section describes several scenarios in which it is beneficial to execute an R1/R2 swap.

Symmetrix array load balancing

In today's rapidly changing computing environments, it is often necessary to deploy applications and storage on a different Symmetrix array without having to lose disaster protection. R1/R2 swap can enable this redeployment with minimal disruption, while offering the benefit of load balancing across two Symmetrix storage arrays.

For example, if you want to reconfigure an SRDF/CE environment after having decided where the R1 and R2 devices will sit, this procedure will allow you to go from an active/passive configuration to active/active.

Primary data center relocation

Sometimes a primary data center needs to be relocated to accommodate business practices. For example, several financial institutions in New York City routinely relocate their primary data center across the Hudson River to New Jersey as part of their disaster drills. R1/R2 swaps allow these customers to run their primary applications in their New Jersey data centers. The Manhattan data centers then acts as the disaster protection site.

Post-failover temporary protection measure

You can regain a measure of protection after failing over to the remote site. If the hosts on the source side are down for maintenance, R1/R2 swap permits the relocation of production computing to the target site without giving up the security of remote data protection. When all problems are solved on the local Symmetrix array, fail over again and swap the personality of the devices to return to the original configuration.

Symmetrix Virtual Provisioning

SRDF/CE supports Symmetrix Virtual Provisioning™ with SRDF/Synchronous and SRDF/Asynchronous. Thin devices can be used in a CE cluster. The *EMC Solutions Enabler Array Controls CLI Product Guide* provides details on how to set up Symmetrix Virtual Provisioning.

Supported devices

SRDF/CE supports the following types of devices in point-to-point, cascaded, and concurrent configurations:

- ◆ Standard
- ◆ RAID-5
- ◆ RAID-6
- ◆ TDEVs(thin devices)
- ◆ Diskless (applicable to cascaded site A to site C configurations only)

SRDF/CE configuration with multiple remote adapters

SRDF/CE can be configured with multiple RDF links and remote adapter (RA) groups. SRDF/CE not only allows multiple RAs, but periodically tests them to ensure they are functioning. Multiple RA groups are also allowed, and these RA groups do not have to be symmetrical across all RDF links; any one RA group can be allocated over a subset of the defined RDF links.

If a situation occurs where an RDF link goes down, an event log message is posted and an entry is placed in the SRDF/CE log.

Monitoring SRDF link status

SRDF/CE provides a health monitoring feature for the SRDF link that allows you to view link status error messages, which are reported in the Windows event log. This feature allows you to monitor various scenarios, such as SRDF link failure.

Refer to [“Windows event log messages”](#) on [page 23](#) for more information.

SRDF composite groups

SRDF/CE supports the use of SRDF composite groups (CG) that span across multiple RDF groups (also called RA groups). SRDF/CE allows the creation and modification of composite groups as CE groups.

Note: SRDF/CE does not support composite groups that span across multiple Symmetrix arrays.

Cascaded SRDF also makes use of composite groups for consistency protection during failover and failback operations..Support for composite groups requires that the Solutions Enabler RDF daemon (`storrdafd`) be enabled. Refer to the *EMC Solutions Enabler Installation Guide* for information on enabling the RDF daemon.

For Engenuity versions earlier than 5875, support for composite groups requires a Solutions Enabler SRDF/Consistency Group license. Refer to the *EMC Solutions Enabler Installation Guide* for information on the appropriate license keys.

[“Create Group Wizard” on page 91](#) explains how to create a composite CE group. [“Modify Group Wizard” on page 93](#) explains how to add or remove devices from a composite CE group.

For additional information on composite groups, refer to the *EMC Solutions Enabler Symmetrix SRDF Family CLI Product Guide*.

Concurrent SRDF

SRDF/CE supports a concurrent SRDF configuration. In a concurrent SRDF configuration, a single source (R1) device is remotely mirrored to two target (R2) devices at the same time. A concurrent SRDF configuration allows you to have two identical remote copies available at any point in time. It is valuable for duplicate restarts and disaster recovery, and provides increased flexibility for data mobility and application migrations.

Concurrent SRDF technology can use two different RA adapters in the interface link to achieve the connection between the R1 device and its two concurrent R2 mirrors. Each of the two concurrent mirrors must belong to a different SRDF (RA) group.

Note: The *EMC Solutions Enabler SRDF Family Product Guide* provides configuration details for setting up a concurrent SRDF configuration.

For SRDF/CE support, if there is one R1 device paired with two R2 devices, only one of the R2 devices should be mapped to the host on the secondary site. [Figure 9 on page 49](#) shows a concurrent SRDF configuration. Note that the Symmetrix at Remote Site C cannot be mapped to the cluster host.

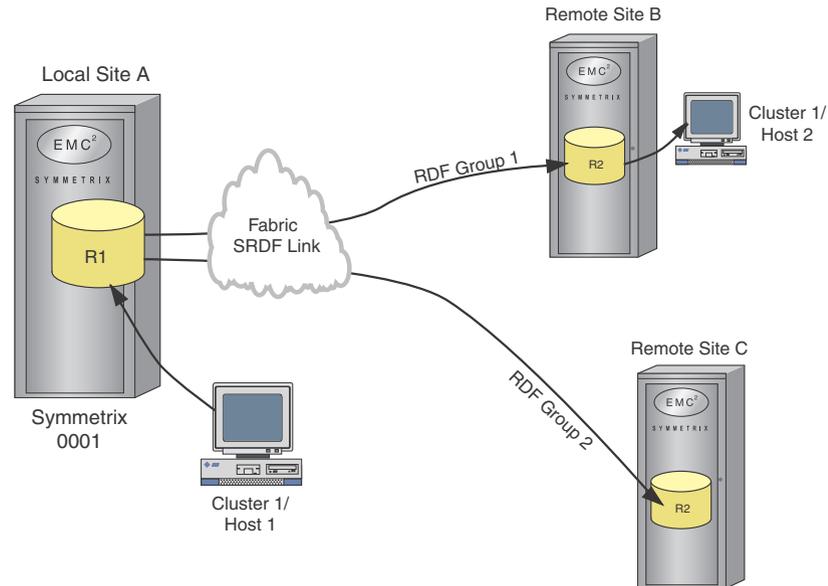


Figure 9 SRDF/CE with concurrent SRDF

Restrictions and limitations

In concurrent configurations, the CE cluster nodes should be present in only two of the three sites. If CE cluster nodes are present in all three sites, the configuration is not supported.

CE configurations where devices in a consistency group are mixed between concurrent, cascaded and point-to-point configurations are not supported.

Supported SRDF modes for concurrent SRDF/CE configurations are:

- Synchronous for Site A to Site B (R11->R2)
- Synchronous or asynchronous for Site A to Site C (R11 -> R2)

Note: Concurrent topology is not supported when Cluster Enabler is configured for use with VMAX 10K/VMAXe arrays.

Failover/Failback behavior

This section describes the failover and failback behavior for concurrent configurations for both planned and unplanned failovers.

Planned Failovers The following results occur for planned failover scenarios in a concurrent configuration:

Failover/Failback between sites A and C in a concurrent configuration (sync A -> B/async A-> C)

Note: CE is installed at sites A and C.

In a concurrent configuration, a planned failover between sites A and C in asynchronous mode, where SRDF does not support swap, results in sites A and B remaining as the R1 and R2. With an RDF pair state of Failed Over, where the R1 becomes write-disabled and the R2 becomes read-write enabled. Sites A and B change to an RDF pair state of invalid.

When the failback between sites A and C in this scenario is initiated, the configuration reverts back to the original concurrent configuration, with sites A and C in asynchronous mode with an RDF pair state of consistent. Sites A and B reverts back to an RDF pair state of synchronized.

Failover/Failback between sites A and B in a concurrent configuration (sync A -> B/async A -> C)

Note: CE is installed at sites A and B.

In a concurrent configuration, a planned failover between sites A and B in synchronous mode, a swap is performed that results in cascaded configuration, where site A becomes an R21, site B becomes an R1, and site C becomes an R2.

Failover/Failback between sites A and B in a concurrent configuration (sync/sync, where SRDF does not support swap)

In a concurrent configuration of sync/sync, where SRDF does not support swap, a planned failover between sites A and B in synchronous mode, results in sites A and B remaining as the R1 and R2, with an RDF pair state of Failed Over. Sites A and C change to an RDF pair state of invalid.

When the failback between sites A and B in this scenario is initiated, the configuration reverts back to the original concurrent configuration, with sites A, B, and C in synchronous mode with an RDF pair state of synchronized.

Unplanned Failovers

The following results occur for unplanned failover scenarios in a concurrent configuration involving storage failure:

Storage failure at site A (async A->B, sync A->C)

In a concurrent configuration where the RDF mode between sites A and B is asynchronous, and synchronous between sites A and C, a storage failure at site A causes the RDF pair state between sites A and B, and A and C to become partitioned. In this failover scenario, Cluster Enabler fails over to either site B or C, depending on which node is configured as a CE site.

Storage failure at site A (sync A->B, sync A->C)

In a concurrent configuration where the RDF mode between sites A and B is synchronous, and synchronous between sites A and C, a storage failure at site A causes the RDF pair state between sites A and B, and A and C to become partitioned. In this failover scenario, Cluster Enabler fails over to either site B or C, depending on which node is configured as a CE site.

Cascaded SRDF

SRDF/CE supports a Cascaded SRDF configuration. Cascaded SRDF is a three-way data mirroring and recovery solution that provides enhanced replication capabilities, greater interoperability, and multiple ease-of-use improvements. Cascaded SRDF support allows replication between three sites without requiring the need for SRDF BCVs on the second Symmetrix array. A cascaded SRDF configuration does not require three separate site locations, although that is the most common configuration for a disaster recovery solution.

The basic cascaded SRDF configuration consists of a primary site (*SiteA*) replicating data to a secondary site (*SiteB*) and replicating the same data to a tertiary site (*SiteC*), as shown in [Figure 10](#). Note that the Secondary SiteB device is labeled R21. This device is the R2 mirror of the Primary SiteA R1 device, and the R1 mirror of the Tertiary SiteC R2 device. The SiteA and SiteB devices have an SRDF pair state and the SiteB and SiteC devices have an SRDF pair state.

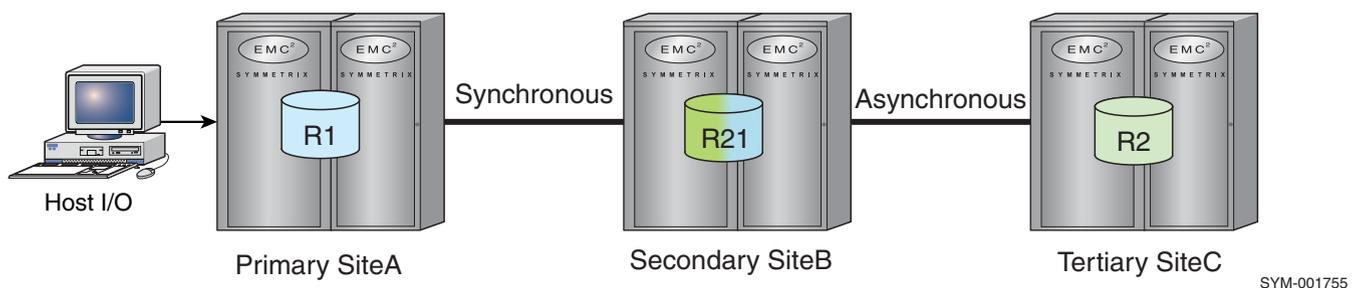


Figure 10 Sample SRDF Cascaded configuration

Note: The *EMC Solutions Enabler SRDF Family Product Guide* provides configuration details for setting up a cascaded SRDF configuration.

Cascaded SRDF/CE requirements

The following requirements are needed for cascaded SRDF/CE support:

- ◆ The secondary site (with the R21 devices) must be running Symmetrix DMX-3 or DMX-4 or higher hardware platforms with Enginuity 5773 and higher.
- ◆ R1 and R2 devices that are paired with R21 devices must be in an array that is running Enginuity 5671 or 5772 and higher.
- ◆ SRDF modes for cascaded SRDF/CE support are:
 - Synchronous for Site A to Site B (R1->R21)
 - Asynchronous for Site B to Site C (R21 -> R2)

Note: Currently ACP disk mode from B to C is not supported.

Restrictions and limitations

In cascaded configurations, the CE cluster nodes should be present in only two of the three sites. If CE cluster nodes are present in all three sites, the configuration is not supported.

The following restrictions and limitations are listed for SRDF/CE cascaded support:

- ◆ In a cascaded configuration, CE can be installed only at the following sites:
 - At Sites A and B; where the replication mode is synchronous or
 - At sites A and C
- ◆ CE configurations where devices in a consistency group are mixed between concurrent, cascaded and point-to-point configurations are not supported.
- ◆ Cascaded disks cannot be used as the quorum disk.
- ◆ Cascaded RDF devices are not discovered in CE if the R21 and R2 Symmetrix arrays are mapped to the same R2 host.
- ◆ Cascaded topology is not supported when Cluster Enabler is configured for use with VMAX 10K/VMAXe arrays.

Failover/Failback behavior

This section describes the failover and failback behavior for cascaded configurations for both planned and unplanned failovers.

Note: Cascaded SRDF makes use of composite groups for consistency protection during failover and failback operations. Support for composite groups requires that the Solutions Enabler RDF daemon (`storrd`) be enabled." c.f. section on composite groups.

Planned Failovers

The following results occur for planned failover scenarios in a cascaded configuration:

Failover/Failback between sites A and B in a cascaded configuration

In a cascaded configuration, a planned failover between sites A and B in synchronous mode, results in a swap and hence a concurrent configuration.

When the failback between sites A and B in this scenario is initiated, the configuration reverts back to a cascaded configuration.

Failover/failback between sites A and C in a cascaded configuration

In a cascaded configuration, a planned failover between sites A and C involves two consecutive failovers between sites A , B, and C as follows:

- ◆ A failover between sites A and B (synchronous mode)
- ◆ A failover between sites B and C (asynchronous mode)

Unplanned Failovers

The following results occur for unplanned failover scenarios in a cascaded configuration involving storage failure:

Failover from site A to site C with a storage failure at site A

In a cascaded configuration where the RDF mode between sites A and B is synchronous, and asynchronous between sites B and C, a storage failure at site A causes the RDF pair state between sites A and B to become partitioned, and consistent between sites B and C. In this failover scenario, CE fails over from B to C and the applications are transitioned (online) to site C.

When site A storage is restored, the RDF pair state between Site A and B become suspended, and consistent between B and C. To transition applications back online to site A, CE performs the following steps:

1. A failover from site A to site B
2. A fallback from site C to site B
3. A fallback from site B to site A

Once these steps are completed, the RDF pair state between sites A and B returns to a synchronized state, and the RDF pair state between sites B and C returns to a consistent state.

Storage failure at sites A and B

In a cascaded configuration where the RDF mode between sites A and B is synchronous, and asynchronous between sites B and C, a storage failure at both sites A and B causes sites A and B to become unreachable, with an RDF pair state of partitioned between sites B and C. If transmit idle is disabled, a failover must be performed to bring applications online at site C. If transmit idle is enabled, applications are automatically brought online at site C. When site A storage is restored, the RDF pair state between sites A and B is either suspended or split, and the RDF pair state between B and C is split.

Note: In a split state scenario, Cluster Enabler requires administrator intervention to return the SRDF pair state to Failed Over.

To transition applications back online to site A, CE performs the following steps:

1. A failover from site A to site B
2. A fallback from site C to site B
3. A fallback from site B to site A

Configuring cascaded SRDF with CE Manager

After all nodes have been discovered, the CE wizard validates the presence of CE cluster nodes in two of the three storage sites. The SRDF configuration is supported for synchronous replication between R1 and R21 devices and asynchronous replication between R21 and R2 devices. [“Viewing cluster dependency” on page 37](#) provides information about viewing existing SRDF configurations. R21 devices at site B display as device type RDF21.

Pre-SRDF/CE clustering considerations

To ensure disaster recovery protection in an SRDF/CE-enabled cluster, consider the following prior to its installation and configuration:

- ◆ Cabling
- ◆ Booting
- ◆ SRDF coexistence

Cabling

Avoid routing all cables through the same path, both in buildings and between sites. To provide an installation with no single point of failure, use a configuration similar to [Figure 11](#).

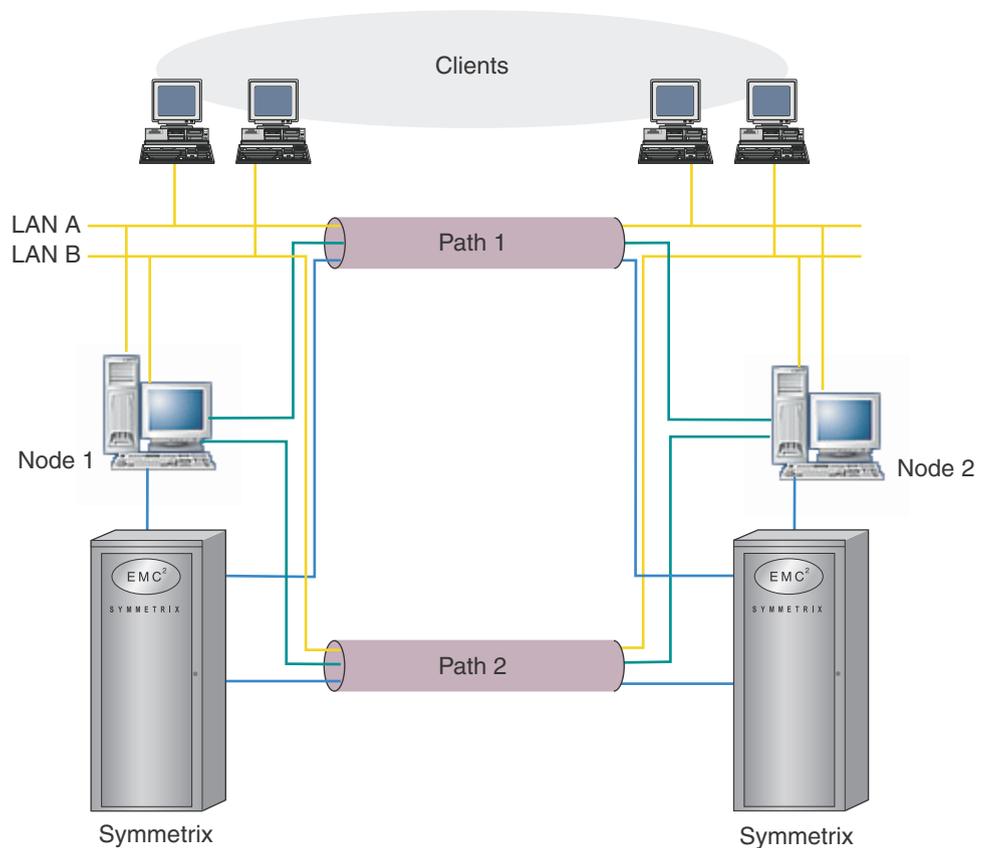


Figure 11 Recommended cabling configuration

Booting

Currently, Microsoft Failover Clusters can only boot from separate private disks (cannot boot off the same bus). Therefore, CE nodes must contain an internal disk for booting or be attached to a nonclustered disk.

SRDF coexistence

Multiple SRDF/CE clusters can share the same SRDF pair. SRDF/CE software can extend the Symmetrix enterprise system to support up to 64 shared quorum disk clusters per Symmetrix pair. There is no limit on the number of MNS clusters per Symmetrix pair.

CHAPTER 3

Clustering Concepts

This chapter describes the various clustering concepts for Microsoft Failover Clusters using a Cluster Enabler cluster solution and the modes of operation:

- ◆ Microsoft Failover Clusters 58
- ◆ CE geographic cluster system 61
- ◆ Application software in a cluster environment 64

Microsoft Failover Clusters

Microsoft Failover Clusters is the clustering extension to Windows Server 2008 and 2012 Enterprise and Datacenter editions. Microsoft Failover Clusters protect against failure of production server hardware or network connections. For data protection, Microsoft Failover Clusters use a protected storage subsystem. The standard failover cluster relies on RAID 1 or RAID 5 array storage to guarantee data protection.

In a typical failover cluster containing one to eight nodes, server nodes share the application workload. Typically, in a node cluster environment with n nodes, each node serves one- n th of the total number of disks and clients connected by a common SCSI bus. If one server node fails, one or several of the remaining nodes take ownership of all the disks and assume all the application workload.

Figure 12 presents a typical two-node failover cluster on Windows Server 2008 or 2012 Enterprise and Datacenter editions.

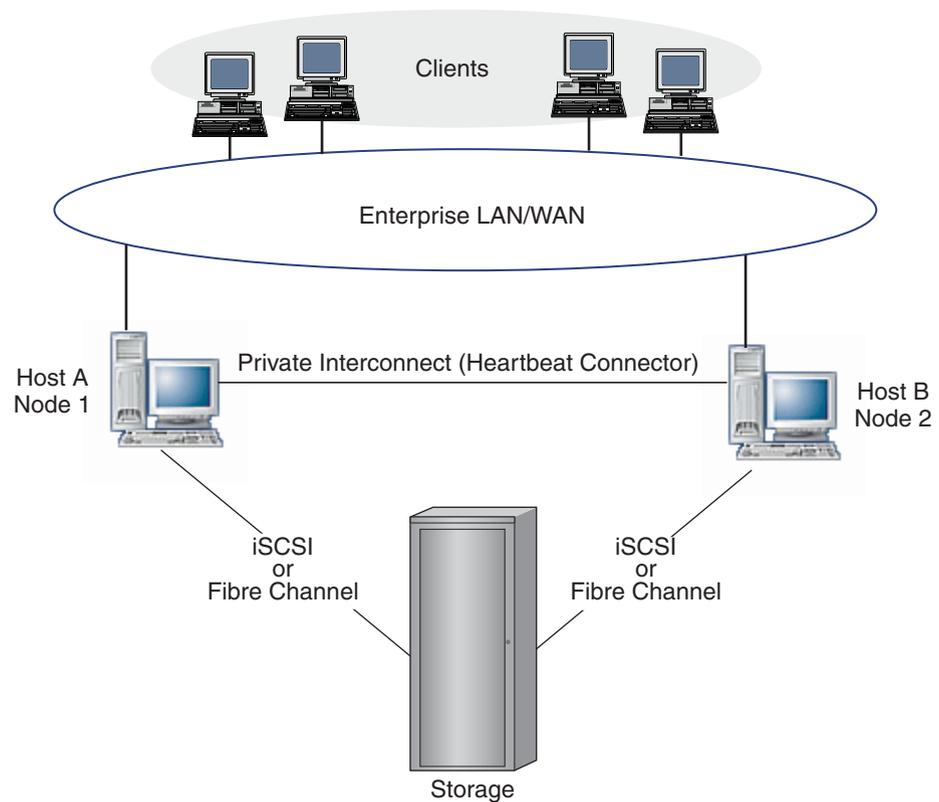


Figure 12 Typical two-node Microsoft Failover Cluster

Figure 13 presents a typical four-node Windows Server 2008 cluster.

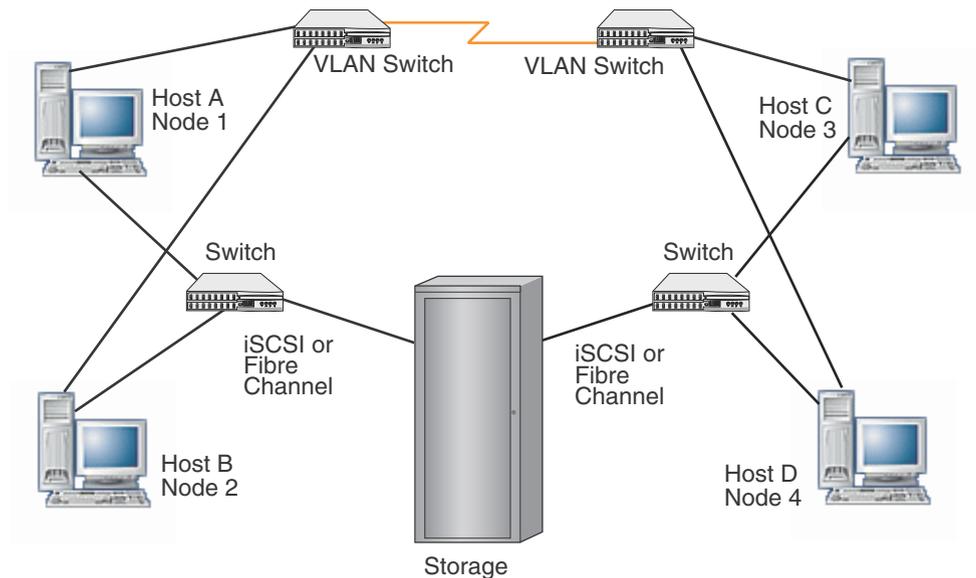


Figure 13 A typical four-node Microsoft Failover Cluster

Microsoft Failover Cluster concepts

Microsoft Failover Cluster is a loosely coupled cluster system. It is not a fault-tolerant, closely coupled system. The concept of a cluster is to take two or more independent computers and set them up to work together to provide higher availability and scalability than what you can obtain using a single system. When failure occurs in the cluster, control of a disk, or resource, moves to another cluster node. This process is called a failover. Failovers can be initiated by a number of events, including the following:

- ◆ Manual failover—The moving of resources from one server to another. Done for system load balancing or for server maintenance.
- ◆ Failover due to hardware failure—The surviving node takes over when a server, iSCSI or Fibre Channel host bus adapter (HBA) card, or network interface card (NIC) fails.
- ◆ Failover due to application failure—The failure of a virtual server or IP resource can initiate the failover.

By contrast, a fault-tolerant system uses special-purpose hardware to run multiple computers in lockstep, which provides nonstop computing with no data loss when a component failure occurs.

There are benefits and limitations to using a cluster architecture.

Benefits

Clustering provides:

- ◆ Improved availability by continuing to provide a service even during hardware or software failure.
- ◆ Increased scalability by allowing new components to be added as the system load increases.
- ◆ Simplified management of groups of systems and their applications by enabling multiple applications on multiple servers to be managed as a single system.

Limitations

Clustering cannot protect against:

- ◆ Software corruption
- ◆ Human-induced failures

Note: Protection of user data through backup (EMC business continuance volumes (BCVs), or other forms of offline data redundancy) remains vitally important to the reliable operation of mission-critical applications.

Microsoft Failover Cluster modes of operation

Microsoft Failover Cluster supports 16 node cluster for Windows Server 2008 and 64 node cluster for Windows Server 2012 Enterprise and Datacenter Editions.

Similar to the modes of operation generally discussed for Cluster Enabler, the configuration for a failover multinode cluster in a geographically distributed cluster environment is either active/passive or active/active. [“Cluster Enabler modes of operation” on page 62](#) provides an example.

Availability

Failover clusters allows active/active application operation. During normal operation, software applications can be running on both nodes. If either node fails, the applications are restarted on the remaining cluster node. This provides high availability by minimizing application downtime. Usually, it takes one to ten minutes to fail over and restart an application on a Microsoft Failover Cluster. Restart time is highly application dependent.

Scalability

In addition to availability protection, cluster technology is scalable. You can add new components to the system and run the same application (accessing the same database) on multiple nodes of a cluster to deliver increased processing power. To provide scalability, data sharing is needed.

CE geographic cluster system

Cluster Enabler provides disaster-tolerant capabilities that enable the cluster servers to be geographically separated¹. [Figure 14](#) illustrates a typical hardware configuration of a two-node CE cluster solution.

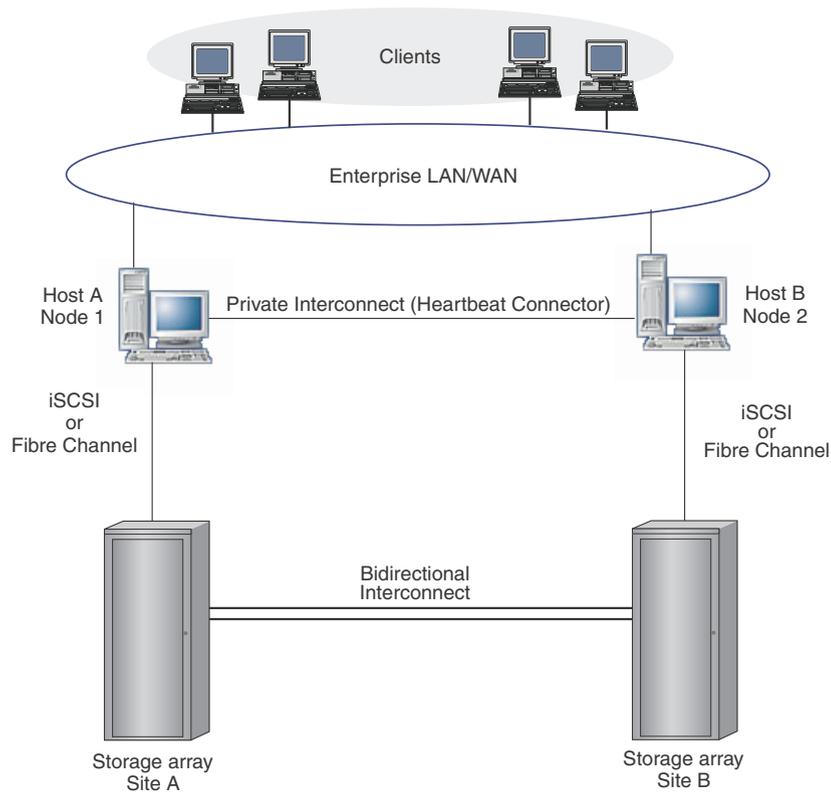


Figure 14 A geographically distributed two-node CE cluster

Network connections can provide a guaranteed maximum round-trip latency between nodes of up to 300 ms. Since many servers can connect to one storage array, it is possible to implement many clusters across this distance.

[Figure 15](#) illustrates a typical hardware configuration of a four-node cluster solution.

1. The *EMC Networked Storage Topology Guide* provides additional information regarding distance restrictions for your specific configuration.

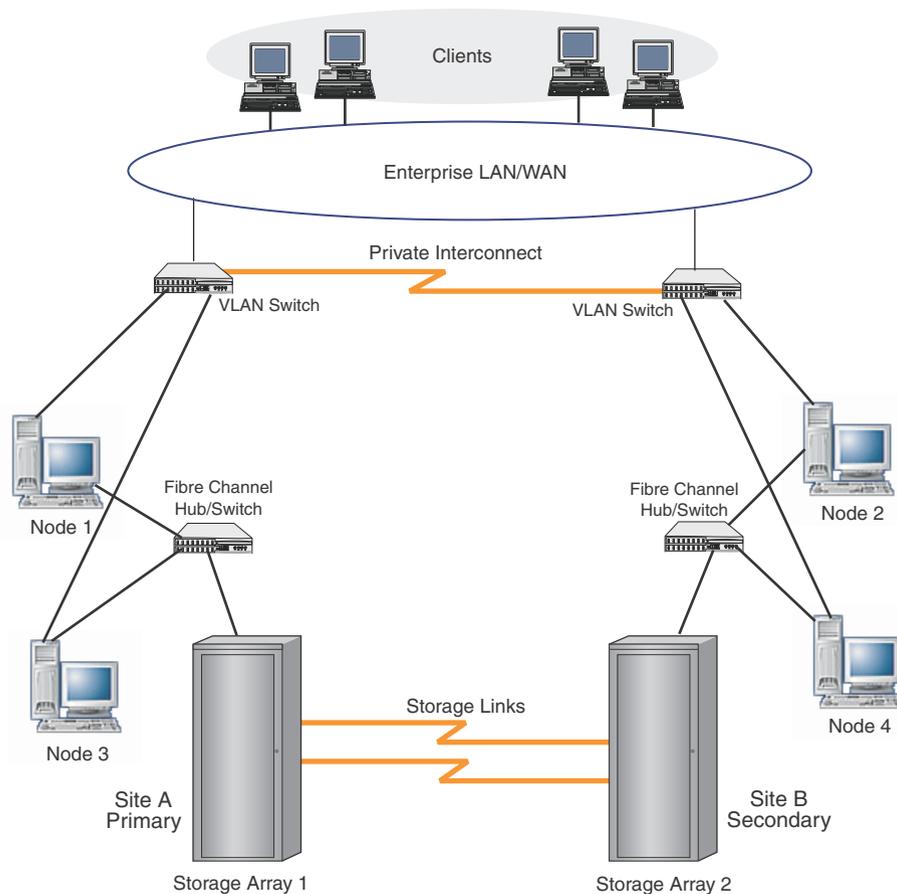


Figure 15 A geographically distributed four-node CE cluster

Cluster Enabler provides disaster-tolerant capabilities by exploiting mirroring and failover capabilities. CE allows two storage arrays to be attached using direct-connect fiber.

Note: For a RecoverPoint/CE four-node storage solution, a RecoverPoint appliance with a replication link would exist between the Fibre Channel switch and the supported storage array. Consult your RecoverPoint/CE product guide for more information.

Cluster Enabler modes of operation

Different cluster designs support different modes of operation and data-sharing mechanisms. The configuration for a CE two-node or multinode cluster in a geographically distributed cluster environment is either active/passive or active/active. EMC defines active/passive and active/active configurations as follows:

- ◆ **Active/Passive**—A cluster of two or more nodes where all processing is done on one node during normal operation, and the work is picked up by a remaining passive node (or nodes) only when a failure occurs on the active node. In a two-node configuration, half of the hardware is normally idle. When failover occurs, the application restarts with full performance.

Note: Active/passive multinode clustering provides greater flexibility than the standard active/passive Microsoft failover cluster two-node cluster by providing more options in resolving failures and load distribution after server failures. For example, in a multinode cluster, your configuration may include one or more passive (idle) servers to take over the load from other servers during a site failure, or you may distribute the load among the surviving active nodes.

- ◆ **Active/Active**—A cluster of two or more nodes where all nodes are running application software during normal operation. When a failure occurs on a node (or nodes), the work is transferred to a remaining node (or nodes) and restarted. The one or more nodes that picks up the work must then handle the processing load of both systems, and performance is usually degraded. However, all the computer hardware is used during normal operation.

Note: The terms active/active and active/passive apply to the cluster and to the applications running on the cluster. Both the cluster software and the application software must be designed for active/active operation.

Figure 16 presents a typical CE two-node two-cluster configuration.

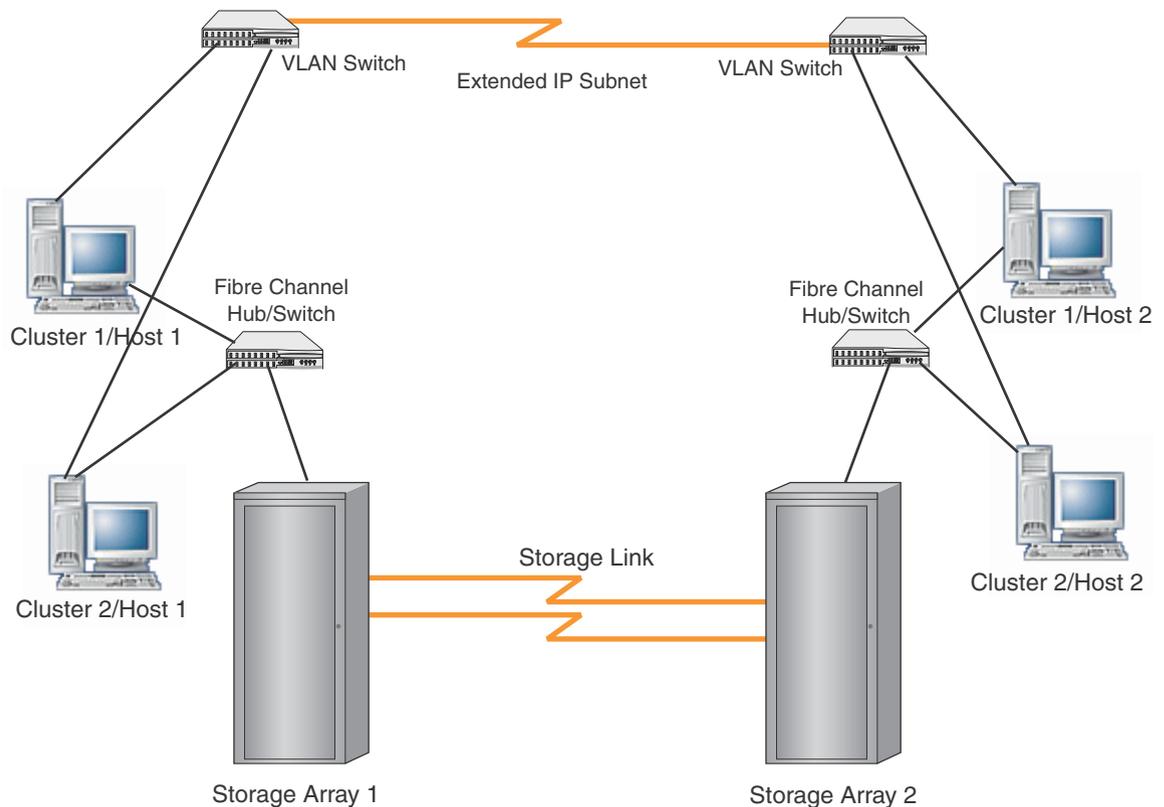


Figure 16 Two-node two-cluster CE configuration

Note: For a RecoverPoint/CE two-node, two-cluster storage solution, a RecoverPoint appliance with a replication link would exist between the Fibre Channel switch and the supported storage array. Consult your RecoverPoint/CE product guide for more information.

Cluster behavior and failover operations

Clusters are designed to overcome failures. There are several possible failure modes in a cluster configuration. Cluster Enabler protects against more failure scenarios than local clusters can. Cluster Enabler protects Microsoft Failover Clusters against disasters by providing geographically dispersed (stretched) cluster capabilities.

Cluster behavior and recovery failover operations depend on the specific scenario of failure, storage configuration, and version of Cluster Enabler plug-in module deployed.

Explanations of specific Cluster Enabler failover and recovery behavior, as well as instructions for site failure and recovery actions, are provided in each EMC Cluster Enabler plug-in module product guide.

Application software in a cluster environment

Software running on a cluster may, or may not, be cluster aware. When software is cluster aware, it provides a restart mechanism that invokes whenever the application resource is moved to another node in the cluster.

Application failover requires a restart of the application whenever failover occurs. Restart is not instantaneous. Unlike a fault-tolerant computer, a distributed cluster does not provide nonstop computing. The time that the restart takes, and the completeness of the recovery, is application dependent.

- ◆ For a transaction-oriented application (such as SQL or Exchange that contain both a database and transaction log files), the application provides a restart mechanism to recover work in progress. Usually a transaction log is used to record all work in progress. When a node fails, the information in host memory is lost, but the work can be reconstructed by applying the transaction log to the database to restart. This mechanism recovers all transactions completed before the failure. Transactions partially complete are lost and must be reentered.
- ◆ Applications such as Microsoft Word or Microsoft Excel provide a checkpoint capability. If the application experiences a failover, all work since the last disk checkpoint is lost.
- ◆ If an application has neither a database nor checkpoint capability, and also retains no information (or state) between client requests (such as a web browser or a Microsoft Outlook client), then it can fail over by reissuing the outstanding request. In this scenario, no work is lost, and no restart is needed on the server.
- ◆ If the application has neither a checkpoint nor restart capability, and it retains the state between client requests to the server, then it must be rerun from the beginning when the node it is running on fails.

CHAPTER 4

Cluster Behavior

This chapter describes SRDF/Cluster Enabler behavior in various operational modes. Unless otherwise noted, Cluster Enabler behavior is described for a standard two-node cluster:

- ◆ Cluster failover operation 66
- ◆ Response to complete site failure..... 70
- ◆ Failure behavior when using MNS with File Share Witness 73

Cluster failover operation

Clusters are designed to overcome failures. There are several possible failure modes in a cluster configuration. Cluster Enabler protects against more failure scenarios than local clusters can. Failure of an individual client affects only one user and is not discussed in this chapter. In an SRDF/CE cluster, eight types of cluster elements can fail (singly or in combination). [Figure 17 on page 66](#) provides a depiction of various cluster failures in a geographically distributed two-node SRDF/CE cluster for Symmetrix arrays.

This section describes the following:

- ◆ “SRDF/CE failover and recovery behavior” on page 67
- ◆ “SRDF/CE unique behavior” on page 68
- ◆ “Complete site failure and recovery” on page 69

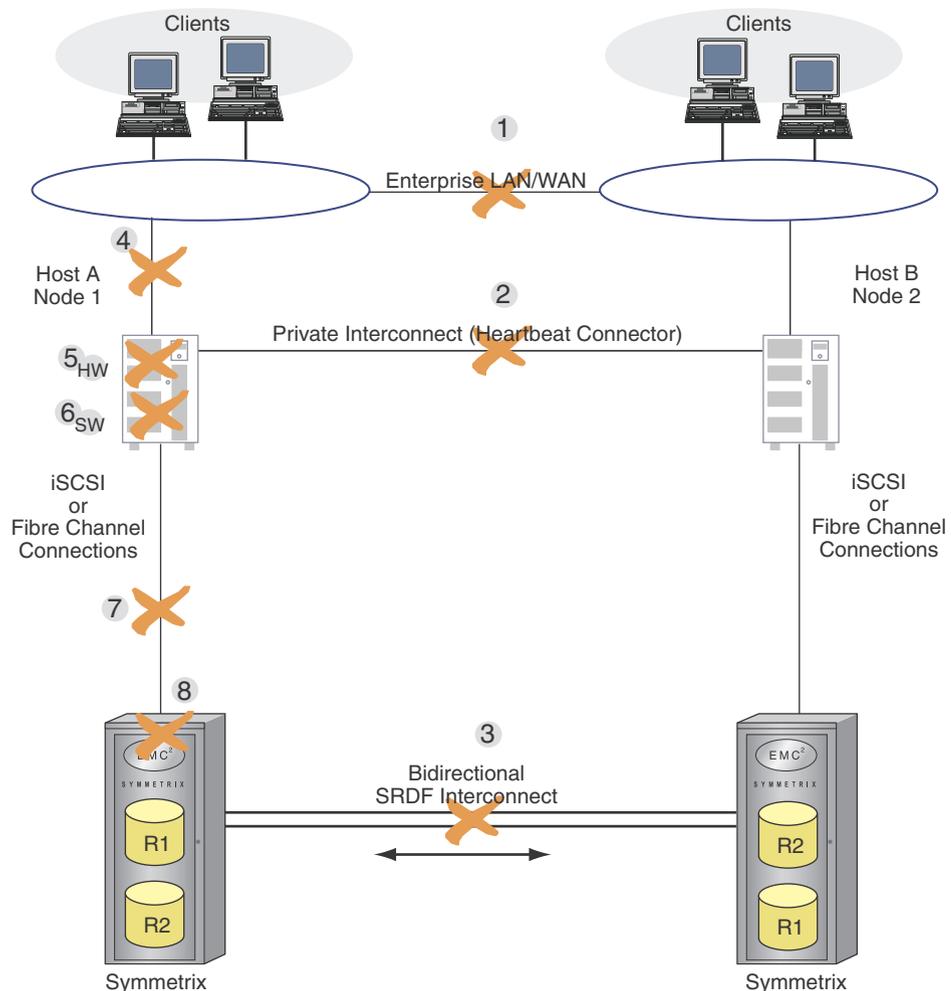


Figure 17 SRDF/Cluster Enabler failover operation

The section that follows discusses how a two-node cluster responds to various combinations of element failures. Cluster response during failure modes is similar in a three- or four-node system, but a standard two-node system is used in this section for discussion purposes. The starting condition for each of these failure scenarios is:

- ◆ Both nodes are operational.
- ◆ Node 1 (N1) owns the quorum disk for Symmetrix.
- ◆ Both the public link (internode LAN link) and the private link (heartbeat link) are configured in Microsoft (MS) failover clusters as *enabled for all network access*.

Simply stated, the failover and recovery operations Cluster Enabler provides can be divided into situations where:

- ◆ The behavior of Cluster Enabler is the same as Microsoft failover local clusters.
- ◆ The geographic separation and disaster tolerance of Cluster Enabler causes unique behavior and provides recovery alternatives.

SRDF/CE failover and recovery behavior

The following sections introduce SRDF/Cluster Enabler failover and recovery behavior common with MS failover clusters. [Figure 17 on page 66](#) shows the numbered callouts to these sections.

LAN link failure (1)

If the LAN connection between nodes fails, both servers are still available and can communicate over the heartbeat link. No failover occurs, current processing continues, and client requests from clients connected to the LAN locally continue to be serviced. Client traffic from clients connected through the LAN link fail.

Heartbeat link failure (2)

If the heartbeat link fails, MS failover clusters routes heartbeat messages across the public LAN. Operation of the cluster continues with no failover of resources.

Storage link failure (3)

[“SRDF link failure\(3\)” on page 68](#) provides a detailed explanation.

Host NIC failure (4)

The host is cut off from all clients. Processing continues uninterrupted on the other host. On the failed host, client input to that host fails, but current processing activities continue. MS failover clusters detects the NIC has failed. The isolated node takes resources offline to halt processing. The other node brings the failed resources online so application failover can occur.

Server failure (5)

If the host node hardware fails, or the operating system crashes, all heartbeat messages to the remaining node cease. The remaining node then uses the quorum disk to discover the first host has failed. The remaining node then brings the resources of the failed node online and starts the applications recovery procedures.

Application software failure (6)

If an application module fails, MS failover clusters initiates a failover to the remaining node. The Cluster Enabler resource monitor is directed to make the storage resource for the failed application available on the other node to allow application failover.

Host bus adapter failure (7)

An HBA failure is a resource failure that triggers a cluster failover operation. If both storage arrays are still running, the failover operation completes normally.

SRDF/CE unique behavior

The following sections introduce SRDF/Cluster Enabler unique behavior which is different from MS failover cluster behavior. [Figure 17 on page 66](#) shows the numbered callouts to these sections.

Storage array failure (8)

When a mirrored disk fails in a storage array, it is not visible to the host because normal operations continue with the mirror, and the failed drive is hot replaced without disturbing the host. However, if an entire storage array fails, it appears to its attached server as a resource failure indistinguishable from an HBA failure. The MS failover cluster on that server triggers a failover operation. However, because the storage array itself has failed, the remaining devices recognize that communication is lost and prevent failover from completing unless automatic failover is set as described in [“Complete site failure and recovery” on page 69](#).

SRDF link failure(3)

If the link between a Symmetrix array fails, the EMC ControlCenter® Symmetrix Manager, or the Symmetrix Management Console application notices the condition and reports an error.

The MS failover cluster server does not notice the change (because access to existing disk resources is not disturbed). However, when the SRDF/CE resource detects an SRDF link failure, the appropriate actions are taken (for example, synchronize the mirror group, swap the personality, etc.) when the SRDF link is restored. SRDF link failures or any failures in performing the restore action are noted in the Event Log, and in the SRDF/CE log.

Note: Upon link recovery, a synchronization operation will be attempted on RDF devices that are in the suspended state. Devices in a split, mixed, or other state will not automatically be synchronized. SRDF/CE detects replication link offline to online transitions when the transition time between the two states is more than one minute. For transition times less than one minute, devices may not automatically synchronize and would require user intervention to manually synchronize them.

If MS failover cluster or a user attempts to fail over or fail back a group, and there is no link available to perform that operation, the operation is not allowed. However, if there are multiple active lateral nodes and the groups in question are on that lateral side, lateral-to-lateral failover is permitted.

Note: This behavior can be overridden by enabling the Automatic Failover feature for a particular group.

Complete site failure and recovery

Local MS failover cluster

In a local MS failover cluster, if an entire site fails (such as from a flood, fire, and so forth) the entire cluster fails. By contrast, with a CE cluster, each site contains only one of the two nodes in the cluster (or only one of the n nodes in a multinode cluster).

CE cluster

A complete site failure can be caused by either a *site failure* or a *total communication failure*. Figure 18 illustrates the two types of complete site failure.

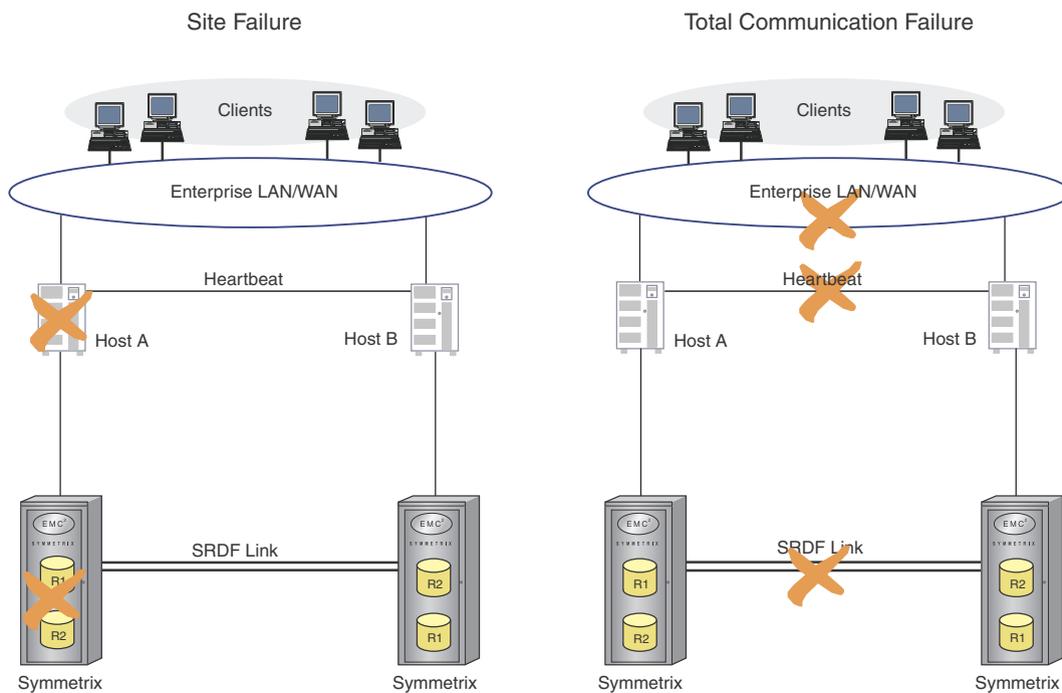


Figure 18 Types of complete site failure

Site (server and storage) failures (5+8)

Site failure occurs when the host and storage array both fail (such as from a natural disaster or human error).

Total communication failure (1+2+3)

A total communication failure can occur while the host and storage array remain operational (such as a backhoe digs up the cable conduit where all communications cables leave a building).

A total communication failure, while both nodes remain operational, is referred to as a *split-brain* condition and is a potential cause of logical data corruption. For example, if both sides assume the other is dead and begin processing new transactions against their copy of the data, two separate and unreconcilable copies of the data can be created.

Both nodes are isolated from each other, but not from local clients. It is impossible to determine if the other node is alive. No remote client processing is possible, but running processes continue.

Note: There is no way for the surviving node to determine which of these two types of failures caused the site failure.

Response to complete site failure

In Cluster Enabler, the site failure modes determine the behavior of a cluster when a failure occurs, separating the two storage arrays and suspending remote data mirroring protection.

If a complete site failure occurs, MS failover cluster on the surviving node first notices that heartbeat messages are no longer being received. MS failover cluster attempts to communicate with the other node using the LAN communication path, to see if communication is still possible.

MS failover cluster then queries the status of the disk resource and decides whether to bring the disk resources on the local node online or to set them offline. The commands to perform this query from MS failover cluster to Cluster Enabler are:

- ◆ **Is Alive?** — Determines whether a currently online resource is still healthy and can continue to be used, or whether it and all dependent cluster resources must be taken offline.
- ◆ **Online Request** — Changes the state of an offline resource to online for a failover.

Each group's failover option setting determines how Cluster Enabler responds to queries from Cluster Service. This setting must be manually configured to select the desired failover and recovery behavior.

Inappropriate user actions that cause groups to *bounce back* act differently. If you attempt to move the quorum group when the SRDF link is down, the MS failover cluster destination node terminates, and the group bounces back. Active/active configurations are obviously affected because any applications on the destination node now move. This behavior is a result of the preceding behavior.

Important: If MS failover cluster cannot write to the quorum disk when it wants to, it terminates.

The Cluster Enabler site failure mode settings are:

- ◆ **Restrict Group Movement** — In an SRDF link failure, this setting will only attempt to move disks laterally. [Figure 19](#) shows lateral and peer nodes. If the SRDF link is up, this setting has no impact.
- ◆ **Automatic Failover** — The Automatic Failover policy sets the group to allow automatic failover to another remote (peer) node in the event of an SRDF link failure.

Whenever a failure occurs such that mirrored data protection between sites is lost (for example, the SRDF link is down or a Symmetrix array is down), Cluster Enabler responds to the failure by not allowing any new disk groups to be brought online until communication with the other node has been reestablished (unless the Automatic Failover feature is set).

WARNING

Data Loss is possible for any group from Nodes 1 and 3 that is brought online with Automatic Failover if outstanding writes were not mirrored to the secondary site.

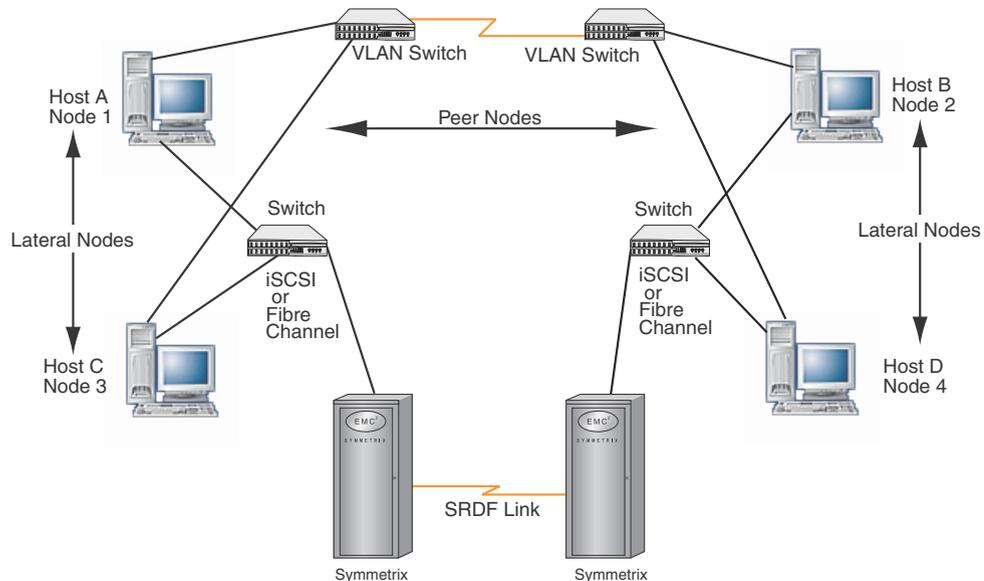


Figure 19 Lateral and peer nodes

Quorum disk-based clusters for SRDF/CE

For quorum disk-based clusters, the side that remains up with respect to a Symmetrix array is based on what node owns the quorum resource. In a site disaster, Failover clusters keep all nodes up on the side owning the quorum. All resources owned by the other side are moved to the surviving side.

In the quorum disk case, SRDF/CE monitors all nodes. If tracks (data) are not owed to the surviving side, then the move proceeds smoothly. If tracks are owed to the surviving side, then the Automatic Failover option is required to make the move successful. Therefore, if SRDF/CE detects a split-brain¹ condition during normal group failover processing, the Automatic Failover option will cause the failing site to successfully transition to the new site.

1. A total communication failure, while both nodes remain operational, is referred to as split-brain condition and is a potential cause of logical corruption. For example, if both sides assume that the other is dead and begin processing new transactions against their copy of data, two separate and unreconcilable copies of the data can be created.

Behavior override

In addition to the site failure mode settings, Cluster Enabler provides the ability to override the mode behavior and bring resources back online under user direction through the Automatic Failover feature. This enables you to decide where processing is allowed to continue.

If you determine that one site is actually down, and the other site remains operational, you can use the Automatic Failover feature to:

- ◆ Override the failure mode.
- ◆ Allow disk resources to be brought online, even though SRDF is not operating and there is no mirror protection of data.



Use the Automatic Failover feature with great care.

EMC does not recommend using the Automatic Failover feature during normal non-disaster operations.

Failure behavior when using MNS with File Share Witness

Failure behavior and recovery

In general, Cluster Enabler behaves similarly to a two-node cluster using a quorum disk.

The following example explains a four-node cluster for Majority Node Set with File Share Witness. [Figure 20](#) provides an illustrated example for Symmetrix arrays. The production nodes, Nodes 1 and 2 are at the primary site. The remote nodes, Nodes 3 and 4, are at the secondary site, and the file share node is at a third site. The cluster is configured with all the described settings.

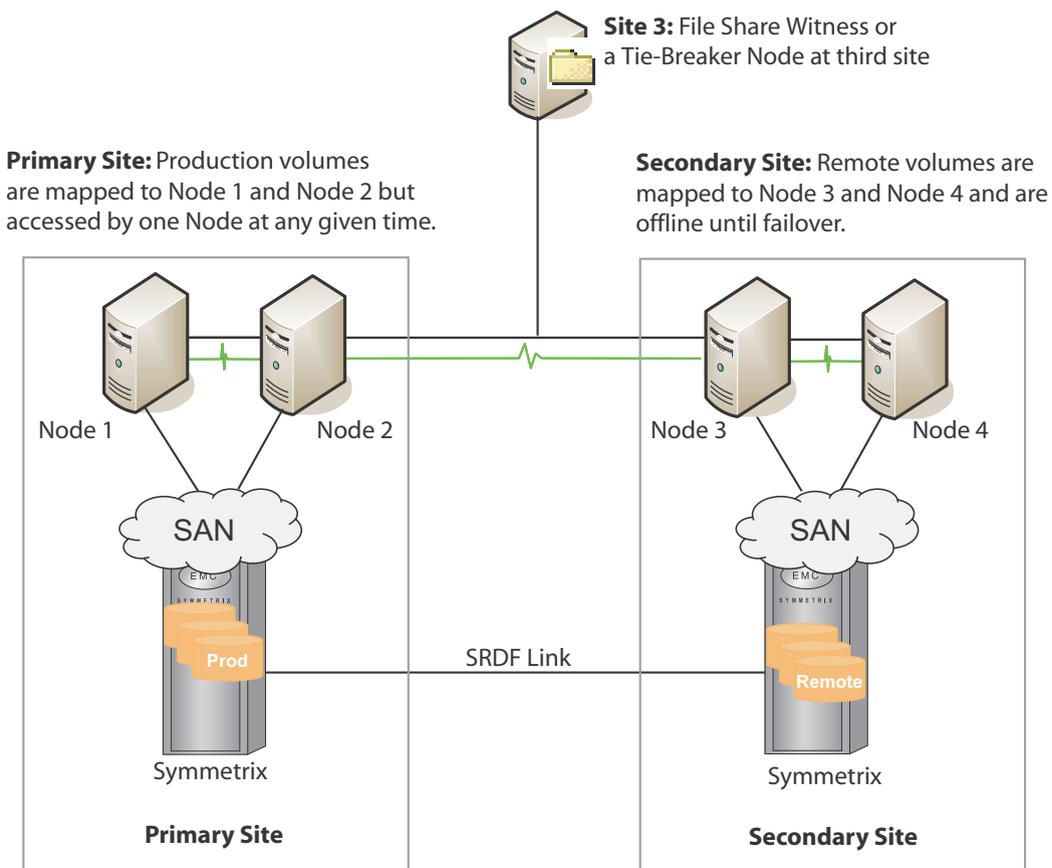


Figure 20 MNS clusters with File Share Witness

In the following examples, groups are cluster groups that contain one or more Cluster Enabler managed physical disk resources. The failover policy has been set to Restrict Group Movement.

Storage failure at primary site

- ◆ Groups on Nodes 3 and 4 remain online but cannot failover.
- ◆ Groups on Nodes 1 and 2 move to Nodes 3 and 4 but stay offline and must be brought online manually by enabling Automatic Failover.

⚠ WARNING

Data Loss is possible for any group from Node 1 and 2 that are brought online with Automatic Failover, if outstanding writes were not mirrored to the secondary site.

SRDF link failure

- ◆ Groups on Nodes 3 and 4 remain online but cannot failover.
- ◆ Groups on Nodes 1 and 2 remain online but cannot failover.
- ◆ To move a group to a different node, enable Automatic Failover on the destination node.

⚠ WARNING

Data Loss is possible for any group that is moved with Automatic Failover if outstanding writes were not mirrored.

Site failure (server and storage) at primary site

- ◆ Groups on Nodes 3 and 4 remain online but cannot failover.
- ◆ Groups on Nodes 1 and 2 move to Nodes 3 and 4 but stay offline and must be brought online manually by enabling Automatic Failover.

⚠ WARNING

Data Loss is possible for any group from Nodes 1 and 2 that are brought online with Automatic Failover if outstanding writes were not mirrored to the secondary site.

Total communication failure

- ◆ If all nodes have connectivity to the file share witness, the cluster will take two of the nodes at one site offline.
- ◆ If only one node has connectivity to the file share witness, the cluster will take the other nodes offline.
- ◆ If no nodes have connectivity to the file share witness, the entire cluster will go offline. (See Microsoft procedures for forcing an MNS cluster node online.)
- ◆ If Nodes 3 and 4 are the surviving node:
 - Groups on Nodes 3 and 4 remain online but cannot failover.
 - Groups on Nodes 1 and 2 move to Nodes 3 and 4 but stay offline and must be brought online manually by enabling Automatic Failover.

⚠ WARNING

Data Loss is possible for any group from Nodes 1 and 2 that are brought online with Automatic Failover if outstanding writes were not mirrored to the secondary site.

CHAPTER 5

SRDF/CE Installation

This chapter explains how to install and uninstall the SRDF/Cluster Enabler Plug-in module.

- ◆ Installation overview 76
- ◆ Installing the SRDF/CE plug-in module 79
- ◆ Uninstalling the SRDF/CE plug-in module 79

Installation overview

This chapter describes how to install the SRDF/CE for Microsoft Failover Clusters plug-in module on the supported Microsoft Windows Server 2008 or 2012 systems. It also describes how to uninstall the Cluster Enabler software.

Note: SRDF/CE version 4.1.4 supports software upgrades from Version 3.x and later. “[Appendix A](#)”, *Base Component Installation and Upgrade* provides instructions on how to upgrade your existing supported SRDF/CE software from Version 3.x and later to Version 4.1.4.

It is recommended that you contact EMC Customer Support for assistance if any of the following issues are applicable:

- ◆ You have applications already layered with dependencies.
- ◆ You need other devices online.
- ◆ You are not confident about installing and configuring new software within the context of Windows Server 2008 or 2012, Microsoft Failover Clusters, and Symmetrix arrays with SRDF.

Before you begin

Before you begin to install SRDF/Cluster Enabler, you should read the following installation requirements and considerations:

- ◆ The Cluster Enabler Base component is a prerequisite for the Cluster Enabler plug-ins, and therefore must be installed prior to or with the plug-ins. For instructions on installing the Base component, refer to “[Appendix A](#)”, *Base Component Installation and Upgrade*.
- ◆ The supported versions of CE that may be upgraded to Cluster Enabler Version 4.1.4 using the InstallShield wizard include only Cluster Enabler for Microsoft Failover Clusters Versions 3.x and later. The *EMC SRDF/Cluster Enabler Version 3.1 Product Guide* provides instructions for upgrading to version 3.1 from prior versions.

Note: For a clean install, all existing clusters will have to be reconfigured and any unique settings in CE will be lost.

- ◆ There are two Windows processor architectures that are supported:
 - x86
 - x64 (AMD64 and Intel EM64T)

Note: Microsoft does not support mixed architecture clusters. All nodes must have the same Windows architecture.

- ◆ The following SRDF/CE license key is required to be installed as a SYMAPI feature on every host that CE manages:

SYMAPI Feature: SRDF/CE (MSCS)/Symmetrix

- ◆ For Enginuity versions earlier than 5875, SRDF/CE requires that the appropriate license keys for Solutions Enabler SRDF/Synchronous and SRDF/Asynchronous be entered to create SRDF pairs. Refer to the *EMC Solutions Enabler Installation Guide* for information on the appropriate license keys.
- ◆ SRDF/CE support for composite groups requires that the Solutions Enabler RDF daemon (`storrdfd`) be enabled. Refer to the *EMC Solutions Enabler Installation Guide* for information on enabling the RDF daemon.
- ◆ Installation requires that all nodes first be installed with the Failover Cluster feature.
- ◆ For Failover Cluster on Windows Server 2008 or 2012, Microsoft Cluster Validate must pass all tests except storage.
- ◆ Cluster Enabler Version 4.1.4 requires that a minimum version of Solutions Enabler 7.3 first be installed.
- ◆ Cluster Enabler Version 4.1.4 requires .Net Framework 3.5 be installed.
- ◆ Upgrade scenarios where the storage is being replaced is not supported.
- ◆ Configurations where the cluster node is zoned to both local and remote storage arrays are not supported.
- ◆ For upgrade scenarios, the cluster quorum type can only be changed before or after the upgrade.

Note: For information on converting existing clusters to CE clusters, refer to [“Using the CE Configuration Wizard” on page 84.](#)

SRDF/CE license key card

A physical EMC SRDF/Cluster Enabler for Microsoft Failover Clusters License Key Card (EMC part number: 100S-001870, Revision A01) is required for product licensing. There are two ways to obtain the SRDF/CE license key card:

- ◆ When ordering new hardware from EMC that includes the SRDF/CE software option, manufacturing packs the key card along with the system.
- ◆ When purchasing SRDF/CE software by downloading the code from <http://support.EMC.com>, you must contact the licensing group at licensing@emc.com or call 800-782-4362 and follow the voice prompts (option 4). You need to provide the following order information: Sales order number, system Serial number, and License Authorization Code (LAC) from the letter emailed to you, and a license key card will be mailed to you.

For help with any errors applying license files through Solutions Enabler, contact the EMC Customer Support Center. If you are missing a LAC letter, or require further instructions on activating your licenses, contact EMC's worldwide Licensing team at licensing@emc.com or call:

- ◆ North America, Latin America, APJK, Australia, New Zealand: SVC4EMC (800-782-4362) and follow the voice prompts.
- ◆ EMEA: +353 (0) 21 4879862 and follow the voice prompts.

Getting started with Symmetrix arrays

The following steps are provided only as a high-level overview to getting started with Symmetrix arrays and SRDF/CE:

1. Prepare the Symmetrix array, RDF, and cluster node hardware.
2. Install any necessary drivers, EMC Solutions Enabler version 7.3 or later, EMC Base component and SRDF/CE Plug-in version 4.1.4 on cluster nodes.
3. Configure the Symmetrix storage and mask LUNs to all cluster nodes.
4. Ensure that all SRDF devices that the cluster uses are in a synchronized or consistent state and write-enabled.

Note: All SRDF devices must also be configured to have the `SCSI3_persist_reserv` device attribute set to enabled. The *EMC Solutions Enabler Symmetrix Array Controls CLI Product Guide* provides instructions for setting Symmetrix device attributes.

5. Map the R1 devices to all lateral nodes and the R2 devices to all peer nodes in the cluster. Symmetrix arrays that are both local and remote to a node are not supported. Reboot the nodes.
6. Open the Microsoft Disk Administrator and initialize all R1 devices. If possible format all R1 devices to NTFS format.
7. Nodes are grouped under Site. All nodes in a site shall have same devices mapped. For example, a given R1 device shall be mapped to all nodes in Site 1 and the corresponding R2 device shall be mapped to all nodes in Site 2. Add the appropriate device mappings to the rest of the nodes.
8. Ensure that all devices in a given group are of the same type (for example, either R1 or R2).
9. Verify that the SRDF link is operational by performing a failover to the R2 side. Open Microsoft Disk Administrator to check that the R2 device labels are the same as the R1 device labels. Then perform a failback and write-enabled the R1 devices on the node.
10. Create at least a single node failover cluster. Preferably create the cluster using all lateral nodes on the R1 side.
11. Ensure that all devices in a cluster group are write-enabled on the node which owns the group in the cluster.
12. Use the SRDF/CE configuration wizard to complete the cluster configuration and add R2 side nodes.

Installing the SRDF/CE plug-in module

There are two methods for installing the SRDF/CE plug-in: as a standalone module, or together with the Base component.

This section describes how to install the plug-in as a standalone module. For information on installing the plug-in together with the Base component, refer to [“Appendix A”](#), *Base Component Installation and Upgrade*.

1. Review [“Before you begin”](#) on page 76.
2. Verify that the Base component is installed, as described in [“Appendix A”](#), *Base Component Installation and Upgrade*.
3. Run the plug-in installation program and complete the steps in the installation wizard, being sure to select the same installation directory as the Base component.
4. When prompted to restart your system, click **Yes** to restart the system, or **No** to restart it at a later time.

Note: SRDF/CE Plug-in V4.1.4 installation sets the WMI MemoryPerHost quota to 1GB if the addressable physical memory size is 8GB or more.

Uninstalling the SRDF/CE plug-in module

This section explains the methods for uninstalling the SRDF/CE plug-in module from a configured cluster.

As an alternative method, you can uninstall the Base component, which will also uninstall the plug-in at the same time. For instructions, refer to [“Appendix A”](#), *Base Component Installation and Upgrade*.

Uninstalling the plug-in from some cluster nodes

To remove some cluster nodes and leave the plug-in on the remaining cluster nodes:

1. Open Microsoft Cluster Administrator.
2. Ensure no cluster resource groups are owned by the nodes you will remove. Move any owned resource groups to a different node.
3. Right-click the nodes to remove and choose **Stop Cluster Service**. Wait for the cluster service to stop on the nodes as indicated by a red **X**.
4. Right-click the nodes you want to remove and choose **Evict**. Evicting a node will uninstall the cluster service on that node and remove that node from the cluster.

Note: If CE Manager is already open, perform a refresh before running the Storage Discover Wizard.

5. After evicting nodes, open CE Manager and right-click the **cluster name**. Choose **Storage Discover** and follow through the procedure steps to complete the Storage Discover Wizard.

6. Uninstall CE from the evicted nodes. Use the **Add/Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler SRDF Plug-in.

Uninstalling the plug-in from all cluster nodes/deconfigure the cluster

To uninstall the plug-in from all nodes of the cluster and deconfigure the CE cluster., while maintaining the Microsoft Failover Cluster:

1. Move all resource groups to the nodes on one site (i.e., Site A).
2. Right-click only the nodes on the remote site (i.e., Site B) and choose **Evict**.

Note: If CE Manager is already open, perform a refresh before running the Storage Discover Wizard.

3. After evicting the nodes on the remote site, open CE Manager on a node at Site A and right-click the **cluster name**. Choose **Storage Discover** and follow through the procedure steps to complete the Storage Discover Wizard.
4. From the CE Manager, select **Deconfigure CE**.
5. Uninstall CE from all nodes. Use the **Add or Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler SRDF Plug-in.

Uninstalling the plug-in from all cluster nodes/destroy the cluster

To uninstall the plug-in from all nodes of the cluster and destroy the cluster:

1. Deconfigure the cluster according to steps 1 through 4 in [“Uninstalling the plug-in from all cluster nodes/deconfigure the cluster” on page 80](#).
2. Destroy the cluster using Microsoft Cluster Administrator.
3. Uninstall CE from all nodes. **Use the Add or Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler SRDF Plug-in.

Note: Uninstallation of the SRDF/CE Plug-in does not revert the WMI MemoryPerHost quota value. This value must be changed manually as desired.

CHAPTER 6

Using Cluster Enabler Manager

This chapter provides instructions for using the SRDF/Cluster Enabler Manager graphical user interface.

- ◆ Getting started using the CE Manager..... 82
- ◆ Using the CE Configuration Wizard 84
- ◆ Managing a CE cluster 87
- ◆ Managing a CE cluster group 91
- ◆ Storage component..... 97
- ◆ Viewing information 98
- ◆ Restore and recovery operations 106
- ◆ Configuring a custom resource 111

Getting started using the CE Manager

The Cluster Enabler (CE) Manager GUI (graphic user interface) allows you to configure your Microsoft Failover Clusters for disaster recovery protection. The CE Manager allows you to set up and configure disk-based resources to automatically move geographically dispersed resource groups back and forth.

The CE Manager provides several wizards to assist you in completing various cluster tasks. The first step towards managing disaster recovery for distributed failover clusters is to run the Configuration Wizard to configure a CE cluster.

The Cluster Enabler Manager window

The CE Manager window shown in [Figure 21](#) contains a menu bar, two views, and a navigation tree. After cluster configuration, the navigation tree can be expanded to show four separate components: Groups, Storage, Sites, and Nodes.

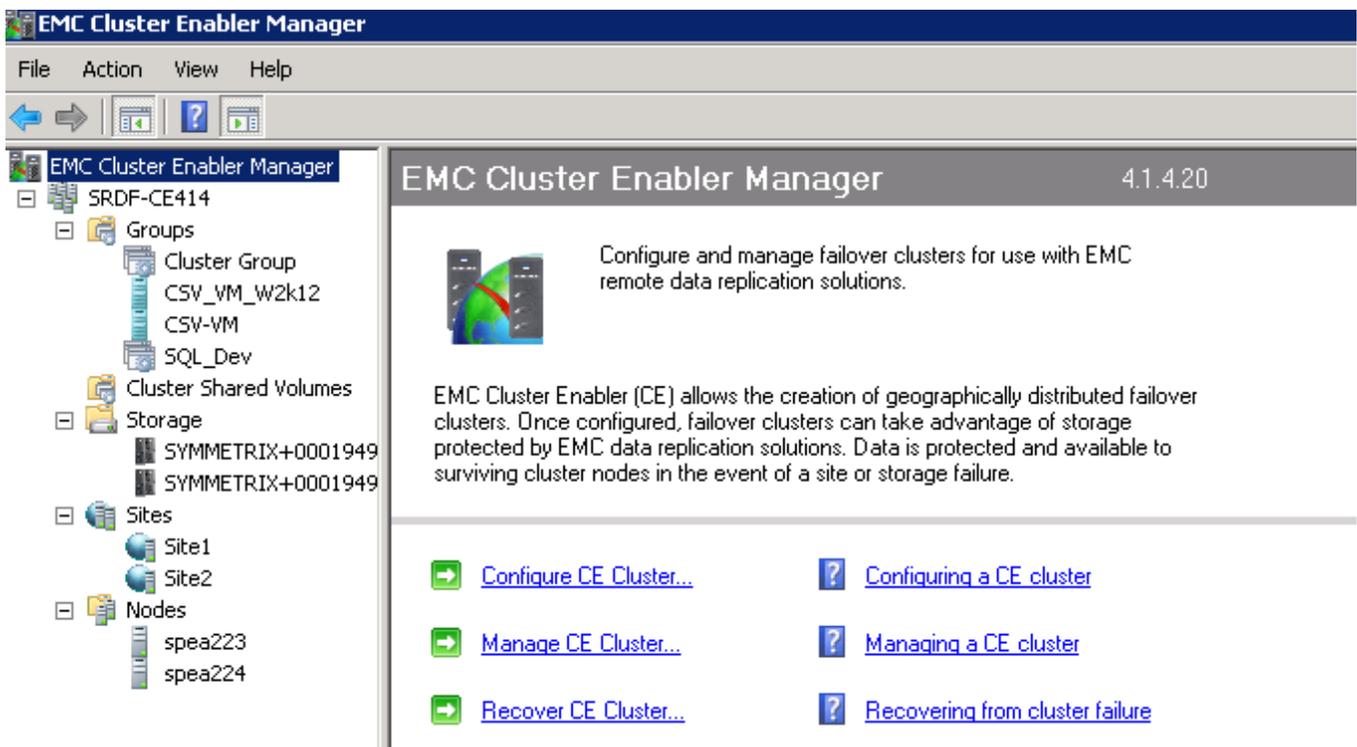


Figure 21 Cluster Enabler Manager window

Cluster Enabler wizards

Wizards are a series of dialog boxes that step you through the completion of a complex task. The Cluster Enabler Manager provides several wizards, as follows:

Configuration Wizard

The Configuration Wizard is used to configure a CE cluster. The configuration process is the first step towards managing disaster recovery for distributed failover clusters. The Configuration Wizard will step you through the process of configuring your failover cluster for management with CE. [“Using the CE Configuration Wizard” on page 84](#) provides detailed instructions for using the wizard.

Storage Discover Wizard

The Storage Discover Wizard automatically discovers and sets up the attached storage. The storage discovery process should be performed after any changes have been made to the storage configuration. [“Storage Discover Wizard” on page 87](#) provides detailed instructions for using the wizard.

Update Mirror Pairs Wizard

The Update Mirror Pairs Wizard steps you through the process of discovering storage, updating the storage configuration, validating the storage groups, and setting up the storage group definitions in the cluster properties database to update the mirrored pairs in a cluster. [“Update Mirrored Pairs Wizard” on page 88](#) provides detailed instructions for using the wizard.

Change Quorum Wizard

The Change Quorum Wizard steps you through the process of changing a cluster's quorum model type. [“Change Quorum Model Wizard” on page 88](#) provides detailed instructions for using the wizard.

Create Group Wizard

The Create Group Wizard steps you through the process of creating a CE Group, adding devices and selecting a group policy. [“Create Group Wizard” on page 91](#) provides detailed instructions for using the wizard.

Modify Group Wizard

The Modify Group Wizard steps you through the process of adding or removing devices in a CE group. [“Modify Group Wizard” on page 93](#) provides detailed instructions for using the wizard.

Recover CE Cluster Wizard

The [“Recover CE Cluster Wizard” on page 109](#) step you through the process of recovering a shared quorum cluster.

Using the CE Configuration Wizard

Cluster Enabler provides a wizard for configuring a CE cluster. The configuration process is the first step towards managing disaster recovery for distributed failover clusters. The Configuration Wizard steps you through the process of configuring your failover cluster for management with CE.

If any of the steps in wizard configuration process fail, the wizard displays a list of the specific errors for each node on a Summary page. Note each error to be corrected and click **Finish** to exit the wizard. After the listed summary problems have been fixed, launch the configuration wizard again to configure the CE cluster.

Note: Whether running Windows Server 2008 or 2012, the applicable Microsoft Failover Clusters must be installed on at least one node prior to configuring a cluster.

Follow these steps to configure a CE cluster using the Configuration Wizard:

1. Select the **EMC Cluster Enabler** icon from the Navigation Tree and click the **Configure CE Cluster** link in the center pane. This opens the first page of the Configuration Wizard. The Configuration Wizard can also be launched using the right-click or action menus.
2. The **Enter cluster name page** appears. Enter a Cluster Name or Node Name in the space provided and click **Configure**. Select a cluster name from the list and click **OK**, then click **Add**. Click **Configure**. If you do not enter a name and click **Configure**, the wizard automatically detects the current clusters on the server and continues. [Figure 22 on page 84](#) shows the first page of the Configuration Wizard.

Note: CE configurations where devices in a consistency group are mixed between concurrent, cascaded and point-to-point configurations are not supported. If CE detects this configuration, an error message displays and the group is excluded from the CE configuration.

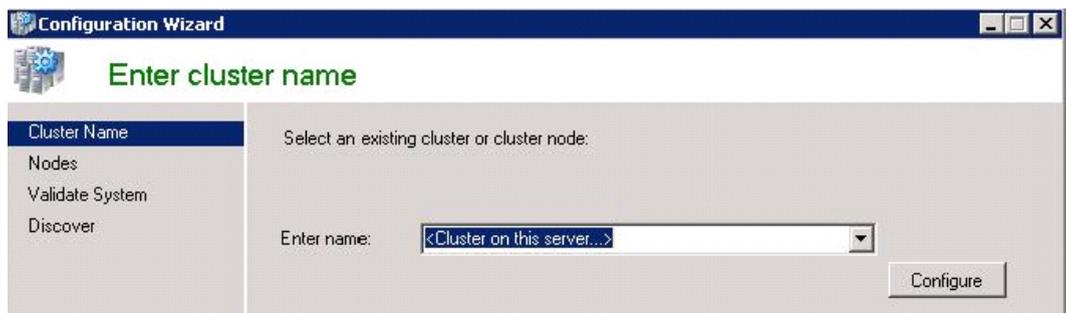


Figure 22 CE Manager Configuration Wizard

3. The Current Nodes page appears listing the current nodes in the cluster. To add a node, enter the node name and click **Add**. If you do not know the node name, you can click **Browse** to browse an active directory of computers. Select a computer name from the list and click **OK**, then click **Add**. Click **Next**.

4. The Validating System Setup process begins. This automated step validates the system configuration by checking that the appropriate versions of Solution Enabler, Cluster Enabler, and Microsoft Failover Clusters are installed and configured. Upon the Validation Complete notification, click **Next**.

Note: If the system validation process fails, the wizard lists the validation errors for each node on the Summary page. Note each error to be corrected and click **Finish** to exit the wizard.

5. The Storage Discovery process begins. This automated step performs a storage discovery for each cluster node to identify the locally-attached and remotely-attached storage. Upon the Discover Completed notification, click **Next**.

Note: If the storage discovery process fails, the wizard lists the storage discovery errors for each node on the Summary page. Note each error to be corrected and click **Finish** to exit the wizard.

6. The Storage Setup process begins. This automated step performs a storage setup for each cluster node. Upon Setup of Storage Configuration Completed, click **Next**.

Note: If the storage setup process fails, the wizard lists the storage setup errors for each node on the Summary page. Note each error to be corrected and click **Finish** to exit the wizard.

Note: In the case of Node and Disk Majority or Disk Only, the storage setup process fails with a warning message if the cluster disk's topology type is not Point to Point and the replication mode is not Synchronous. If you are re-configuring an existing CE cluster, CE will manage the Cluster upon clicking the Finish button. Change the quorum configuration using the Change Quorum Wizard and relaunch the Configuration Wizard. If the Cluster is being configured for the first time, change the quorum disk using the Microsoft Failover Manager and relaunch the CE Configuration Wizard for CE to manage the cluster.

7. The Validating Groups process begins. This automated step performs a group validation for each converted failover cluster group. Upon Validated Groups, click **Next**.

Note: If the validating groups process fails, the wizard lists the validation errors for each node on the Summary page. Note each error to be corrected and click **Finish** to exit the wizard.

8. The Summary page appears. Upon Configuration Wizard Completed Successfully, click **Finish** to exit the wizard.

9. After exiting the CE Configuration Wizard, Cluster Enabler connects to the newly configured cluster. Once connected to the cluster, you will notice that the configured cluster node is now visible in the navigation tree, located in the left pane.

10. Double-click the **cluster** icon to expand the cluster and view the following folders: Groups, Storage, Sites, and Nodes. You are now ready to begin managing your cluster. [Figure 23 on page 86](#) shows an example view of the expanded CE Manager navigation tree.

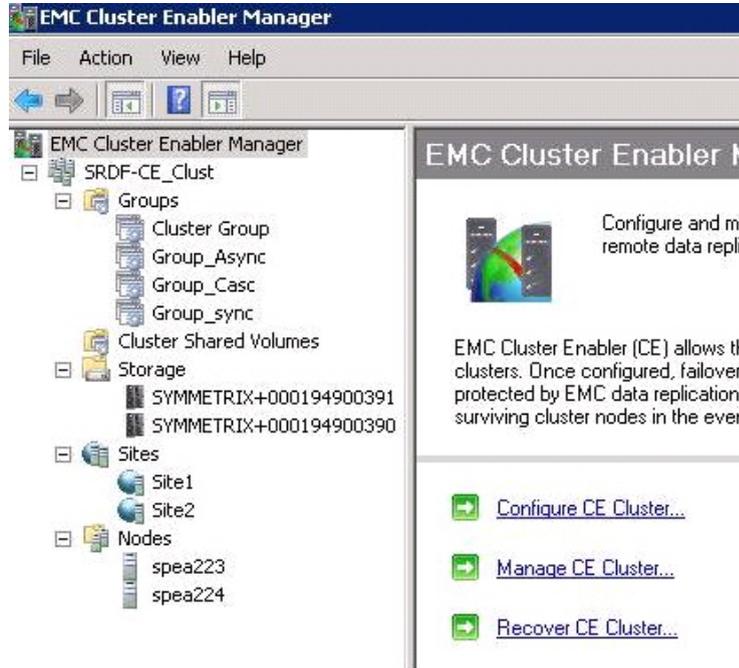


Figure 23 CE Manager expanded navigation tree

Adding nodes

Adding new nodes is also accomplished through using the CE Configuration Wizard. The CE Configuration Wizard steps you through the process of adding a cluster node for management with CE. New nodes must be added using CE Manager and not Microsoft Failover Cluster. Step 3 in the [“Using the CE Configuration Wizard” on page 84](#) provides instructions.

Managing a CE cluster

Once your CE cluster has been configured using the CE Configuration Wizard, you are ready to begin managing the cluster. Even after you exit Cluster Enabler and close the application, your cluster will remain configured unless you perform a deconfigure or delete action on the cluster.

Follow these instructions to begin managing your cluster.

1. Select the **EMC Cluster Enabler** icon from the Navigation Console Tree and click the **Manage CE Cluster** link in the center pane. The Manage CE Cluster option can also be launched by using the right-click or Action menus.
2. Enter the cluster name in the selection box. If you do not enter a name, the default will automatically connect to a cluster accessible on the server. Click **OK**.
3. Once connected to the cluster, you will notice that the configured cluster node is now visible in the navigation tree located in the left pane. Double-click the **cluster icon** to expand the cluster and view the following folders: Groups, Storage, Sites and Nodes. You are now ready to begin managing your cluster.

Cluster Enabler Manager allows you to perform the following cluster management actions on configured clusters. When a particular action is accessed, the appropriate wizard process is launched to assist you in completing the task.

Note: When a CE cluster is managed remotely from a server (management station) which is not part of that cluster, ensure that the installed CE version on all the cluster nodes is the same as on the management station.

Storage Discover Wizard

The Storage Discover Wizard helps you to automatically discover and setup the attached storage. The storage discovery process should be performed after any changes have been made to the storage configuration.

Follow these steps to automatically discover and set up your attached storage using the Storage Discover Wizard:

1. Select the **Cluster** icon from the Navigation Console Tree and select **Action** and **Storage Discover** from the menu bar. This opens the first page of the Storage Discover Wizard. The Storage Discover Wizard can also be launched by using the right-click or Action menus.
2. The Storage Discovery page appears. Upon the Discover Completed notification, click **Next**.

Note: If any storage discovery process fails, the wizard will list the discovery errors for each node on the Summary page. Note each error to be corrected and click **Finish** to exit the wizard.

3. The Storage Setup page appears. Upon the Set up of Storage Configuration Completed notification, click **Next**.
4. The Summary page appears. Upon the Discovered all Nodes notification, click **Finish** to exit the wizard.

5. Cluster Enabler will then refresh the CE cluster to reflect any storage changes.

Update Mirrored Pairs Wizard

The Update Mirrored Pairs Wizard helps you update the mirrored pairs in a cluster. This wizard steps you through the various processes of discovering storage, updating the storage configuration, validating the storage groups, and setting up the storage group definitions in the cluster properties database to update the mirrored pairs in a cluster.

Follow these steps to update the mirrored pairs in a cluster using the Update Mirror Pairs Wizard:

1. Failover CE groups to an R1 node before modifying their R1-R2 relationship.
2. Shutdown the R2 (passive) Node.
3. From R1 node, modify the R1-R2 pairing with a new R2 device.
4. Establish the pairing and wait for the RDF state to be synchronized.
5. Bring up the R2 (passive) node.
6. Manage the cluster from CE manager on the R1 node and perform update mirror pair operation.
7. Select the **Cluster** icon in the navigation tree and select **Action, More Actions...** and **Update Mirror Pairs** from the menu bar. The Update Mirror Pairs Wizard can also be launched using the right-click or Action menus.
8. The first page of the Update Mirror Pairs Wizard opens and begins the Storage discovery process. Upon the Discover Complete notification, click **Next**.
9. The Storage setup process begins setting up the storage configuration. Upon the Setup of Storage Configuration Completed notification, click **Next**.
10. The Validating Groups process begins validating each group in the cluster. Upon the Validated Groups notification, click **Next**.
11. The Updating Storage Mirror Relationships process begins updating the mirrored pairs in the groups. Upon the Update mirror pairs for groups notification, click **Next**.
12. The Summary page appears. Upon Update Mirror Pairs Completed Successfully, click **Finish** to exit the wizard.

Once the Update Mirror Pairs Wizard completes successfully, the internal configuration database updates to reflect the new R1/R2 relationship. Once updated, groups can be failed over between these Symmetrix arrays with the new R2 pairs.

Change Quorum Model Wizard

The Change Quorum Wizard changes the quorum model of a cluster. This wizard will step you through the various processes of changing a cluster's quorum model type.

For Windows Server 2008 and 2012, SRDF/Cluster Enabler allows you to change the cluster model type between No Majority: Disk Only, Node Majority, Node and Disk Majority, and Node and File Share Majority. A descriptive list of all quorum model types is provided in [“Supported model type descriptions” on page 24](#).

Using the Change Quorum Wizard

Once your Microsoft cluster has been configured as a CE cluster, you must use this wizard for all quorum model changes. If your configured CE clusters are No Majority: Disk Only model type, you can use this wizard to change the selected quorum disk. You can also use this wizard to change the file share for configured CE clusters of Node and File Share Majority model types.

Note: To change the quorum model to "Node and File Share Majority" in Windows Server 2008, you must first update the FileShare permissions to add the Cluster Name and allow "Change" and "Read" permissions for the file share. Your windows documentation provides instructions on changing permissions for FileShare.

Figure 24 shows the first page of the Change Quorum Wizard for Windows Server 2008.

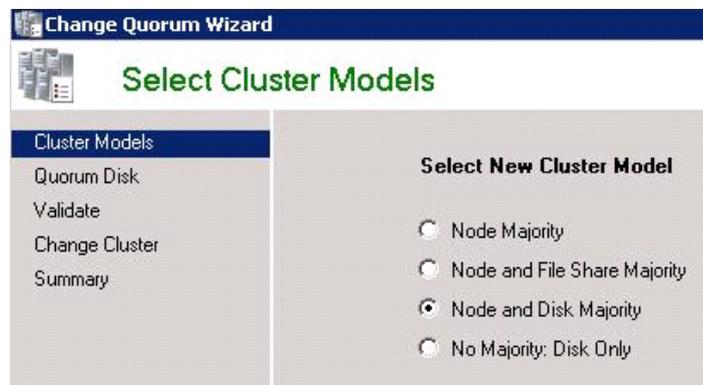


Figure 24 Windows Server 2008 supported quorum models

Change Quorum Model example

The following example steps through the processes of changing a Node Majority cluster to a Majority:Disk Only cluster for Windows Server 2008.

Follow these steps to change the quorum model type of a cluster using the Change Quorum Wizard:

1. Select the **Cluster** icon in the navigation Console tree and select **Action, More Actions...** and **Change Quorum Model** from the menu bar. The Change Quorum Model Wizard can also be launched by using the right-click or Action menu.
2. Cluster Enabler automatically begins by reading the storage configuration.
3. Once the storage configuration has been read, the first page of the Change Quorum Wizard opens. Select the **New Cluster Model** radio button for the model you want to convert the cluster to. In this example, we are changing the cluster model from a Node Majority cluster to a No Majority:Disk Only cluster. Click **Next**.
4. The Select Quorum Disk page opens. Select the quorum disk that you want to use for the quorum disk-based cluster. Click **Next**.

Note: If the selected disk does not have a topology type of Point to Point and a replication mode of synchronous, the Next button is disabled and an error message is displayed. You will not be able to change the quorum disk until the supported disk type is selected.

5. The Select Cluster Number page appears. The wizard will automatically generate a list all of the available cluster numbers. From the Select a Cluster Number scroll box, select the **Cluster Number** that you want to use for the cluster. Click **Next**.

Note: Check the **Show All Cluster Numbers** box to view all of the cluster numbers both used and unused for the system. Do not select a number that is already used.

6. The Validate Cluster Model process automatically begins validating the chosen cluster model. Upon the Validation of Cluster Model Successfully notification, click **Next**.
7. The Change Cluster Model process automatically begins changing the cluster settings to the new model. Upon Change Cluster Model Successfully, click **Next**.
8. The Summary page appears. Upon the Changed Cluster Model Successfully notification, click **Finish** to exit the wizard.
9. In the Cluster Enabler Manager, select the **Cluster** icon, notice that the Cluster Type is now No Majority:Disk Only.

Managing a CE cluster group

Cluster Enabler Manager provides several group actions for managing CE cluster groups. There are two automation wizards available for groups, the Create Group Wizard and the Modify Group Wizard. The Create Group Wizard steps you through the process of creating a CE cluster group. The Modify Group Wizard allows you to edit an existing cluster group by adding or removing devices to and from the group. The group action features also allow you to deconfigure a CE group to convert it back to a cluster group or to delete a group.

Create Group Wizard

The Create Group Wizard steps you through the process of creating a CE Group, adding devices and selecting a group policy.

Follow these steps to create a CE Group using the Create Group Wizard:

1. Select the **Groups** icon from the Navigation Console Tree and select **Action** and **Create Group** from the menu bar. This begins the process of reading the storage configuration. After the storage has been read, the first page of the Create Group Wizard opens. The Create Group Wizard can also be launched using the right-click or Action menus.

Note: A mirrored pair needs to be present on the array before attempting to create a group. Run the Storage Discover Wizard to detect a newly created mirrored pair by right-clicking on the **cluster name** or clicking the **Discover** button in the Select Devices page of the Create Group Wizard.

2. The Enter Group Name page appears. Enter a unique **Group Name** in the space provided and click **Create**. Click **Cancel** to abort the operation and close the wizard.
3. The next wizard page prompts you to select devices for inclusion in the new group. Symmetrix RA group pairs and the devices contained within the RA group are shown in a tree view. Select the desired devices from the list shown by clicking in the **select boxes**. Selected devices are identified by the checked box. Selecting the RA group, automatically selects all devices in that group. After your selections have been made, click **Next**. [Figure 25 on page 92](#) shows the wizard page for selecting devices.

By default, all available configured Symmetrix storage is shown in collapsed view (Collapse All). The tree view can be expanded by selecting **Expand All**. There are three types of devices that can be displayed by checking the selection boxes: Async, Cascaded, and Concurrent. For example, selecting the Async checkbox displays all SRDF asynchronous capable devices within in the same RA group, mapped to the nodes. An error message displays if the selected type of devices are used up or not available. If you select devices from a single Symmetrix RA group, a device group will be created. If you select devices from multiple Symmetrix RA groups, a composite group will be created. Any devices other than the replication mode synchronous or asynchronous will not be listed for device selection.

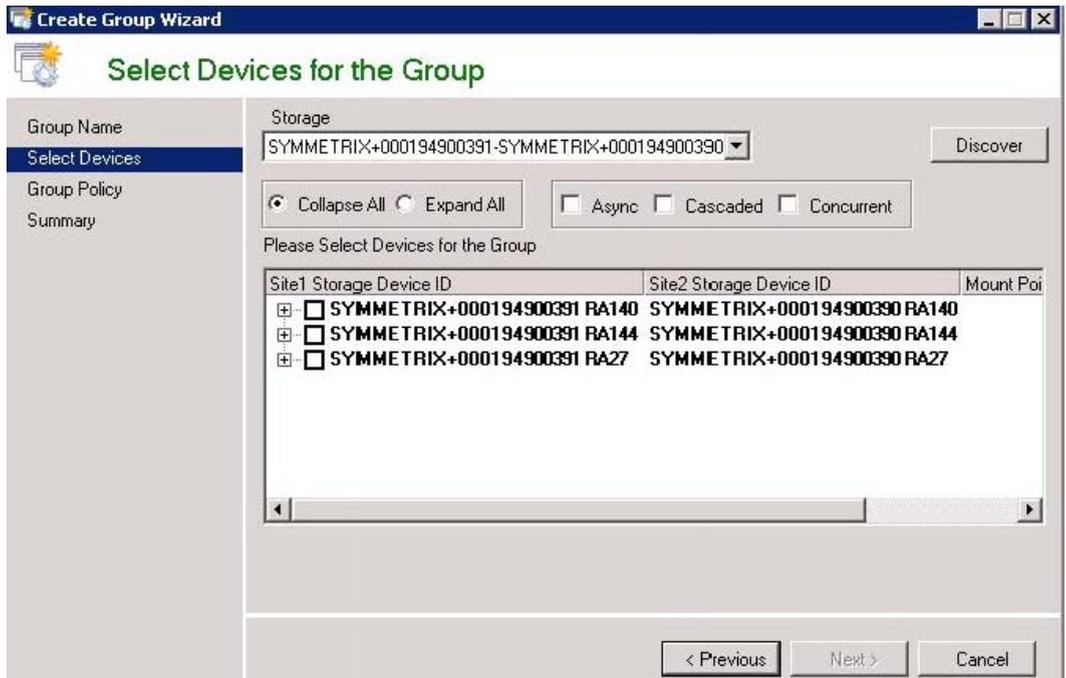


Figure 25 Create Group Wizard, Select Devices for the Group

Note: CE configurations where devices in a consistency group are mixed between concurrent, cascaded and point-to-point configurations are not supported.

Note: Cluster Enabler will prevent Composite Group creation if configured on VMAX 10K/VMAXe arrays.

- The Select Group Policy page appears. From the pull-down menu, select your desired policy for the group. You can select either the **Restrict Group Movement** or **Automatic Failover**. Once selected, click **Next**. [Figure 26](#) shows an example of the Select Group Policy wizard page.

The **Restrict Group Movement** selection restricts the group from failing over to a peer node. In an SRDF link failure, this setting will only attempt to move disk laterally. If the link is up, this setting has no impact.

The **Automatic Failover** policy sets the group to automatically failover to another node in the event of a node or network failure.

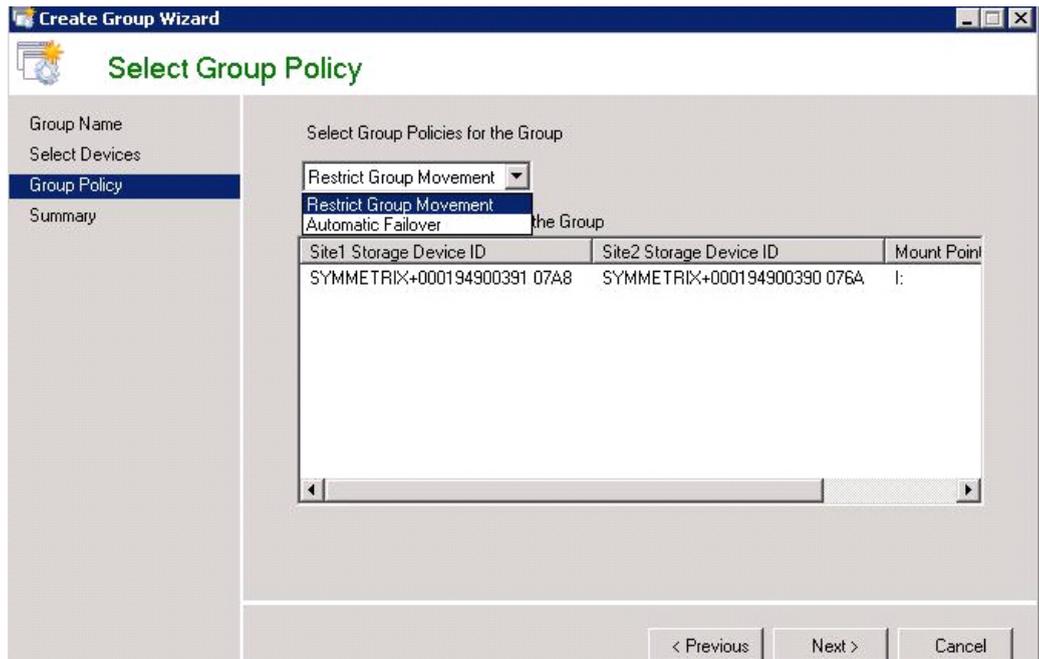


Figure 26 Create Group Wizard, Select Group Policy

5. The Summary page appears. Upon Group Created Successfully, click **Finish** to exit the wizard.

Cluster Enabler automatically begins refreshing the CE Cluster. Upon completion of the refresh, you should see the group that you created listed under Groups. If you do not see the newly created group, select **Action** and **Refresh** from the menu bar. The Refresh action can also be accessed from the right-click or Action menus.

Note: When using the CE Create Group wizard, if you select the Asynchronous check box to convert a synchronous capable device to an asynchronous capable device (or vis versa) and the appropriate SRDF/Asynchronous or SRDF/Synchronous license does not exist, then group creation fails with one of the following messages:

The following Array Licenses are required on Symmetrix: Sync
 The following Array Licenses are required on Symmetrix: Async
 The following Solutions Enabler Licenses are required on node:
 < nodename>

Modify Group Wizard

The Modify Group Wizard steps you through the process of adding or removing devices in a CE group.

Follow these steps to add or remove devices from a CE group using the Modify Group Wizard:

1. Select the **Groups** icon in the navigation tree and select **Action** and **Modify Group** from the menu bar. This begins the Storage Synchronization process. After the storage has finished synchronizing, the first page of the Modify Group Wizard opens. The Modify Group Wizard can also be launched using the Right-click or Action menus.

Note: A mirrored pair needs to be present on the array before attempting to modify a group. Run the Storage Discover Wizard to detect a newly created mirrored pair by right-clicking on the **cluster name** or clicking the **Discover** button in the Select Devices page of the Modify group wizard.

2. From the Select Devices page, select the Action from the pull-down menu for either **Add Devices** or **Delete Devices**. Depending on your selection, a list of available devices that can be added or removed displays. Symmetrix RA group pairs and the devices contained within the RA group are shown in a tree view. By default, the RA Groups are shown in collapsed view (Collapse All). The tree view can be expanded by selecting **Expand All**.

Select the desired devices from the list shown by clicking in the select boxes. Selected devices are identified by the checked box. Selecting the RA group, automatically selects all devices in that group. After your selections have been made click **Next**. [Figure 27 on page 94](#) shows this wizard page.

[“Modify CE group notes” on page 94](#) provides additional information on modifying CE groups.

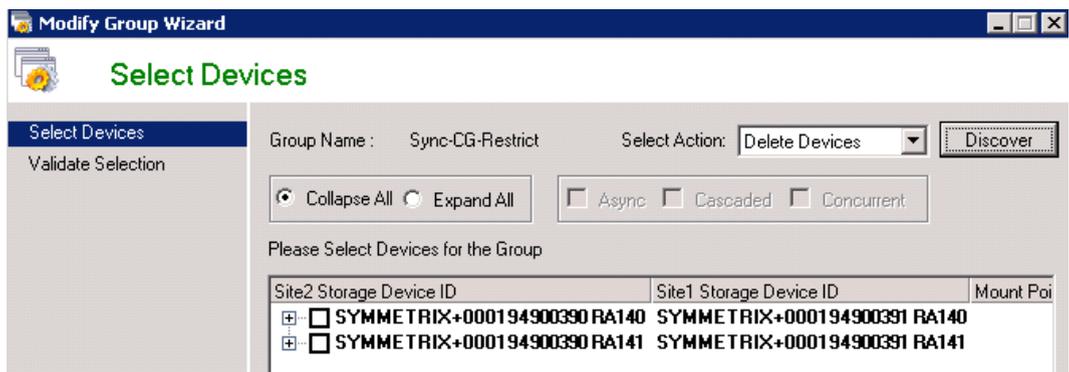


Figure 27 Modify Group Wizard, Select Devices

Modify CE group notes

This section provides additional information for modifying CE groups.

- Selecting Add Devices displays ungrouped devices from all RA group pairs.
- Selecting Delete Devices for a composite group displays all devices under each RA group pair in this group.
- If deleting devices in a composite group, and you delete all devices from within one RA group, the modified group automatically converts to a device group and RDF consistency is disabled.
- If the group is a device group, and you add devices from another RA group, the group automatically converts to a composite group and RDF consistency is enabled. A warning message displays in this case.

- CE configurations where devices in a consistency group are mixed between concurrent, cascaded and point-to-point configurations are not supported.
 - Any devices other than the replication mode synchronous or asynchronous will not be listed for device selection.
 - Cluster Enabler will prevent Composite Group creation if configured on VMAX 10K/VMAXe arrays.
3. The Validate Selection page appears, click **Next** to validate your selection or click **Cancel** to abort the action.
 4. The Summary page appears. Upon Group Modified Successfully, click **Finish** to exit the wizard.

Cluster Enabler automatically begins refreshing the CE cluster. Upon completion of the refresh, you should see the updated group information reflecting the devices added or deleted. If you do not see the updated group information, select **Action** and **Refresh** from the menu bar. The Refresh action can also be accessed from the right-click or Action menus.

Note: When using the Modify Group wizard, if you select the Asynchronous check box to convert a synchronous capable device to an asynchronous capable device (or vis versa) and the appropriate SRDF/Asynchronous or SRDF/Synchronous license does not exist, then group creation fails with one of the following messages:

The following Array Licenses are required on Symmetrix: Sync
 The following Array Licenses are required on Symmetrix: Async
 The following Solutions Enabler Licenses are required on node:
 < nodename >

Configure a CE Group

To configure a single group using the Configure CE Group option, select a group listed under the Groups icon located in the Navigation Tree and select **Action** and **Configure CE Group** from the menu bar. The Configure CE Group action can also be accessed by using the right-click menu from the group name. A dialog box pop-up appears asking you to confirm the action. Click **Yes** to configure the Microsoft Failover Cluster group to a CE group or **No** to abort the action. The Configure CE Group action is enabled for non-converted Microsoft Failover Cluster groups and disabled for all existing CE Groups.

If the group is configured as a CE Group, CE failover support to the remote nodes will be added. The CE resource is added, but the Microsoft Failover Cluster physical disk resources are not changed. Converting a Microsoft Failover cluster group to a CE group enables the Delete, Modify Group, and Deconfigure Group CE operations.

When configuring a VM Group through CE, if the underlying CSV disk group is not configured, it will be configured first before the VM Group is configured. You can also configure the CSV disk group by itself through CE.

Refer to [Figure 28, “Configure CE Group option,”](#) on page 96.

CAUTION

The Configure CE Group Option fails with the following error message if the Microsoft Failover Cluster Group does not have a disk in the group:

The Group <group_name> does not have the Clustered Disk

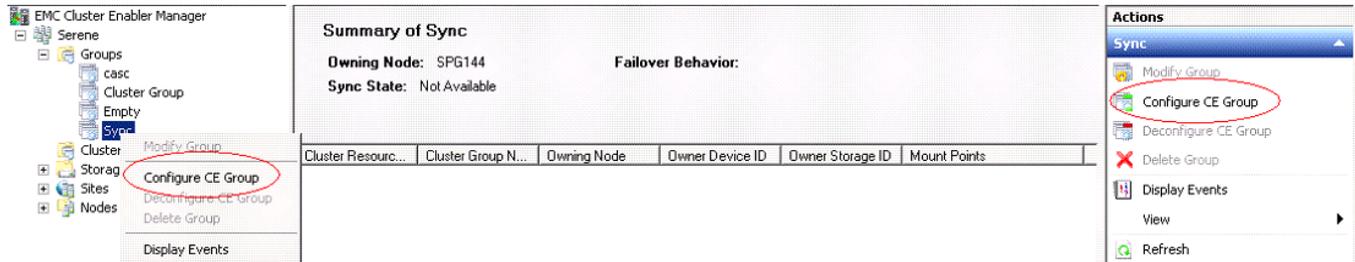


Figure 28 Configure CE Group option

Deconfigure a CE group

To deconfigure a CE group, select a **group** listed under the Groups icon located in the Navigation Console Tree and select **Action** and **Deconfigure CE Group** from the menu bar. The deconfigure option can also be accessed by using the Right-click or Action menus. A dialog box pop-up will appear asking you to confirm the action. Click **Yes** to convert the group or **No** to abort the action.

Note: A VM group can be de-configured only if its underlying CSV disk group is de-configured from Cluster Enabler. De-configuring the CSV disk group de-configures all dependent VMs.

CAUTION

If the group is deconfigured, CE failover support to the remote nodes will no longer be operational. To make group failover operational again, you will need to reconfigure the cluster group using the CE Configuration Wizard in the CE Manager.

Delete a CE group

To delete a CE group, select a **group** listed under the Groups icon located in the Navigation Console Tree and select **Action** and **Delete Group** from the menu bar. The delete group option can also be accessed using the Right-click or Action menus. A dialog box pop-up will appear asking you to confirm the action. Click **Yes** to delete the group or **No** to abort the action.

Note: Deleting a CE group deconfigures the group and then removes it from the cluster.

Storage component

Selecting the Storage icon from the navigation tree allows you view the attached storage device information for each storage array. Select a storage array to view the summary information columns in the center pane. [Figure 29](#) shows the CE Manager Storage component view for a Symmetrix array.

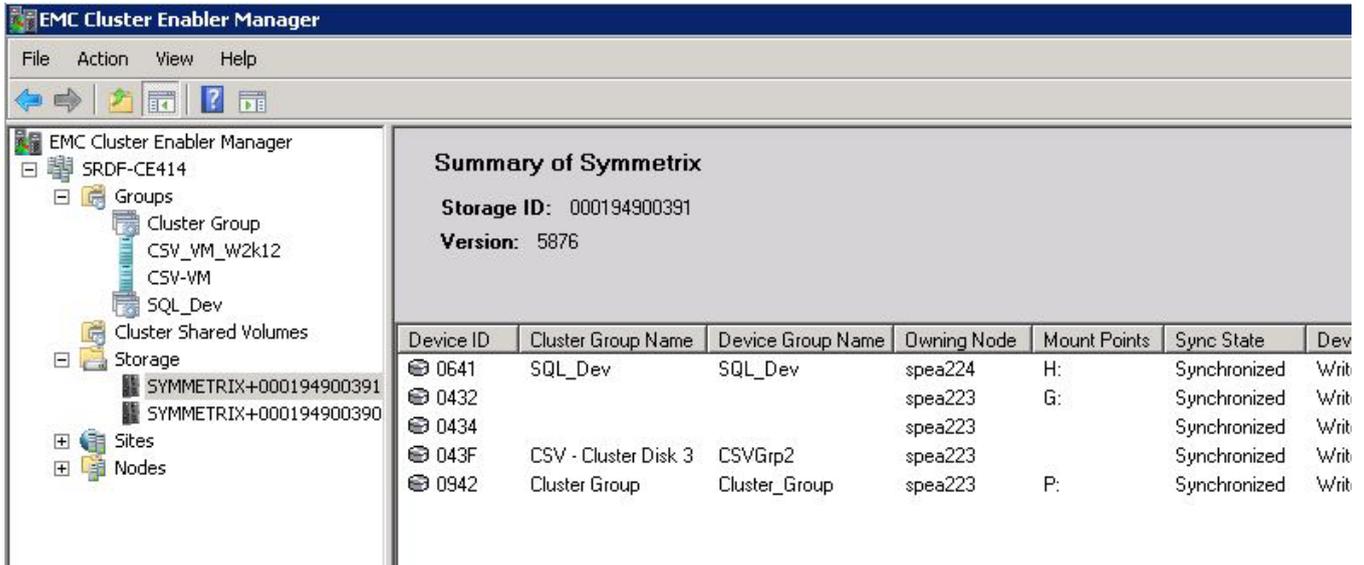


Figure 29 Example of Symmetrix storage array view

[Table 7 on page 97](#) lists the summary and column heading information displays for Symmetrix storage arrays.

Table 7 Storage component display information (page 1 of 2)

Summary and column heading information	Description
Name	The Symmetrix array ID (such as, Symmetrix+00187900830).
Version	Displays the Enginuity version.
Device ID	Shows all SRDF R1/R2 device IDs that are mapped to any cluster member node.
Cluster Group Name	Indicates the CE Group name to which the device belongs.
Device Group Name	Indicates the SYMAPI device group or composite group name to which the device belongs; derived from Cluster Group name.
Owning Node	If a device belongs to a cluster group, the owning node information is obtained directly from Microsoft Failover Cluster. Otherwise, the owning node is a node where the device is write-enabled.
Mount Points	Indicates the mount point of the physical drive on the owning node.
Sync State	Indicates the RDF state for the group. The <i>EMC Solutions Enabler SRDF Family CLI Product Guide</i> provides a listing of all possible RDF states.
Device Status	Indicates the SRDF R1/R2 device status. The possible device status states are Ready, Not Ready, and Write-Disabled.
Capacity MB	Indicates the device capacity in megabytes.

Table 7 Storage component display information (page 2 of 2)

Summary and column heading information	Description
Swap Capable	Indicates whether the device is swap capable (True or False).
Async Capable	Indicates whether the device is asynchronous capable (True or False).
WWN	Displays the devices World Wide Name (WWN).
Logical Device	Displays the logical device name (if applicable).
RDF Type	For Symmetrix arrays, shows the RDF device type of R1 or R2.
RA Group	For Symmetrix arrays, indicates the RA group to which the group belongs.
R1 Invalid Tracks	For Symmetrix arrays, indicates the number of invalid R1 track (if any).
R2 Invalid Tracks	For Symmetrix arrays, indicates the number of invalid R2 track (if any).
RDF Async Lag Time	For Symmetrix arrays, indicates the lag time between the target (R2) device and the source (R1) device in an SRDF/Asynchronous environment.
Invista WWN	Displays the Invista devices World Wide Name (WWN).

Adding and removing devices from a group

From the storage view, when you select a device listed in the center pane, you can add or remove it from the group by clicking the **Add to Group** or **Remove from Group** icons displayed in the right Action pane. [Figure 30](#) displays the storage actions located in the right action pane.

**Figure 30** CE Manager storage actions

When you select the action, Cluster Enabler opens the **“Modify Group Wizard”** at the validation step. Click **Next** to add or remove your selection. The **Summary** dialog box appears. Upon Group Modified Successfully, click **Finish** to exit the wizard.

Viewing information

Cluster Enabler allows you to view certain summary information about CE Groups, Storage, Sites and Nodes. This information is displayed in the center pane as you select each of the icons located in the navigation tree. [“Storage component” on page 97](#) describes the storage information and available actions. Summary information for Groups, Nodes, and Sites are described below.

Displaying group information

Selecting the Groups icon from the navigation tree allows you view group information for all configured CE cluster groups in the center pane. Double-clicking on a specific group displays summary information for that group. [Figure 31](#) displays the CE Manager group component.

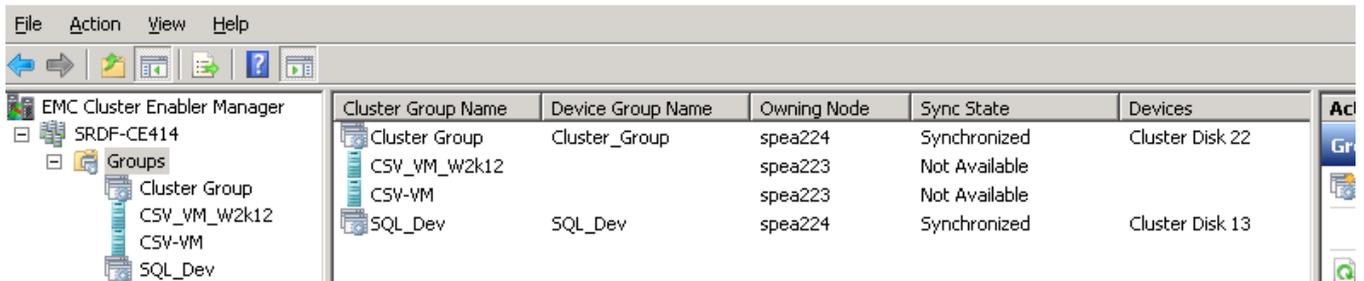


Figure 31 CE Manager Groups component

Selecting a specific group icon from the navigation tree allows you view the group summary information columns for each configured CE cluster group in the center pane. [Figure 32](#) displays the CE Manager group information for a specific group.

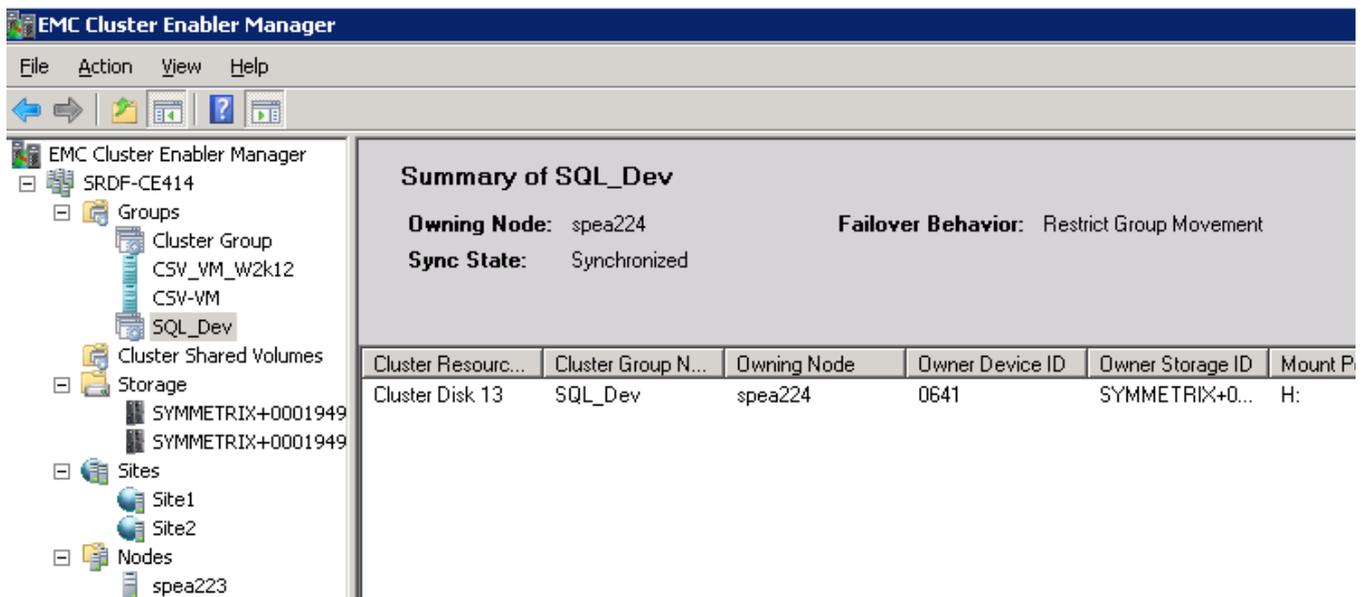


Figure 32 CE Manager groups information

[Table 8](#) lists the summary and column heading information that displays for CE Groups.

Table 8 Groups component displayed information

Summary and column heading information	Description
Cluster Group Name	Indicates the CE Group name to which the device belongs.
Device Group Name/ Consistency Group Name	Indicates the SYMAPI device group or composite group name to which the Symmetrix device belongs; derived from Cluster Group name.
Owning Node	Shows the failover cluster node name that owns the particular group. This information is obtained directly from MS Failover Cluster. Only groups that are part of the cluster display.
Sync State	Indicates the RDF state for the group. The <i>EMC Solutions Enabler SRDF Family CLI Product Guide</i> provides a listing of all possible RDF states,
Devices	Listed by disk resource name in the cluster.
Cluster Resource Name	Listed by physical disk resource name.
Owner Device ID	The Symmetrix device ID mapped to the owning node (such as, ODEC, ODED).
Owner Storage ID	The Symmetrix array ID (such as, Symmetrix+00187900830).
Mount Points	Indicates the mount point of the physical drive on the owning node.

Clicking the **Display Events** icon option in the action pane displays event information in the lower tier of the center pane. [Table 9](#) lists the heading information that displays for CE Groups events.

Table 9 Groups event information

Column heading	Description
Date/Time	Shows the date and time that the recorded event occurred.
Computer Name	Indicates the computer name on which the event occurred.
Group Name	Indicates the group name to which the event occurred.
Message	Displays a detailed message of the event type.

If the group is a virtual machine (VM), the icon for the group changes. CSV-related information about the VM is displayed as shown in [Figure 33 on page 101](#). Refer [Table 10 on page 101](#) for CSV Group Information.

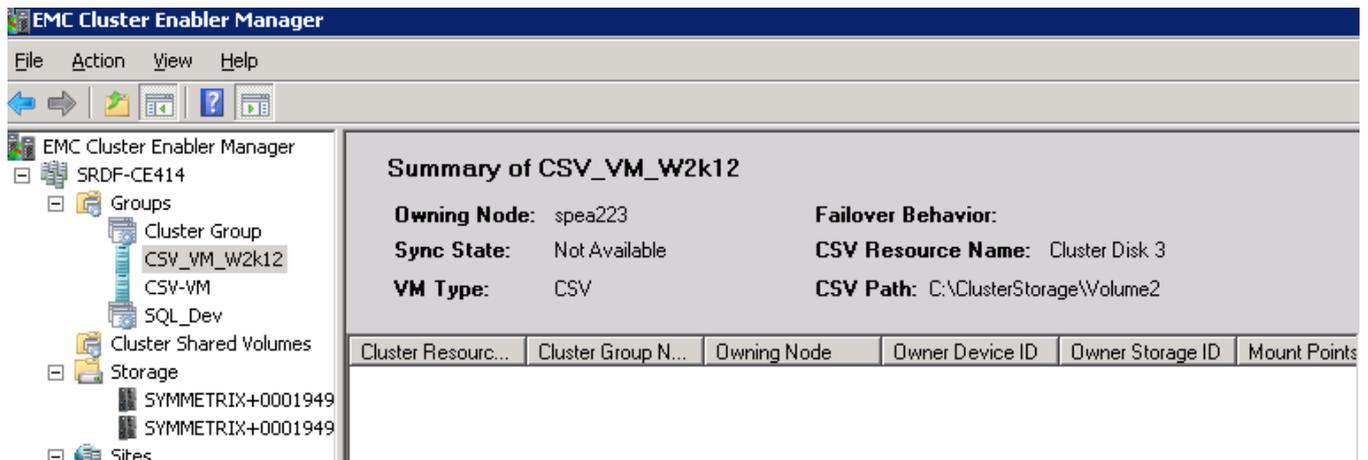


Figure 33 CE Manager VM groups information

Table 10 CSV Group information

Summary and column heading information	Description
Cluster Resource Name	Listed by physical disk resource name.
CSV Path	The Windows path where the CSV disk is mounted. Usually shown as c:\ClusterStorage\Volume.
Owning Node	Shows the failover cluster node name that owns the particular group. This information is obtained directly from MS Failover Cluster. Only groups that are part of the cluster display.
Device Group	Indicates the SYMAPI device group name to which the device belongs; derived from Cluster Group name.
Sync State	Indicates the RDF state for the group. For a listing of all possible RDF states, refer to the EMC Solutions Enabler SRDF Family CLI Product Guide.

Displaying node information

Selecting the **Nodes** icon from the navigation tree allows you view node information for all nodes in the center pane. Double-clicking on a specific node displays summary information for that node. [Figure 34](#) displays the CE Manager nodes component.

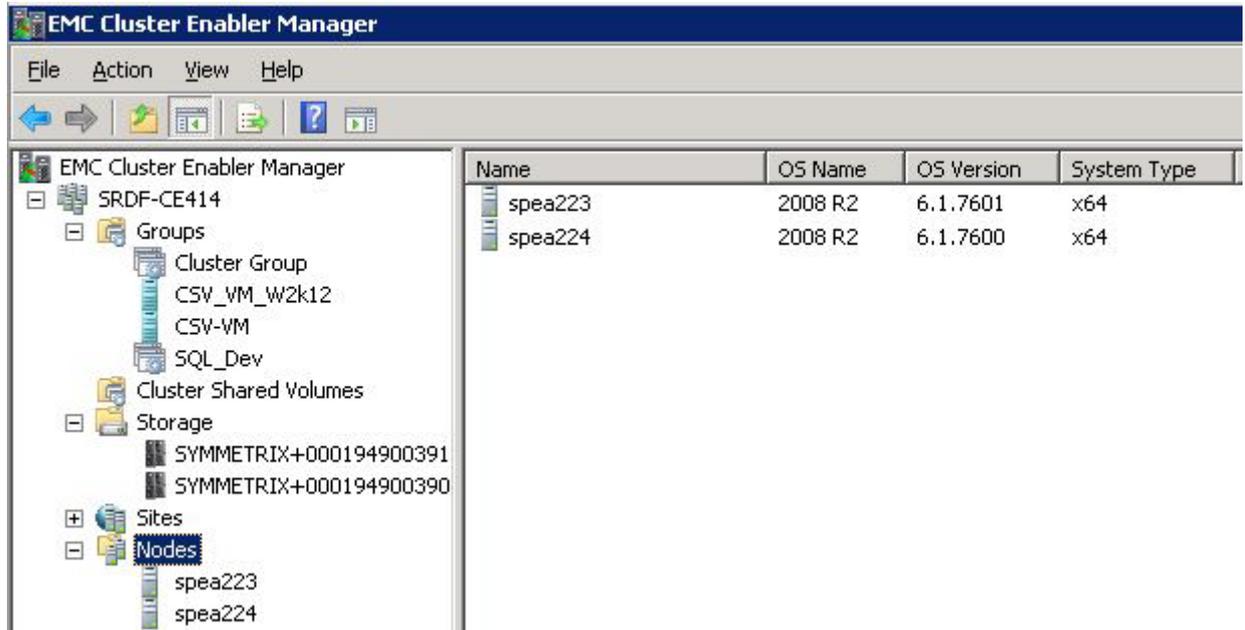


Figure 34 CE Manager Nodes component

Selecting a specific node icon from the navigation tree allows you view the node summary information columns for each node in the center pane. [Figure 35](#) displays the CE Manager node information for a specific node.

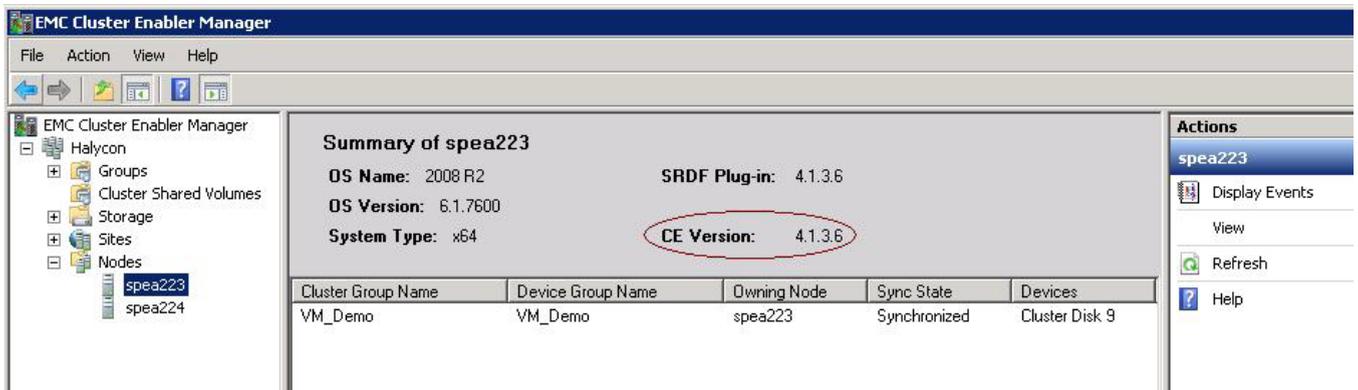


Figure 35 CE Manager node information

[Table 11](#) lists the summary and column heading information that displays for Nodes.

Table 11 Nodes component displayed information

Summary and column heading information	Description
Name	Displays the node name.
OS Name	Displays the Windows operating system (such as, 2008 SP2).
OS Version	Displays the Windows operating system version (such as, 5.2.3790).
System Type	Displays the Windows system type (such as, X86).
CE Plug-in	Displays the CE Plug-in version (i.e. 4.0.0.22).
CE Version	Displays CE Base version installed on the selected node.
Cluster Group Name	Indicates the CE Group name to which the device belongs.
Device Group Name	Indicates the SYMAPI device group or composite group name to which the device belongs; derived from Cluster Group name.
Owning Node	Shows the failover cluster node name that owns the particular group. This information is obtained directly from MS Failover Cluster. Only groups that are part of the cluster will display.
Sync State	Indicates the RDF state for the group. The <i>EMC Solutions Enabler SRDF Family CLI Product Guide</i> provides a listing of all possible RDF states,
Devices	Listed by cluster resource name.

Clicking the **Display Events** icon option in the action pane displays event information in the lower tier of the center pane. [Table 12](#) lists the heading information that displays for CE Node events.

Table 12 Node event information

Column heading	Description
Date/Time	Shows the date and time that the recorded event occurred.
Computer Name	Indicates the computer name on which the event occurred.
Group Name	Indicates the group name to in which the event occurred.
Message	Displays a detailed message of the event type.

Displaying site information

Selecting the Sites icon from the navigation tree allows you view site information for all sites in the center pane. Double clicking on a specific site displays summary information for that site. [Figure 36 on page 104](#) displays the CE Manager Site component.

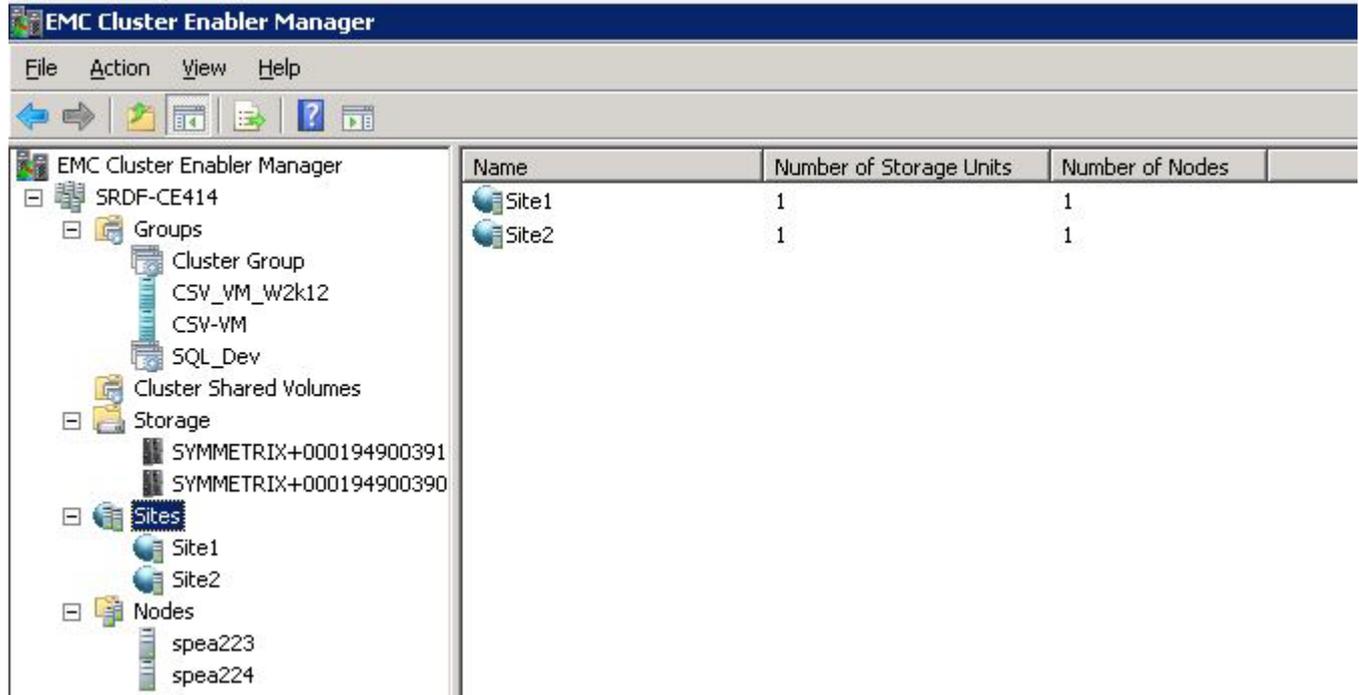


Figure 36 CE Manager Sites component

Selecting a specific site icon from the navigation tree allows you view the site summary information columns for each site in the center pane. [Figure 37 on page 105](#) displays the CE Manager site information for a Symmetrix site.

You can change the name of a site by using the right-click menu and selecting **Rename**. The site name will then become editable. The site name can also be changed by using the right-click menu and selecting **Properties**. You may then change the site name in the Properties tab.

The screenshot displays the EMC Cluster Enabler Manager interface. The left-hand tree view shows the hierarchy: SRDF-CE414 > Groups > Cluster Group > CSV_VM_w2k12, CSV-VM, SQL_Dev; Cluster Shared Volumes; Storage > SYMMETRIX+000194900391, SYMMETRIX+000194900390; Sites > Site1, Site2; and Nodes > spea223, spea224. The main pane shows the 'Summary of Site2' with the following information:

- Number of Storage Units: 1
- Number of Nodes: 1

Below the summary are three tables:

Storage Unit Name and Id	Version Information
SYMMETRIX+000194900390	5876

Node Name	OS Name	OS Version	System Type
spea223	2008 R2	6.1.7601	x64

Cluster Group Name	Device Group Name	Owning Node	Sync State	Devices
Cluster Group	Cluster_Group	spea223	Synchronized	Cluster Disk 22
CSV-VM		spea223	Not Available	
CSV_VM_w2k12		spea223	Not Available	
SQL_Dev	SQL_Dev	spea224	Synchronized	Cluster Disk 13
CSV - Cluster Disk 3	CSVGrp2	spea223	Synchronized	Cluster Disk 3

Figure 37 CE Manager Symmetrix site information

[Table 13](#) lists the summary and column heading information that displays for Symmetrix sites.

Table 13 Site component displayed information (page 1 of 2)

Summary and column heading information	Description
Name	Displays the Site name.
Number of Storage Units	Displays the number of storage units for this site.
Number of Nodes	Displays the number of nodes for this site.
Storage Unit Name & ID	The Symmetrix array ID (such as, Symmetrix+00187900830).
Version Information	Displays the Symmetrix Engenuity version.
Node Name	Displays the node name.
OS Name	Displays the Windows operating system (such as, 2008 SP2).
OS Version	Displays the Windows operating system version (such as, 5.2.3790).
System Type	Displays the Windows system type (such as, X86).
Cluster Group Name	Indicates the CE Group name to which the device belongs.
Device Group Name	Indicates the SYMAPI device group or composite group name to which the device belongs; derived from Cluster Group name.

Table 13 Site component displayed information (page 2 of 2)

Summary and column heading information	Description
Owning Node	Shows the failover cluster node name that owns the particular group. This information is obtained directly from MS Failover Cluster. Only groups that are part of MS Failover Cluster display.
Sync State	Indicates the RDF state for the group. The <i>EMC Solutions Enabler SRDF Family CLI Product Guide</i> provides a listing of all possible RDF states.
Devices	Listed by cluster resource name.

Restore and recovery operations

This section details some of the restore and recovery operations that should be performed for different types of failures.

The following SRDF/CE restore and recovery operations are provided:

- ◆ [“Restoring a failed SRDF site” on page 106](#)
- ◆ [“Recover SRDF backup site in case of primary site failures” on page 107](#)
- ◆ [“Recovery from SRDF link failure” on page 108](#)
- ◆ [“Restrict group movement and recovery” on page 108](#)
- ◆ [“Recovery from a corrupt quorum log” on page 109](#)
- ◆ [“Symmetrix array replacement” on page 109](#)
- ◆ [“Recover CE Cluster Wizard” on page 109](#)

SRDF/CE recovery procedures

This section details SRDF/CE restore and recovery operations that should be performed for different types of failures.

Restoring a failed SRDF site

The following procedure describes how to restore your storage system after a site failure occurs with all links lost:

1. Restore SRDF and IP links.
2. Restart all nodes.
3. Open CE Manager and connect to the cluster.
4. Perform Storage Discover from CE Manager.

Any groups that are failed over to a secondary site are in a split state. Groups that are not failed over are in suspended state. You can safely bring the groups that did not failover to a secondary site online at this point.

5. To restore groups that failed over to a secondary site, follow these steps:

⚠ WARNING

Choosing the wrong option for restore could cause data loss. Contact EMC support if you have any question about the commands that should be issued.

Assuming that the secondary site has good data and you want to copy this data to primary site, follow these steps:

- a. Open the command line prompt.
- b. Change the directory to **<CE Install Directory>**.
- c. Set the `SYMCLI_DB_FILE` environment variable to `SRDFCESymapi.db`.
- d. Run the following command for every failed-over group.

```
symrdf -g <failed over group name> restore -incr
```

- e. Monitor the group state by running the following command:

```
symrdf -g <groupname> query
```

Assuming that the primary site has good data and you want to copy this data to secondary site:

- a. Open the command line prompt.
- b. Change the directory to **<CE Install Directory>**.
- c. Set the `SYMCLI_DB_FILE` environment variable to `SRDFCESymapi.db`.
- d. Run the following command for every failed over group:

```
symrdf -g <failed over group name> establish -incr
```

- e. Monitor the group state by running the following command:

```
symrdf -g <groupname> query
```

Recover SRDF backup site in case of primary site failures

The following procedure describes how to recover a backup site when the primary site fails. Cluster Enabler lets you set the failover option on a group basis.

For MNS clusters

You can restart the backup site using the `/forcequorum` option that is described in Microsoft clusters manual.

For Shared Quorum models

1. If the Cluster Group (the group in which the quorum disk is a member) does not have the failover option set to "Automatic Failover", then the group will not failover to secondary node and therefore the cluster cannot be started. On one of the secondary nodes, use the Recover CE Wizard to start the cluster in "Safe Mode". This starts the cluster service on this node with just the Cluster Group.

Once you have cluster service running on the secondary site:

2. Open CE Manager and connect to the cluster.

3. For each group that you want to failover, change the failover policy to “Automatic Failover”.
4. From MS Cluster Administrator/Failover Cluster Manager, bring all of these groups online.

At this point the cluster is running with required services at the back up site.

Recovery from SRDF link failure

The following two procedures describe how to recover from an SRDF link failure:



Choosing the wrong option for restore could cause data loss. Contact EMC support if you have any question about the commands that should be issued.

- ◆ For groups that failed-over on the RDF link, when the link was in a failed state:
 - a. Choose the remote mirror that has valid user/application data.
 - b. Move the group to a node that has a valid mirror mapped.
 - c. Restore the SRDF link.
 - d. Open CE Manager and perform Storage Discover.
 - e. Open the command line prompt.
 - f. Set the `SYMCLI_DB_FILE=<CE InstallDir>\SRDFCESymapi.db`
 - g. If the R1 has valid data, enter the following command:


```
symrdf -g <groupname> establish -incr
```

 If the R2 has valid data, enter the following command:


```
symrdf -g <groupname> restore -incr
```
- ◆ For groups that remained online on the same side as before the link failure:
 - a. Restore SRDF link.
 - b. Open CE Manager and run Storage Discover.

Restrict group movement and recovery

If a CE group has the "Restrict Group Movement" policy set and the RFD link is down, it may take a long time for the resource to come online if the user manually tries to move the group to a node that is connected to a different storage array. For example, if the user tries to move group G1 from the R1 side to the R2 side when the RDF link is down, then Microsoft's preferred owner logic will attempt to bring the group online on the R2 side as expected.

But since the restrict group movement policy is set for the CE group, Microsoft will fail the resource on the R2 side nodes. This is correct behavior and is expected, but it may take a long time for the resource to fail on all the R2 nodes before coming back online on one of the R1 side nodes. This is because by default Microsoft will try to bring the group online 3 times on each node. The more nodes you have in the cluster the longer it will take for

Microsoft to complete the preferred owner logic. To minimize this undesirable effect you can change the property of the resources to "Do not Restart". This will minimize number of retries and reduce the time required to bring the group online.

Recovery from a corrupt quorum log

The complete Microsoft article can be found on Microsoft Knowledge Base Article 172951 at the following site:

<http://support.microsoft.com/kb/q172952>

Symmetrix array replacement

The following process can be used to replace a Symmetrix array. This process assumes that all RDF groups are Dynamic and that all failover cluster groups are configured for Swapping RDF personalities (SwapEnabled) during failover:

1. Change the Microsoft Failover Cluster service start up to Manual on all cluster nodes.
2. Failover all groups to the Symmetrix array that you are NOT replacing. Now the groups are online on R1 side of RDF device.
3. Shutdown all nodes that are attached to the Symmetrix array that is being replaced.
4. Replace the R2 Symmetrix array and establish new R1/R2 relations.
5. Bring the SRDF link up and synchronize R2 with R1 data.
6. Wait for the synchronization to complete.
7. Adjust device masks on all nodes connected to new Symmetrix array, so that the devices are correctly mapped to these hosts.
8. Reboot the nodes attached to new Symmetrix array.
9. Open CE Manager on one of the nodes connected to R1 side.
10. Choose UpdateMirrorPair and step through the wizard processes.
11. Once UpdateMirrorPair wizard completes successfully, CE updates its internal configuration database to reflect new R1/R2 relations. At this point you should be able to failover groups between these Symmetrix array.
12. Reset the Cluster Service Startup type to Automatic.

Recover CE Cluster Wizard

Cluster Enabler provides a wizard to help you recover a failed shared quorum cluster by bringing the cluster online on a single node. The shared quorum cluster will fail to come online in a site failover scenario where the failover option for a quorum group is set to "Restrict Group Movement". The Recover CE Cluster Wizard changes the failover policy on quorum group to "Automatic Failover" and then brings the cluster online on the node. You can then use the Create Group Wizard to change other groups failover policies and bring them online appropriately.

Note: The Recover CE Cluster Wizard is useful for shared quorum clusters. To force a Majority Node Set (MNS) cluster node to form a cluster use the /forcequorum option as documented in your Microsoft Clusters documentation.

Follow these steps to automatically recover and restore a shared quorum cluster using the Recover CE Cluster Wizard:

1. Select the **EMC Cluster Enabler Manager** icon from the navigation tree and select **Action** and **Recover CE Cluster** from the menu bar. This opens the first page of the Recover CE Cluster Wizard. The wizard can also be launched by using the right-click or Action menus.

Note: When running the Recover CE Wizard to recover a CE cluster, you should only run the wizard on a failed node when the entire cluster is down.

2. The Enter Node Name page appears. Type the **Cluster Name** and **Node Name**, click **Validate**. The Recover CE Wizard should only be run on a single node. [Figure 38 on page 110](#) shows the Enter Node Name page.

Figure 38 Recover CE Cluster Enter Node Name

3. The Choose Tasks page appears. To restart a cluster in safe mode and bring the cluster online using previous CE cluster settings, select **Start Cluster in Safe Mode**. To resolve a cluster number for a Shared Quorum model cluster and recover the cluster, select **Resolve Cluster Number**. Click **Next**. [Figure 39](#) shows the Choose Tasks page.

Figure 39 Recover CE Cluster Choose Tasks

- If you selected **Resolve Cluster Number** in step 3 and are recovering a shared quorum model cluster, the following screen appears.

Figure 40 Recover CE Cluster Change Cluster Number

The wizard will automatically generate a list all of the available cluster numbers. From the Select a Cluster Number scroll box, select the Cluster Number that you want to use for the cluster, click **Next**. [Figure 40](#) shows the Change Cluster Number page.

Note: Check the Show All Cluster Numbers box to view all of the cluster numbers both used and unused for the system. Do not select a number that is already used.

- The Summary page appears and Cluster Enabler begins to restart the cluster service for the CE cluster. Upon "Started cluster service successfully", click **Finish**.

Configuring a custom resource

This section provides examples of creating and modifying a custom resource using the CE Manager. A custom resource could be a Veritas volume or other third-party resource. Once a CE Group is created and the custom resource volumes are added, the storage resource will be comprised of Symmetrix disks.

Before you can configure a custom resource using the CE Manager, you must set up the custom resource using the vendor's resource software (for example, Veritas Volume Manager). Then you must manually add the custom resource to Microsoft Failover

Clusters. For example, a custom resource is of the Resource Type “Generic Application”. Figure 41 shows a custom resource named “test” in “Group 4” as displayed from the Microsoft Cluster Administrator application.

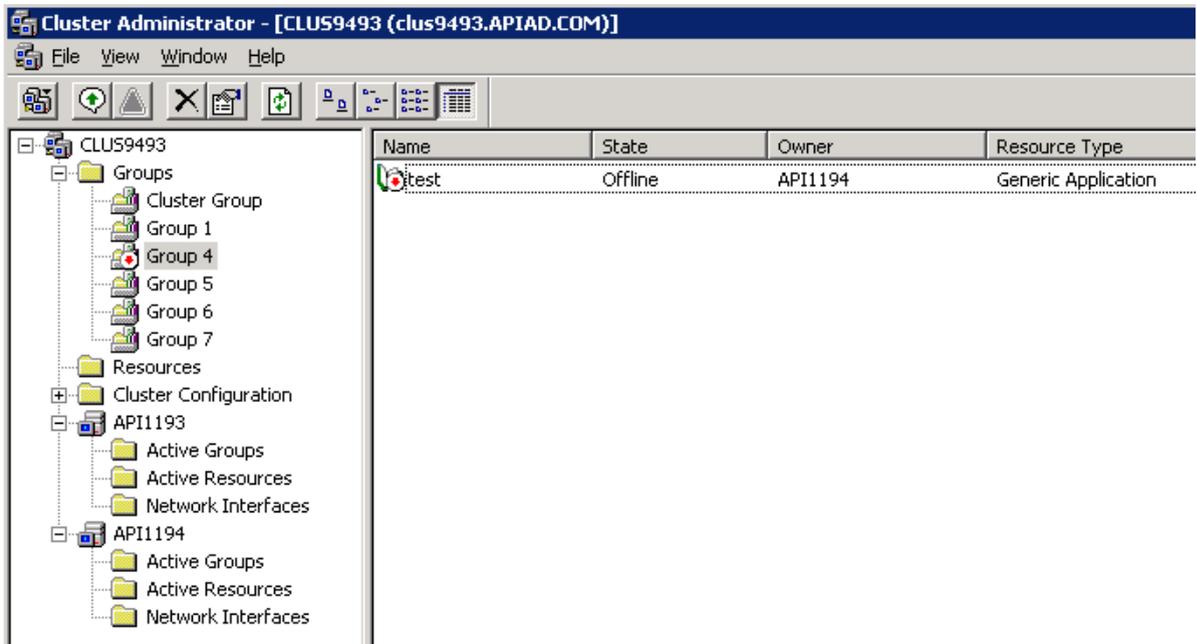


Figure 41 Microsoft Cluster Administrator, Generic Application Resource Type

For Cluster Enabler to recognize a third-party resource, it must be added to cluster properties. Figure 42 displays the “CustomResourceTypes” as listed in the EMC Cluster Enabler cluster properties, which is viewable from the command line.

```
C:\Program Files\EMC\Cluster-Enabler>cluster /priv
Listing private properties for '':
T Cluster          Name              Value
-----
M CustomResourceTypes <?xml version="1.0"?>
M CustomResourceTypes <CustomResTypeList>
M CustomResourceTypes <CustomResType>Volume Manager Disk Group</C
CustomResType>
M CustomResourceTypes </CustomResTypeList>
M SiteList          SiteList          <?xml version="1.0"?>
<SiteList>
<Site>
<SiteName>Site1</SiteName>
<Node>API1194</Node>
<Storage>SYMMETRIX+000187900830</Storage>
</Site>
<Site>
<SiteName>Site2</SiteName>
<Node>API1193</Node>
<Storage>SYMMETRIX+000187900848</Storage>
</Site>
</SiteList>
D ClusterNumber    38 (0x26)
D ISECLuster       1 (0x1)
```

Figure 42 Cluster properties

If you would like to use another third-party resource (for example, Generic Application), you need to run the following command string from the command line:

```
cluster /priv CustomResourceTypes="<?xmlversion=1.0?>,"<CustomResTypeList>,
" <CustomResType>Volume Manager Disk Group</CustomResType>,
" <CustomResType>Generic Application</CustomResType>,
"</CustomResTypeList>:MULTISTR
```

Figure 43 displays the changed cluster properties with Generic Application added to CustomResourceTypes.

```
C:\Program Files\EMC\Cluster-Enabler>cluster /priv CustomResourceTypes="<?xml version="1.0"?>","<CustomResTypeList>","<CustomResType>Volume Manager Disk Group</CustomResType>","<CustomResType>Generic Application</CustomResType>","</CustomResTypeList>":MULTISTR
```

```
C:\Program Files\EMC\Cluster-Enabler>cluster /priv
Listing private properties for '':
```

T	Cluster	Name	Value
M		CustomResourceTypes	<?xml version="1.0"?>
M		CustomResourceTypes	<CustomResTypeList>
M		CustomResourceTypes	<CustomResType>Volume Manager Disk Group</CustomResType>
M		CustomResourceTypes	<CustomResType>Generic Application</CustomResType>
M		CustomResourceTypes	</CustomResTypeList>
M		SiteList	<?xml version="1.0"?>
		<SiteList>	
		<Site>	
		<SiteName>Site1</SiteName>	
		<Node>API1194</Node>	
		<Storage>SYMMETRIX+000187900830</Storage>	
		</Site>	
		<Site>	
		<SiteName>Site2</SiteName>	
		<Node>API1193</Node>	
		<Storage>SYMMETRIX+000187900848</Storage>	
		</Site>	
		</SiteList>	
D		ClusterNumber	38 <0x26>
D		ISCECluster	1 <0x1>

Figure 43 Cluster properties with Generic Application

After you have configured your custom resource for MS failover clusters, you can use the CE manager Create Group Wizard to create a custom resource CE Group. “Using CE Manager to create a custom resource CE Group” on page 113 explains the process.

Using CE Manager to create a custom resource CE Group

Follow these steps to create a CE Group using custom resources for management with Cluster Enabler.

1. Open the CE Manager and select the **Groups** icon from the Navigation tree and select **Action** and **Create Group** from the menu bar. This begins the process of reading the storage configuration. After the storage has been read, the first page of the Create Group Wizard opens. The Create Group Wizard can also be launched using the Right-click or Action menus.

Note: A mirrored pair needs to be present on the array before attempting to create a group. Run the Storage Discover Wizard to detect a newly created mirrored pair by right-clicking on the **cluster name** or clicking the **Discover** button in the Select Devices page of the Create Group Wizard.

2. The **Enter Group Name** dialog box appears. Enter the exact same **Group Name** as displayed in the MS Cluster Administrator in the space provided and click **Create**. For this example, the Group Name is “Group 4”. Click **Cancel** to abort the operation and close the wizard.
3. The next wizard page prompts you to select devices for inclusion in the new group. The wizard recognizes that this is a custom resource group and displays a warning that a custom resource is being configured.

Note: A SYMAPI device group for Symmetrix will be created by Cluster Enabler. The corresponding CE resource will also be created and the custom resource will be made dependent on the CE resource. Physical disk resources will not be created in the failover cluster by Cluster Enabler.

Select the appropriate devices from the list shown by clicking in the **select boxes**. Selected devices are identified by the checked box. Click **Next**.

Note: The tree view can be expanded by selecting **Expand All**. There are three types of devices that can be displayed by checking the selection boxes: Async, Cascaded, and Concurrent. For example, selecting the Async checkbox displays all SRDF asynchronous capable devices within in the same RA group. An error message displays if selected type of devices are used up or not available. If you select devices from a single Symmetrix RA group, a device group will be created. If you select devices from multiple Symmetrix RA groups, a composite group will be created.

4. The Select Group Policy page appears. From the pull-down menu, select your desired policy for the group. You can select either the **Restrict Group Movement** or **Automatic Failover**. Once selected, click **Next**. [Figure 44](#) shows the select group policy for the devices in mount point N.

The **Restrict Group Movement** selection restricts the group from failing over to a peer node. In an SRDF link failure, this setting will only attempt to move disk laterally. If the link is up, this setting has no impact.

The **Automatic Failover** policy sets the group to automatically failover to another node in the event of a node or network failure.

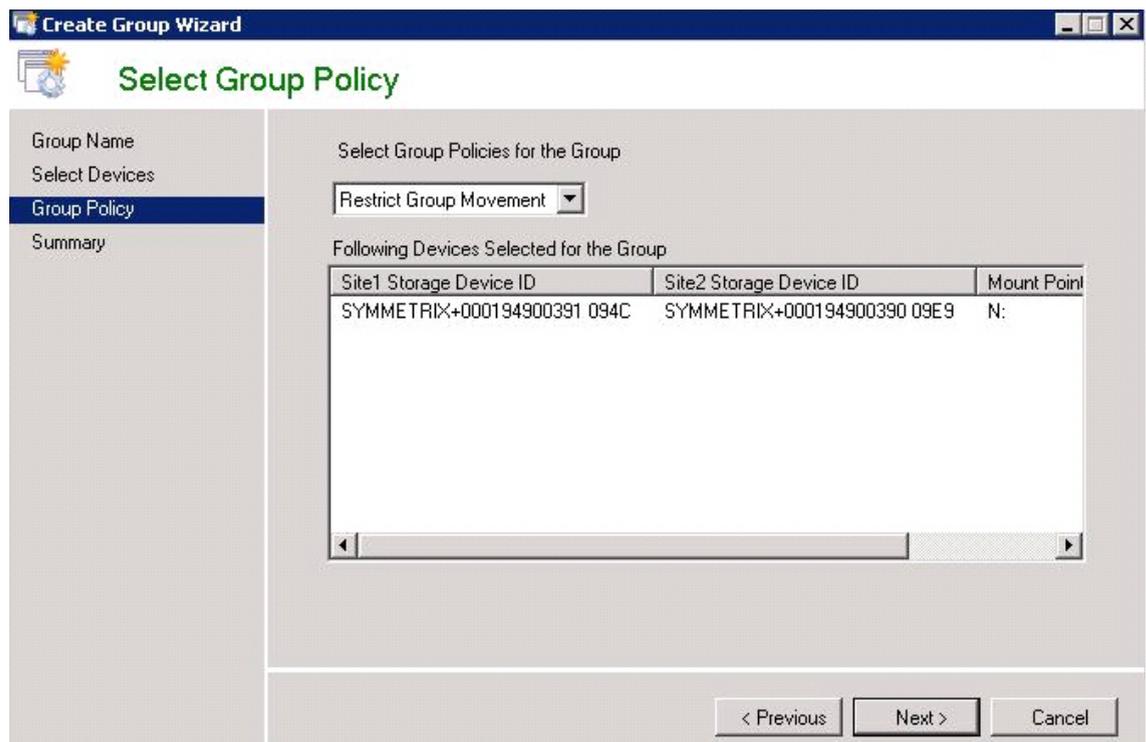


Figure 44 Select Group Policy, custom resource

5. The Summary page appears. Upon Group Created Successfully, click **Finish** to exit the wizard.
6. Cluster Enabler automatically begins refreshing the CE cluster. Upon completion of the refresh, you should see the group that you created listed under Groups. If you do not see the newly created group, select **Action** and **Refresh** from the menu bar. The Refresh action can also be accessed from the right-click or Action menus.
7. Open the Microsoft Cluster Administrator application and select Group 4. A resource named “EMC_Group 4” of resource type “EMC Cluster Enabler” is now visible in Group 4. [Figure 45](#) displays the new group in the Microsoft Cluster Administrator application.

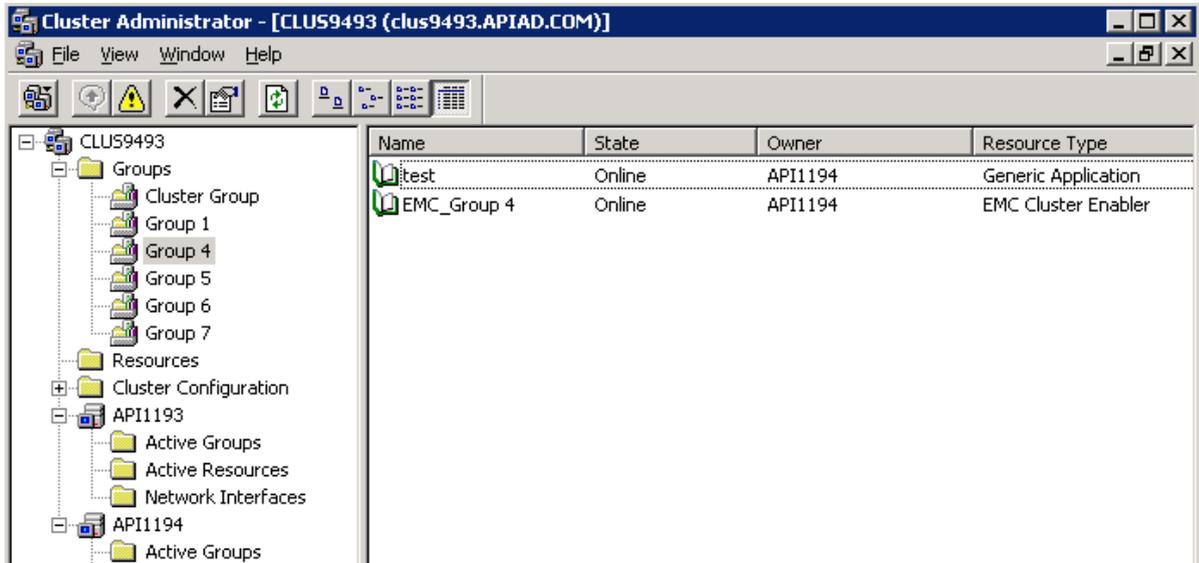


Figure 45 Microsoft Cluster Administrator, EMC_Group 4

Using CE Manager to edit a custom resource CE Group

If the composition of an underlying custom resource changes, you should make the same changes to the CE Group custom resource by adding or deleting devices from the group. Changes to a custom resource group can be made by using the CE Manager Modify Group Wizard. The following example adds devices to the custom resource CE Group in a Symmetrix array.

Follow these steps to add or remove devices from a custom resource CE Group using the Modify Group Wizard.

1. Select the **Group** icon in the navigation tree and select **Action** and **Modify Group** from the menu bar. This begins the process of reading the storage configuration. After the storage configuration has been read, the first page of the Modify Group Wizard opens. The Modify Group Wizard can also be launched using the Right-click or Action menus. In this example, Group 4 is selected.

Note: A mirrored pair needs to be present on the array before attempting to modify a group. Run the Storage Discover Wizard to detect a newly created mirrored pair by right-clicking on the **cluster name** or clicking the **Discover** button in the Select Devices page of the Modify group wizard.

- From the Select Devices page, select the Action from the pull-down menu for either **Add Devices** or **Delete Devices**. Depending on your selection, a list of available devices that can be added or removed displays. Symmetrix RA group pairs and the devices contained within the RA group are shown in a tree view. By default, the RA Groups are shown in collapsed view (Collapse All). The tree view can be expanded by selecting **Expand All**.

Select the desired devices from the list shown by clicking in the select boxes. Selected devices are identified by the checked box. Selecting the RA group, automatically selects all devices in that group. After your selections have been made click **Next**.

- The Validate Selection page appears, click **Next** to validate your selection or click **Cancel** to abort the action. The wizard recognizes that this is a custom resource group and displays a warning that a custom resource is being modified. [Figure 46 on page 116](#) displays the Validate Selection page for Group 4.

Note: Only the storage group and the corresponding CE resource will be modified. No physical disk resources will be added to the failover cluster.

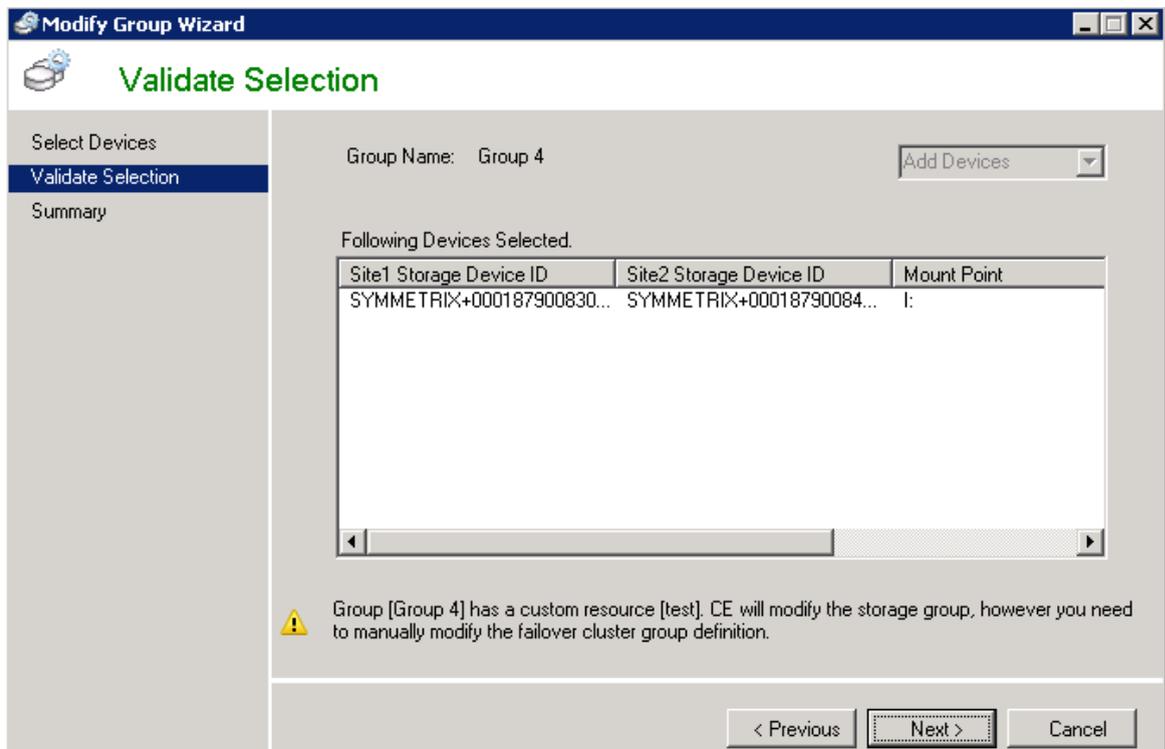


Figure 46 Validate selection, custom resource

- The Summary page appears. Upon Group Modified Successfully, click **Finish** to exit the wizard.
- Cluster Enabler automatically begins refreshing the CE cluster. Upon completion of the refresh, you should see the updated group information reflecting the devices added or deleted. If you do not see the updated group information, select **Action** and **Refresh**

from the menu bar. The Refresh action can also be accessed from the right-click or Action menus. [Figure 47 on page 117](#) displays a summary of the modified Group 4 from this example.

Summary of Group 4

Owning Node: API1194 **Failover Behavior:** Restrict Group Movement

Sync State: Synchronized

Cluster Resource	Cluster Group Name	Owning Node	Owner Device ID	Owner Storage ID
test	Group 4		ODD6	SYMMETRIX+0...
test	Group 4		ODD7	SYMMETRIX+0...

Figure 47 Summary of Group 4, custom resource

Note: The **Deconfigure CE Group** option removes the storage group definition and CE resource but does not change the Microsoft failover cluster physical disk resources.

Appendix A

Base Component Installation and Upgrade

This chapter explains how to install and uninstall the Cluster Enabler Base component and provides configuration instructions for Windows Server 2008 (including R2) and 2012.

- ◆ [Installation overview](#) 120
- ◆ [Before you begin](#) 120
- ◆ [Installing the Cluster Enabler Base Component](#)..... 121
- ◆ [Uninstalling the Cluster Enabler Base Component](#) 123
- ◆ [Configuring a CE cluster on Server Core](#) 124
- ◆ [Upgrading Windows Server 2008 to Windows Server 2008 R2](#) 125

Installation overview

The Base component InstallShield Wizard allows you to install the Base component by itself or with one or more plug-ins.

Note: Starting with Cluster Enabler V4.0, the Cluster Enabler Base component is a prerequisite for the Cluster Enabler plug-ins, and therefore must be installed along with or prior to the plug-ins.

It is recommended that you contact EMC Customer Support for assistance if any of the following issues apply:

- ◆ You have applications already layered with dependencies.
- ◆ You need other devices online.
- ◆ You are not confident about installing and configuring new software within the context of Windows Server 2008 or 2012, or Microsoft Failover Clusters.

IMPORTANT

Upgrading to Cluster Enabler version 4.1.4 is supported for versions 3.x and higher. Refer to [“Before you begin” on page 120](#) for additional requirements before upgrading.

Before you begin

Before you begin to install the Cluster Enabler Base component, you should read the following requirements and considerations:

- ◆ The following Windows processor architectures are supported:
 - x86
 - x64 (AMD64 and Intel EM64T)

Note: Microsoft does not support mixed architecture clusters. All nodes must be the same Windows architecture.

- ◆ To use Remote desktop in console Mode, follow these steps:
 1. Click **start** and select **Run**.
 2. Enter:


```
mstsc/admin/v:<host name>
```
 3. Click **OK**.
- ◆ Installation requires that Microsoft Windows Installer version 4.5 first be installed.
- ◆ Installation requires that all nodes first be installed with the Failover Cluster feature.
- ◆ For Failover Cluster on Windows Server 2008 or 2012, Microsoft Cluster Validate must pass all tests except storage.
- ◆ Upgrade scenarios where the storage is being replaced is not supported.
- ◆ Configurations where the cluster node is zoned to both local and remote storage arrays are not supported.

- ◆ Installation on Windows Server 2008 R2 Core requires additional configuration steps. [“Configuring a CE cluster on Server Core” on page 124](#) provides configuration instructions.

Installing the Cluster Enabler Base Component

This section explains the methods for installing the Cluster Enabler Base Component:

- ◆ Installing the Base Component separate from the plug-ins (clean install).
- ◆ Installing the Base Component along with the plug-ins (clean install).
- ◆ Installing the Base Component while upgrading.

Note: Before starting either of the procedures in this section, be sure to review [“Before you begin” on page 120](#) of this guide. In addition, if you are installing a plug-in along with the Base component, be sure to review [“Installation overview” on page 120](#).

Installing the Base Component separate from the plug-ins (clean install)

To install the Base Component separate from the plug-ins:

1. Run the Base Component installation program (EMC_CE_BASE_4.1.4.zip) downloaded from EMC online support.
2. Complete the steps in the InstallShield Wizard.
3. When prompted to restart your system, click **Yes**.

You have now finished installing the Base Component.

Installing the Base Component along with the plug-ins (clean install)

To install the Base Component along with the plug-ins:

1. Create a temporary directory. Download (save) the Base component (EMC_CE_BASE_4.1.4.zip) downloaded from EMC online support.
2. Download the plug-ins from EMC online support to the temporary directory you just created, being sure not to rename it.
3. In the temporary directory, navigate to your operating system's directory (either x64 or x86), and run the EMC_CE_Base.msi file to launch the InstallShield Wizard.
4. Complete the steps in the InstallShield Wizard.
5. When prompted to restart your system, click **Yes** to restart the system, or **No** to restart it at a later time.

Upgrading the Base Component along with the plug-ins

To upgrade the Base Component along with the desired plug-in modules:

Note: On Windows XP SP3, Windows Vista SP1, Windows 7, and Windows Server 2008 (including R2) select **Start ->Run**, and type in **mstsc /admin /v: <host name>** to use Remote Desktop in the Console Mode.

1. Move all cluster groups to node A.
2. Perform the following actions on all other cluster nodes:
 - a. Copy the `setup.exe`, `EMC_CE_Base.msi`, and `.msi` files for the plug-ins to the same local folder on your host.
 - b. Click `setup.exe` to launch the installation.
 - c. A Plug-in Selection dialog box displays the available plug-in modules. Select your desired plug-in modules to be installed.
 - d. Complete the steps in the InstallShield wizard, being sure to select the Upgrade path.
 - e. When prompted to restart your system, click **Yes**.
 - f. After the node has finished rebooting, log onto the node. Using the Cluster Manager verify that the cluster service is up.
3. After all other nodes are up, move all groups from node A to one of the other nodes. If using a shared quorum cluster model, verify that the quorum group comes online on the other node before continuing.
4. Repeat step 2 on node A.

Upgrading only the Base Component

To upgrade only the CE Base Component:

Note: On Windows XP SP3, Windows Vista SP1, Windows 7, and Windows Server 2008 (including R2), select **Start ->Run**, and enter **mstsc /admin /v: <host name>** to use Remote Desktop in the Console Mode.

1. Move all cluster groups to node A.
2. Perform the following actions on all other cluster nodes:
 - a. Copy the `setup.exe` and `EMC_CE_Base.msi` to the same local folder on your host.
 - b. Click `setup.exe` to launch the installation.
 - c. Complete the steps in the InstallShield Wizard, being sure to select the Upgrade path.
 - d. When prompted to restart your system, click **Yes**.
3. After all other nodes are up, move all groups from node A to one of the other nodes. If using a shared quorum cluster model, verify that the quorum group comes online on the other node before continuing.
4. Repeat step 2 on node A.

Uninstalling the Cluster Enabler Base Component

This section explains the methods for uninstalling the Cluster Enabler Base Component from a configured cluster:

- ◆ Uninstalling the Base Component from some cluster nodes
- ◆ Uninstalling the Base Component from all cluster nodes/deconfigure the cluster
- ◆ Uninstalling the Base Component from all cluster nodes/destroy the cluster

IMPORTANT

Uninstalling the Base component will also uninstall the Cluster Enabler plug-ins.

Uninstalling the Base component from some cluster nodes

To remove some cluster nodes and leave Cluster Enabler on the remaining cluster nodes:

1. Open Microsoft Cluster Administrator.
2. Ensure no cluster resource groups are owned by the nodes you will remove. Move any owned resource groups to a different node.
3. Right-click the nodes to remove, and choose **Stop Cluster Service**. Wait for the cluster service to stop on the nodes as indicated by a red X.
4. Right-click the nodes you want to remove and choose **Evict**. Evicting a node uninstalls the cluster service on that node and removes that node from the cluster.
5. After evicting nodes, open CE Manager, and right-click the **cluster name**. Choose **Storage Discover**, and follow through the procedure steps to complete the Storage Discover Wizard.

Note: If CE Manager is already open, perform a refresh before running the Storage Discover Wizard.

6. Uninstall CE from the evicted nodes. Use the **Add/Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler Base Component. Reboot when prompted to complete the uninstall.

Uninstalling the base component from all cluster nodes/deconfigure the cluster

Follow these steps to uninstall Cluster Enabler from all nodes of the cluster and deconfigure the CE cluster. (The Windows Server failover cluster will be maintained.):

1. Move all resource groups to the nodes on one site (that is Site A).
2. Right-click only the nodes on the remote site (that is Site B), and choose **Evict**.
3. After evicting the nodes on the remote site, open CE Manager on a node at Site A, and right-click the **cluster name**. Choose **Storage Discover**, and follow through the procedure steps to complete the Storage Discover Wizard.

Note: If CE Manager is already open, perform a refresh before running the Storage Discover Wizard.

4. From the CE Manager, select **Deconfigure CE**.
5. Uninstall CE from all nodes. Use the **Add or Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler Base Component. Reboot when prompted to complete the uninstall.

Uninstalling the base component from all cluster nodes/destroy the cluster

To uninstall Cluster Enabler from all nodes of the cluster and destroy the cluster:

1. Deconfigure the cluster according to steps 1 through 4 in “[Uninstalling the base component from all cluster nodes/deconfigure the cluster](#)”.
2. Destroy the cluster using Microsoft Cluster Administrator.
3. Uninstall CE from all nodes. **Use the Add or Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler Base Component. Reboot when prompted to complete the uninstall.

Configuring a CE cluster on Server Core

The following instructions are provided as guidelines for configuring and managing a CE cluster for a Windows Server 2008 R2 and 2012 Core installation.

Requirements and considerations

Before you begin to install the Cluster Enabler Base component on a Windows Server 2008 R2 Core edition, you should read the following requirements and considerations:

- ◆ Managing an R2 Server Core cluster can be done from a Windows Server 2008 or Windows Server 2008 R2 host.
- ◆ The remote host where you are managing the R2 Server Core cluster must be on the same domain as the R2 Server Core hosts.
- ◆ The following website provides useful information for Server Core installations:

<http://technet.microsoft.com/en-us/library/cc753802.aspx>

R2 Server Core configuration

To configure and install the Base component on a Windows Server 2008 R2 Core edition:

1. Create a failover cluster on an R2 Server Core host.

To enable Failover Clustering, enter the following from a command line on a R2 Server Core host:

```
Ocsetup FailoverCluster-Core Hyper-V
```

To enable Hyper-V, enter the following:

```
Ocsetup Microsoft-Hyper-V
```

Create a failover cluster using `cluster.exe`. You can create a cluster from Failover Cluster Manager on the remote host.

Open Failover Cluster Manager on a remote host. If you have not created a cluster on the R2 Server Core hosts, create it using Failover Cluster Manager. Once the cluster is created, connect to this cluster.

2. Install .Net Framework, Microsoft Visual C++ 2005 Redistributable x86, Solutions Enabler and CE on R2 Server Core hosts.

Note: Microsoft Visual C++ 2005 Redistributable x86 (VCREDI~3.EXE) must be installed before attempting to install CE V4.1.4.

To install .Net Framework 3.5 for Windows Server 2008, enter the appropriate Server Core command from a command line on a R2 Server Core host:

```
#To install NetFx2-ServerCore
Dism /online /enable-feature /featurename:NetFx2-ServerCore

#To install NetFx2-ServerCore-WOW64
Dism /online /enable-feature /featurename:NetFx2-ServerCore-WOW64

#To install NetFx3-ServerCore
Dism /online /enable-feature /featurename:NetFx3-ServerCore

#To install NetFx3-ServerCore-WOW64
Dism /online /enable-feature /featurename:NetFx3-ServerCore-WOW64
```

For Windows Server 2012 Core, use the following command to install .Net Framework 3.5, where source drive points to the sxs folder from the Windows Server 2012 setup DVD:

```
dism /online /enable-feature /featurename:NetFX3 /all
/Source:d:\sources\sxs /LimitAccess
```

Follow the installation instructions for [“Installing the Base Component along with the plug-ins \(clean install\)” on page 121](#).

3. Manage an R2 Server Core cluster from a remote host.
 - Install EMC CE Base Component on the remote host and reboot after prompted.
 - Run CE Configuration Wizard on the remote host to convert the R2 Server Core cluster to a CE cluster.

Upgrading Windows Server 2008 to Windows Server 2008 R2

Upgrading Windows Server 2008 SP1 or SP2 to Windows Server 2008 R2 while Cluster Enabler is installed on the host is not supported.

IMPORTANT

Attempting to upgrade to Windows Server 2008 R2 with Cluster Enabler version 4.1.4 installed causes undesirable results.

Follow these steps to prepare your host before upgrading from Windows Server 2008 SP1 or SP2 to Windows Server 2008 R2:

1. From the CE Manager, select **Deconfigure CE**.

2. Uninstall the CE Base Component from all nodes. Use the **Add or Remove Programs** utility in the **Control Panel** to remove EMC Cluster Enabler Base Component. Reboot when prompted to complete the uninstall.
3. Follow the Windows Server 2008 R2 upgrade instructions for upgrading your operating system. Refer to the Microsoft Technet article titled, *Understanding the Process of Migrating to a Cluster Running Windows Server 2008 R2*, available at:

<http://technet.microsoft.com/en-us/library/cc731812.aspx>

4. Install the Cluster Enabler Base Component and any plug-ins on all nodes.
5. From CE Manager, use the Configuration Wizard to configure the CE cluster.

GLOSSARY

This glossary contains terms related to the Cluster Enabler software.

A

- active state** The state in which a MirrorView remote mirror is running normally. See also, [“remote mirror states”](#).
- agent** An installed program designed to control a particular resource type. Each type of resource supported in a cluster is associated with an agent.
- Application Program Interface (API)** A language and message format used by an application program to communicate with another program that provides services for it. APIs are usually implemented by writing function calls. Examples of APIs are the calls made by an application program to such programs as an operating system, messaging system, or database management system. See also, [“SYMAPI”](#) and [“CLARAPI”](#).
- asynchronous mode** See [“SRDF Asynchronous \(SRDF/A\)”](#).
- attention state** The MirrorView mirror's secondary image is fractured, and the mirror is configured to generate an alert in this case. The mirror continues to accept server I/O in this state.
- Auto recovery** A MirrorView option to have synchronization start as soon as a system-fractured secondary image is determined to be reachable.
- availability** The ability to continue to provide a service even during hardware or software failure.

B

- BCV device** A Symmetrix business continuance volume (BCV) that functions as a mirrored media to a standard device for a protected storage environment.
- BCV mirror** A Symmetrix BCV device upon establishing or reestablishing a BCV pair.
- BCV pair** A standard Symmetrix device and a BCV device that provide a protected storage environment.
- business continuance** An SRDF function that ensures business applications continue running despite possible disk failures.

C

- cache** Random access electronic storage used to retain frequently used data between the CPU and either a hard disk or slower RAM. It speeds up general data flow because a cache can be accessed quickly.
- CDP** See [“continuous data protection \(CDP\)”](#).

channel director	The component in the Symmetrix array that interfaces between the host channels and data storage. It transfers data between the channel and cache.
CLARAPI	CLARiiON Application Program Interface. See “Application Program Interface (API)” .
client	<p>A computer using services or resources provided by a remote machine, called a server. Often, communications software has a separate version for the client, or guest, and the server, or host.</p> <p>Clients create a TCP/IP session with a service in the cluster using a known IP address. This address appears to the cluster software as a resource in the same group as the application providing the service. In a failure, the Cluster Service moves the entire group to another system.</p>
client failover	The response of a client machine after resource failure on the server for the client caused a resource failover. A client detects a failure in the session and reconnects in exactly the same manner as the original connection. The IP address is now available on another machine, and the connection is quickly reestablished. In this simple case, all information related to the original session not committed to disk is lost. This provides higher availability, but no fault tolerance for the service. Applications can use transactions to guarantee the client request is committed to the server database to gain fault-tolerant semantics.
CLR	See “continuous local and remote replication (CLR)” .
cluster	A group of two or more independent computers addressed and used as a single system.
cluster-aware software	Software that provides a restart mechanism invoked whenever the application resource is moved to another node in the cluster.
cluster service	The collection of software on each node that manages all cluster-specific activity.
Cluster Shared Volumes	Cluster Shared Volumes (CSV) is a Microsoft Failover Clustering feature that allows all nodes in a cluster concurrent access to data on every CSV-enabled shared disk.
consistency group	<p>A set of MirrorView logical units that are mirrored in a way that allows a recoverable copy in the event of a disaster.</p> <p>For RecoverPoint, a consistency group is a data set consisting of the production source and its replicas. A consistency group comprises the production source volumes and either a local replica, remote replica, or both. Each consistency group contains as many replication sets as there are volumes in the production storage to replicate.</p>
consistency group condition	Displays more detailed information about the MirrorView consistency group, including whether the group is active, inactive, admin fractured, system fractured, waiting on admin, or invalid.
consistency group state	Indicates the current state of the MirrorView consistency group: synchronized, consistent, synchronizing, out-of-sync, scrambled, empty, incomplete, or local only.
consistent state (of image)	State in which a MirrorView secondary image is identical to either the current primary image or to some previous instance of the primary image.

continuous asynchronous	A RecoverPoint replication mode where each write transaction is acknowledged locally at the source side and then sent to the target side. The primary advantage of continuous-asynchronous replication is its ability to provide synchronous-like replication without degrading the performance of host applications.
continuous data protection (CDP)	A RecoverPoint configuration that uses a methodology that continuously captures or tracks data modifications and stores changes independent of the primary data, enabling recovery points from any point in the past. CDP provides fine granularities of restorations to infinitely variable recovery points.
continuous local and remote replication (CLR)	A RecoverPoint configuration that includes both a CDP and a CRR copy, providing concurrent local and remote data protection. In RecoverPoint, the CDP copy is normally used for operational recovery, while the CRR copy is normally used for disaster recovery.
continuous remote replication (CRR)	A Recover Point configuration where data is transferred between two sites over Fibre Channel or a WAN. In this configuration, the RPAs, storage and splitters exist at both the local and the remote site.
continuous synchronous	A RecoverPoint replication mode. In continuous synchronous replication, the host application that initiates the write waits for an acknowledgment from the replica before continuing. Replication in synchronous mode produces a replica that is 100% up to date with the production source.
create mirror	To establish a remote mirror, that is, use the remote mirror software to create data structures on one or more LUNs on specific storage systems, such that one is the primary image and the other is a secondary image.
CRR	See “continuous remote replication (CRR)” .
D	
data center migrations	A function that reduces application outage to minutes instead of hours.
dependency	The requirement of one resource needing another resource to function properly. The Cluster Enabler resource becomes a dependency for physical disk resources in the cluster. Therefore, any operations performed on the disk resource cannot be completed until the Cluster Enabler resource has been invoked.
device	A uniquely addressable part of the storage array consisting of a set of access arms, the associated disk surfaces, and the electronic circuitry required to locate, read, and write data. Also called a LUN (logical unit number).
device group	A grouping of several devices established to provide configuration, status, and performance data on the collective devices within the group.
device number	The value that logically identifies a disk device in a string. See also “LUN” .
director	The component in the Symmetrix array that allows the Symmetrix array to transfer data between the host channels and disk devices. See also <i>channel director</i> and <i>disk director</i> .
disaster recovery	A function that recovers data at the disaster recovery site in minutes rather than days.

discover A discover action performed in the Cluster Enabler Configuration Wizard scans the storage array connected to the current node and gathers device information.

disk director The component in the Symmetrix array that interfaces between cache and the disk devices.

E

establish A BCV process that assigns a BCV device as the next available mirror of a standard device.

established The BCV pair condition where the BCV device and standard device are synchronized and functioning as a Symmetrix mirror. A BCV pair is established by the BCV commands establish, reestablish, restore, or incremental restore.

F

failback The action of moving a resource back to the cluster member designated to be the resource's Preferred Owner. By default, resources are owned by their Preferred Owner, so a failback would only occur if the resource moved from its Preferred Owner. This is likely the result of a failover.

failover The process of taking one or more resources offline on one cluster member and bringing them online on another cluster member.

fault-tolerant Continuous operation in case of failure. A fault-tolerant system can be created using two or more computers that duplicate all processing, or having one system stand by if the other fails. It can also be built with redundant processors, control units, and peripherals. Fault-tolerant operation requires backup power in a power failure. It may also imply duplication of systems in disparate locations in the event of natural catastrophe or vandalism.

FDDI An acronym for Fiber Distributed Data Interface.

Fibre Channel A high-speed serial interface capable of data transfer rates of up to 400 MB/s.

Fibre Channel Director The Fibre Channel adapter (FA) in the Symmetrix subsystem that interfaces between the host Fibre Channel interface and data storage. It transfers data between the channel and cache.

forced failover A CE feature allowing you to automatically keep a cluster up on a particular array or arrays in a total site disaster.

forced quorum Software functionality allowing the cluster to be forced up in the event that total communication is lost between nodes and Microsoft Failover Cluster. Microsoft Failover Cluster wants to shut down the cluster to avoid a split-brain condition.

[See “split-brain condition”.](#)

fracture A condition in which I/O is not mirrored to the MirrorView secondary image and can be caused when you initiate the fracture (Admin Fracture) or when the system determines that the secondary image is unreachable (System Fracture). An admin fracture may also occur if the MirrorView software detects an error condition that requires administrative intervention to correct.

fracture log A bitmap, maintained in SP memory, that indicates which portions of the MirrorView primary image might differ from the secondary images. The fracture log is used to shorten the synchronization process after fractures. The bitmap is maintained in SP memory, so if the mirror is not configured to use the optional write intent log (which stores the data on disk), and the SP that controls the primary fails while the secondary image is fractured, the fracture log is lost, and full synchronization of the secondary image is required.

G

graphical user interface (GUI) A method that allows users to interact with the computer and its special applications based on graphics instead of text. GUIs use icons, pictures, and menus and use a mouse as well as a keyboard to accept input.

group A collection of resources to be managed as a single unit. Usually, a group contains all elements needed to run a specific application and for client systems to connect to the service provided by the application. Groups allow an administrator to combine resources into larger logical units and manage them as a unit. Operations performed on a group affect all resources contained within that group.

H

HBA See “host bus adapter (HBA)”.

heartbeat A polling communication mechanism used by the cluster processes to determine whether the other members of the cluster are alive and working or have failed. If the heartbeat is not functioning, a failover is initiated, and another node in the cluster takes over the services.

high availability The characteristic of a computer system/computing environment that allows it to continue to provide applications and access to data if a single component or resource fails. Service is interrupted for only a brief time, and may or may not be apparent to the end users.

host bus adapter (HBA) A device circuit board that provides an interface between the SCSI bus and the computer I/O bus (for example, PCI, EISA, microchannel).

hyper-volume The term used by the Symmetrix array to make a physical disk appear as multiple physical disks. Each hypervolume has its own unique SCSI address.

I

I/O Input/output.

identifier (ID) A sequence of bits or characters that identifies a program, device, controller, or system.

image condition The condition of a MirrorView secondary image provides additional information about the status of updates for the image. Values include normal, administratively fractured, system fractured, queued to be synchronized, synchronizing, or waiting-on-admin.

image state Indication of the relationship between a MirrorView secondary image and the primary image of a mirror. The image states are: synchronized, consistent, synchronizing, and out-of-sync.

incremental establish A TimeFinder BCV or SRDF control operation.

For BCV control operations, an *incremental establish* causes the BCV device to be incrementally synchronized and functioning as a Symmetrix mirrored device. (The devices must have been previously paired.) This is the same as an *establish* operation except an incremental establish is much faster: It copies only the differences or new storage data from the standard device to the BCV device. Any changed tracks on the BCV device are overwritten by the data on the corresponding tracks from the standard device.

For SRDF control operations, an *incremental establish* causes the target (R2) device to be incrementally synchronized and established as a Symmetrix mirrored device. (The devices must have been previously paired.) This is the same as an *establish* operation except that an incremental establish is much faster: It copies only the differences or new storage data from the source (R1) device to the target (R2) device. Any changed tracks on the device are overwritten by the data on the corresponding tracks from the source (R1) device.

incremental restore A TimeFinder BCV or SRDF control operation.

In BCV control operations, an *incremental restore* is a control operation that reassigns a BCV device as the next available mirror of the standard device in the pair. However, the standard devices are updated with only the data written to the BCV device during the time of the original pair split. The data written to the standard device during the split is overwritten with data from the BCV mirror.

In SRDF control operations, an *incremental restore* is a control operation that reassigns a target (R2) device as the next available mirror of the source (R1) device in the pair. The source (R1) devices are updated with only the data written to the target (R2) device during the time of the original pair split. The data written to the source (R1) device during the split is overwritten with data from the target (R2) mirror.

individual mirror group When there exists only a single mirror in the mirror group, the group is called an individual mirror group. For CLARiiON arrays, the creation of a mirror group does not require the creation of a corresponding consistency group if there exists only one mirror in the mirror group. Consistency groups are automatically created and deleted when individual mirror groups are converted to/from multiple mirror groups.

L

lateral node Nodes connected to the same Symmetrix array.

LUN A logical unit number (LUN) is a unique identifier used on a SCSI bus that enables it to differentiate between up to eight separate storage devices (each of which is a logical unit). See also, “[device number](#)”.

M

Microsoft Management Console (MMC) A Microsoft user interface (UI) framework for use in administering different components of the Microsoft Windows operating platform. This framework is used to host-specific UI/control extensions called *snap-ins*. Use snap-ins to administer both local and remote computers. Third-party snap-ins can be written for use with MMC.

mirrored pair A device comprising two hypervolumes with all data recorded twice—once on each disk drive.

mirroring A device comprising two hypervolumes with all data recorded twice—once on each disk drive. The Symmetrix array maintains two or more identical copies of a set of data on separate disks. Each copy automatically updates during a write operation. If one disk device fails, the Symmetrix array automatically uses one of the other copies from another disk drive.

Also, a MirrorView feature that provides disaster recovery by maintaining one or more mirrors of LUNs on other storage systems. MirrorView can work in conjunction with, but is independent of, the other major CLARiiON software options such as PowerPath software and SnapView software. MirrorView works with LUNs in SAN storage systems, and thus can be used to mirror one or more LUNs that may compose a SAN storage group.

MMC See “[Microsoft Management Console \(MMC\)](#)”.

MSCS Microsoft Cluster Service. A shared-nothing cluster solution for Windows Server 2003. In Windows Server 2008, this is now called Microsoft Failover Cluster.

N

network interface card (NIC) A device that provides network communication capabilities to and from a computer system.

Node Majority A quorum-capable resource based on replicating data to local disks associated with a majority of cluster nodes. MNS enables you to create a server cluster without shared disk for the quorum resource. Cluster Enabler allows you to configure an MNS cluster on Windows Server 2008 and 2012 Enterprise and Datacenter Editions.

nodes Members of a cluster. Also referred to as systems. A node contains a CPU, disk, and network resource.

O

offline The state of a resource or group that classifies it as unavailable. When used in context with a cluster member, offline implies the cluster member may not be booted, or the cluster service on the node in question may not be functioning properly.

online The state of a resource or group that classifies it as available. When used in context with a cluster member, online implies the other cluster members are receiving heartbeats from the cluster member in question. See also “resource”.

out-of-sync state In MirrorView, a remote mirror state in which the software does not know how the primary and secondary images differ; therefore, a full synchronization is required to make the secondary images usable for recovery. See also, “[image state](#)”.

P

peer node Nodes connected to different Symmetrix arrays located across the link from each other.

primary image The LUN on the MirrorView production storage system that contains user data and is the source for data copied to the secondary image. For MirrorView/CE there is one primary image and one secondary image. A remote mirror is ineffective for recovery unless it has at least one secondary image. This manual also refers to primary image as primary or primary mirror image.

promote (to primary) The operation by which the administrator changes a MirrorView image's role from secondary to primary. As part of this operation, the previous primary image becomes a secondary image. If the previous primary image is unavailable when you promote the secondary image (perhaps because the primary site suffered a disaster), the software does not include it as a secondary image in the new mirror. A secondary image can be promoted if it is in either the synchronized state or the consistent state. An image cannot be promoted if it is out-of-sync or synchronizing.

Q

query A command reporting the state of all the BCV devices in the system, as well as the status of SRDF states.

quiesce threshold The time period after which, without I/O from the server, any MirrorView secondary image in the consistent state and not fractured is marked as being in the synchronized state (the default is 60 seconds).

quorum disk An ordinary disk volume used as a special communication mechanism between server systems. In a Microsoft failover cluster, a small amount of cluster system data (a few megabytes) is stored on this volume. The SCSI-3 `Reserve` and `Reset` commands are used to move quorum-disk ownership back and forth between nodes. If the heartbeat mechanism fails, the quorum disk is used for each node to verify whether the other node is still functioning. Because not all disk products implement these multihost SCSI-3 commands, not all disk products will work in a failover cluster environment. Thus, Microsoft is very rigorous in providing the Cluster/RAID category of tests to qualify disks (refer to Microsoft's Hardware Compatibility List) capable of running with Microsoft failover cluster software).

R

R1 device See "[source \(R1\) device](#)".

R2 device See "[target \(R2\) device](#)".

RA Remote adapter. An RA provides the link connection and fiber optic protocol support between the local and remote Symmetrix arrays. The RA cable connection is ESCON fibre (ESCON protocol).

RAID Redundant array of independent disks. Data is stored on multiple magnetic or optical disk drives to increase output performance and storage capacities and to provide varying degrees of redundancy and fault tolerance. Instead of storing valuable data on a single hard disk that could fail at any time, RAID ensures a backup copy of all information always exists by spreading data among multiple hard disks.

RDF1/RDF2 A type of SRDF device group. Only RDF devices can be assigned to an RDF group. An RDF1 group type contains source (R1) devices and an RDF2 group type contains target (R2) devices.

RDF State The SRDF state information displayed in a column in the middle pane.

The possible RDF states are the following:

Invalid — The device and link are in an unrecognized combination.

SynclnProg — Synchronizing in progress.

Synchronized — The source and target have identical data.

Split — The source is split from the target and the target is enabled.

Suspended — The link is suspended.

Failed Over — The target is write-enabled, the source is write-disabled and the link is suspended.

Partitioned — The communication link to the remote Symmetrix array is down and the device is write-enabled.

R1 Updated — The target is write-enabled, the source is write-disabled and the link is up.

R1 UpdInProg — The target is write-enabled, the source is write-disabled, the link is up, but there are invalid tracks between the target and the source.

Mixed — This state is only set for a `SymDgShow()` call when the RDF states of the devices in the group are different from each other, thereby making the RDF state of the group *mixed*.

N/A — Not applicable.

Consistent — R2 data is consistent.

recovery policy	In MirrorView, the policy for recovering the secondary mirror image after a system fracture. If the recovery policy is set to Auto, then the secondary starts re-synchronizing as soon as the primary image determines that the secondary mirror image is once again accessible. If the policy is set to Manual, then an administrator must explicitly start a synchronization operation to recover the secondary mirror image.
reestablish	A business continuance process that reassigns a BCV device as the next available mirror of the standard device with which it was previously paired. The BCV mirror is updated with the data written to the standard device during the period the BCV pair was split. The data written to the BCV device during the split is overwritten by data from the standard device.
Remote Link Director (RLD)	<p>RLDs create the data link paths between two data storage units. Each Symmetrix array requires a minimum of two, up to a maximum of eight RLDs, depending on the Symmetrix model in use. Each RLD manages two ESCON fibre link connections. Each RLD can perform a single I/O at a time to its paired RLD in the remote Symmetrix array.</p> <p>RLDs have either an RA1 or RA2 designation. RA1s reside in the source Symmetrix array. RA2s reside in the target Symmetrix array. These RLDs can also be assigned to an RA group.</p> <p>See also “RA”.</p>
remote mirror	<p>For Symmetrix, the remote mirror refers to a target (R2) device located in a remote Symmetrix array. When a source (R1) device is participating in SRDF operations with a target (R2) device, all writes to the R1 device are mirrored to a target (R2) device in a remote Symmetrix array.</p> <p>For MirrorView, a remote mirror is the combination of a LUN on one storage system, called the primary image, and another LUN on a different storage system, called the secondary image. The software maintains the secondary image as an exact copy of the primary image</p>

at some (possibly previous) point in time. If the server and/or storage system at the primary site fails, you can promote the secondary image to take over the role of the primary, thus allowing continued access to your production data.

remote mirroring	A feature that provides the means for disaster recovery by maintaining one or more copies (images) of LUNs at separate locations.
remote mirror states	<p>There are three types of MirrorView mirror states. The mirror states are active, inactive, and attention.</p> <p>Active — The remote mirror is running normally.</p> <p>Inactive — I/O is rejected. This can be a temporary state during some consistency group operations or a result of an error during a consistency group operation.</p> <p>Attention — The state to alert you that the minimum number of images required is not currently met. A fracture or the removal of an image can cause this. The mirror will continue to accept I/O in this state.</p>
Replication set	A RecoverPoint term. A storage volume in the production source that is replicated must have a corresponding volume at each copy. A replication set is production volume and its associated volume at the local copy, the remote replica, or both.
resource	An object managed by the Cluster Service that sees all resources as identical opaque objects. Resources may include physical hardware devices, such as disk drives and network cards, or logical items, such as disk partitions, TCP/IP addresses, entire applications, and databases. A resource is said to be online on a node when it is providing its service on that specific node.
resource failback	The movement of resources back to their preferred location in the cluster. This is usually done under manual user control to avoid a situation where a resource is failed back, and then immediately fails over again because of an unresolved node problem. Microsoft Failover Cluster also allows automatic failback and provides a timing window to try to avoid repeated failovers.
resource failover	The process where control of a resource moves to another node of a cluster. Failover can be initiated automatically or manually. When initiated automatically, the cluster management software detects a failure of server node hardware or an application. When manually initiated, the cluster administrator uses the Cluster Administrator software application.
resource group	A collection of resources to be managed as a single unit. Usually a group contains all elements needed to run a specific application, and for client systems to connect to the service provided by the application. Groups allow an administrator to combine resources into larger logical units and manage them together. Operations performed on a group affect all resources contained within that group.
restore	<p>A TimeFinder BCV or SRDF control operation.</p> <p>In BCV control operations, a restore copies a full BCV mirror back to the standard device in the pair and reassigns the BCV device as the next available mirror to the standard device.</p> <p>In SRDF control operations, a restore copies the full target (R2) device back to the source (R1) device in the pair and reassigns the target (R2) device as the next available mirror to the source (R1) device.</p>

See also *incremental restore*.

RF A remote adapter that provides the link connection and fiber optic protocol support between the local and remote Symmetrix arrays. The RF cable connection is SCSI fibre (SCSI protocol). An RF differs from an RA only in the type of connection; an RA uses an ESCON fibre connection, and an RF uses a SCSI fibre connection.

See also [“RA”](#).

S

scalability The ability to add new components to a storage system as system load increases.

SCSI Small Computer System Interface. SCSI is a high-speed parallel interface used to connect microcomputers to SCSI peripheral devices, such as disks, printers, and other computers and local area networks.

secondary image For MirrorView, a LUN that contains a mirror of the primary image LUN.

secondary image state The secondary image states are synchronized, consistent, synchronizing, and out-of-sync. They describe the data on the secondary storage system in relation to the data on the primary storage system.

semisynchronous mode An SRDF mode of operation that provides an asynchronous mode of operation. Applications are notified an I/O (or I/O chain) is complete once the data is in the cache of the local RA1 Symmetrix array. Any new data is then written to cache in the remote RA2 Symmetrix array. The remote Symmetrix array acknowledges receipt of the data once it is secure in its cache. If source tracks are pending transfer to a target (R2) device, and a second write is attempted to the source (R1) device, the Symmetrix array disconnects (non-immediate retry request), and waits for the pending track to transfer to the remote Symmetrix array.

snap-in See [“Microsoft Management Console \(MMC\)”](#).

snapshot A RecoverPoint term. A snapshot is the difference between one consistent image of stored data and the next. Snapshots are taken seconds apart. The application writes to storage; at the same time, the splitter provides a second copy of the writes to the RecoverPoint appliance.

snapshot replication mode A RecoverPoint replication mode that only transfers data that has changed between one consistent image of the storage subsystem and the next. By definition, snapshot replication produces a replica that is not up to date.

Solutions Enabler Also known as SYMCLI, an application written using the Symmetrix Application Programming Interface (SYMAPI) that retrieves data from a Symmetrix array using special low-level SCSI commands.

Solutions Enabler allows you to run commands on the host to obtain configuration, status, and performance data from the Symmetrix arrays attached to hosts that are running in an open systems environment.

SYMCLI SRDF and TimeFinder components allow you to perform control operations on RDF and BCV devices.

source (R1) device	<p>A Symmetrix source (R1) device that is participating in SRDF operations with a target (R2) device. All writes to this device are mirrored to a target (R2) device in a remote Symmetrix array. An R1 device must be assigned to an RDF1 group type.</p> <p>See also <i>RDF1/RDF2</i>.</p>
source unit	<p>In an SRDF configuration, it is the primary data storage subsystem. It initiates many of the SRDF synchronization activities. An SRDF configuration must have at least one source unit and one target unit. See also <i>target unit</i>.</p>
split	<p>A business continuance process that removes the BCV mirror from the existing BCV pair and assigns the BCV mirror back to its original device address. The BCV device then holds an instant copy of the data from the standard device.</p>
split-brain condition	<p>A total communication failure while both nodes remain operational. A split-brain condition is a potential cause of logical data corruption. For example, if both sides assume the other is dead and begin processing new transactions against their copy of the data, two separate and unreconcilable copies of the data can be created.</p>
SRDF	<p>Symmetrix Remote Data Facility. SRDF consists of the microcode and hardware required to support Symmetrix remote mirroring.</p>
SRDF Asynchronous (SRDF/A)	<p>A high-performance, extended-distance asynchronous replication using a delta set architecture for reduced bandwidth requirements and no host performance impact.</p> <p>Asynchronous mode provides a point-in-time image on the target (R2) device only slightly behind the source (R1) device. SRDF/A session data is transferred to the remote Symmetrix system in delta sets, eliminating the redundancy of same-track changes being transferred over the link, reducing the required bandwidth. SRDF/A only needs enough bandwidth to support the average production workload versus peak workloads.</p> <p>SRDF/A is intended for users who require no host application impact while maintaining a consistent, restartable image of their data on the R2 side at all times.</p>
SRDF link	<p>Fiber optic connections and channels between two Symmetrix arrays. A minimum of two to a maximum of eight links can exist between the two units.</p>
stretch cluster	<p>A Microsoft cluster that is geographically distributed across multiple physical locations.</p>
SYMAPI	<p>Symmetrix Application Program Interface. See “Application Program Interface (API)”.</p>
SYMCLI	<p>See “establish”.</p>
synchronize	<p>For MirrorView, the process of updating each secondary image with changes from a primary image. There are several levels of synchronization: synchronization based on a fracture log, synchronization based on the optional write intent log, and full synchronization (a complete copy). Synchronization based on the fracture or write intent log requires copying only part of the primary image to the secondary images.</p>
synchronized state	<p>For SRDF, the state in which the data in the R1 device is identical to that of the R2 device.</p> <p>For MirrorView, the state in which the data in the secondary image is identical to that in the primary. On the next write to the primary, the image state will change to consistent. See also, “secondary image state”.</p>

synchronizing state For MirrorView, when a secondary image is in the process of synchronizing. The data in the secondary image is not usable for recovery until the synchronization operation completes. Thus, an image in the synchronizing state cannot be promoted to the primary image. See also, “[secondary image state](#)”

synchronous mode An SRDF mode of operation that ensures 100 percent synchronized mirroring between the two Symmetrix arrays. This is a synchronous mode of operation. Applications are notified that an I/O (or I/O chain) is complete when the RA2 Symmetrix array acknowledges that the data has been secured in cache.

T

target (R2) device A Symmetrix target (R2) device participating in SRDF operations with a source (R1) device. It resides in the remote, or target, Symmetrix array. It is paired with a source (R1) device in the local Symmetrix array and receives all write data from its mirrored pair. This device is not accessed by user applications during normal I/O operations. An R2 device must be assigned to an RDF group type. See also *RDF1/RDF2*.

target unit In an SRDF configuration, this subsystem maintains synchronization with the devices it is paired with in the source unit. It can serve as a source unit during disaster recovery. An SRDF configuration must have at least one source unit and one target unit. See also *source unit*.

Thin-LUN A CLARiiON MirrorView mirror that has at least one image that is a thin-LUN.

V

virtual servers See “[nodes](#)”.

W

write intent log (WIL) For MirrorView, the WIL is a record of recent changes to the primary image. This record is stored in persistent memory on a private LUN reserved for the mirroring software. If the primary storage system fails (not catastrophically; that is, the WIL LUNs or the persistent mirror storage was lost), the optional write intent log can be used to quickly synchronize the secondary images when the primary storage system becomes available. This eliminates the need for full synchronization of the secondary images, which can be a lengthy process on very large LUNs.

workload migrations Similar to data center migrations; especially useful for minimizing outages during preventative maintenance of hardware or software.

