

Symantec™ High Availability Agent for EMC MirrorView Configuration Guide

Windows

6.0

Symantec High Availability Agent for EMC MirrorView Configuration Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Agent version: 6.0

Document version: 6.0 Rev 1

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, the Veritas Logo, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation

350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apj@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4	
Chapter 1	Introducing the Symantec High Availability agent for EMC MirrorView	9
	About the agent for EMC MirrorView	9
	Supported software	10
	Supported hardware	10
	Typical EMC MirrorView setup in a VCS cluster	10
	EMC MirrorView agent functions	12
	About the agent's online function	12
Chapter 2	Installing and removing the agent for EMC MirrorView	14
	Before you install the agent for EMC MirrorView	14
	Installing the agent for EMC MirrorView	14
	Removing the agent for EMC MirrorView	15
Chapter 3	Configuring the agent for EMC MirrorView	17
	Configuration concepts for the EMC MirrorView agent	17
	Resource type definition for the MirrorView agent	17
	Attribute definitions for the MirrorView agent	18
	Sample configuration for the MirrorView agent	19
	Before you configure the agent for EMC MirrorView	21
	About cluster heartbeats	21
	About configuring system zones in replicated data clusters	21
	About preventing split-brain	22
	Configuring the agent for EMC MirrorView	23
	Configuring the agent manually in a global cluster	23
	Configuring the agent manually in a replicated data cluster	24
	Executing the addArrayUser action	24

Chapter 4	Managing and testing clustering support for EMC MirrorView	26
	Typical test setup for the EMC MirrorView agent	26
	Testing service group migration	27
	Testing host failure	28
	Performing a disaster test	29
	Performing the failback test	29
	Failure scenarios for EMC MirrorView	30
	Site disaster	30
	All host or all application failure	31
	Replication link failure	31
	Split-brain in a MirrorView environment	31
Chapter 5	Setting up fire drill	32
	About fire drills	32
	About the MirrorViewSnap agent	33
	MirrorViewSnap agent functions	33
	Resource type definition for the MirrorViewSnap agent	33
	Attribute definitions for the MirrorViewSnap agent	34
	Sample configuration for a fire drill service group	35
	Additional considerations for running a fire drill	36
	Before you configure the fire drill service group	36
	Configuring the fire drill service group	36
	Creating the fire drill service group using Cluster Manager (Java Console)	36
	Adding resources to the fire drill service group	37
	Configuring resources for the fire drill service group	38
	Verifying a successful fire drill	38
Index	40

Introducing the Symantec High Availability agent for EMC MirrorView

This chapter includes the following topics:

- [About the agent for EMC MirrorView](#)
- [Supported software](#)
- [Supported hardware](#)
- [Typical EMC MirrorView setup in a VCS cluster](#)
- [EMC MirrorView agent functions](#)

About the agent for EMC MirrorView

The Symantec High Availability agent for EMC MirrorView provides support for application failover and recovery. The agent provides this support in environments that use MirrorView to replicate data between CLARiiON arrays.

The agent monitors and manages the state of replicated CLARiiON array LUNs that are attached to VCS nodes. The agent ensures that the system that has the MirrorView resource online also has safe and exclusive access to the configured arrays.

You can use the agent in replicated data clusters and in global clusters that run VCS.

The agent supports EMC MirrorView configured in the synchronous or asynchronous modes.

In both modes (synchronous and asynchronous), you can replicate both, consistency groups and a set of individual Remote Mirrors.

Supported software

For information on the software versions that the agent for EMC MirrorView supports, see the Symantec Operations Readiness Tools (SORT) site:

<https://sort.symantec.com/agents>.

Supported hardware

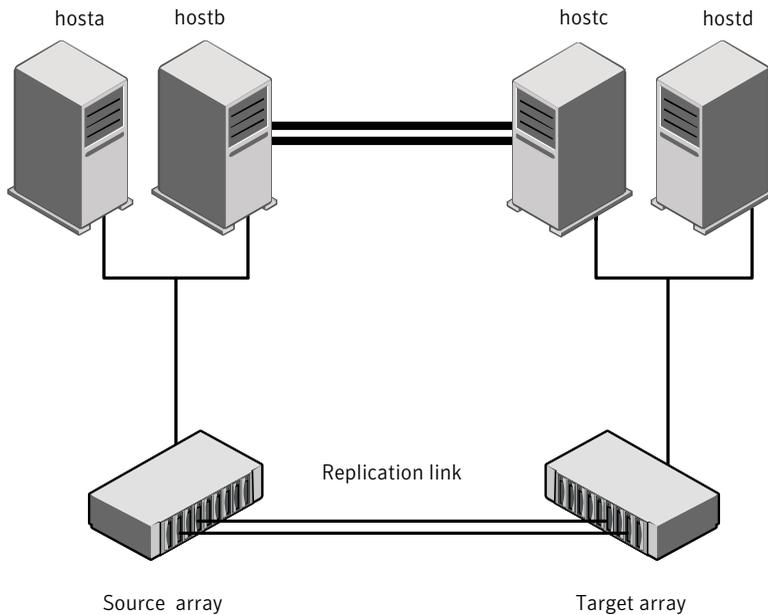
The Symantec High Availability agent for EMC Mirrorview supports Navisphere Secure CLI, naviseccli.

To determine the supported versions of NaviCLI and FLARE code that are on CLARiiON arrays, consult the EMC hardware compatibility list.

Typical EMC MirrorView setup in a VCS cluster

[Figure 1-1](#) displays a typical cluster setup in a MirrorView environment.

Figure 1-1 Typical clustering setup for the agent



Clustering in a MirrorView environment typically consists of the following hardware infrastructure:

- The source array (array1) has one or more hosts. The hosts have a direct connection to a CLARiiON array. The array contains the mirror that is the primary image and the direct connection uses either SCSI or Fibre Channel.
- The target array (array2) consists of one or more hosts. These hosts have a direct connection to another CLARiiON array. The array contains the mirror that is the secondary image and the connection uses either SCSI or Fibre Channel. The secondary image LUNs pairs with the mirrored LUNs in the source array. The target hosts and the array must be at a significant distance from the source side to survive a source-side disaster.
- Network heartbeating between the two data centers to determine their health could be LLT or TCP/IP.
See “[About cluster heartbeats](#)” on page 21.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT. In a global cluster environment, you must attach all hosts in a cluster to the same CLARiiON array.

EMC MirrorView agent functions

The VCS agent for EMC MirrorView monitors and manages the state of replicated CLARiiON LUNs attached to VCS nodes. Agent functions bring resources online, take them offline, and perform different monitoring actions.

The agent performs the following functions:

Table 1-1 Agent functions

Function	Description
online	Creates a lock file on the local host. This lock indicates that the resource is online and makes the mirrors available for the application to use. The agent performs specific actions depending on the state of the mirrors.
offline	Removes the lock file on the local host.
monitor	Verifies that the lock file exists. If the lock file exists, the monitor entry point reports the status of the resource as online. If the lock file does not exist, the monitor entry point reports the status of the resource as offline.
open	Prevents potential concurrency violation if the service group fails over to another node. Note: The agent does not remove the lock file if the agent was started after the <code>hastop -force</code> command.
clean	Removes the lock file.
info	The info function gives the information about the mirrors (in case of synchronous mode of replication). It also gives information about the mirrors/groups in case of asynchronous mode of replication. It uses the <code>-sync listsyncprogress</code> and <code>-async -list</code> or <code>-async listgroups</code> commands to get this information.
addArrayUser	This action entry point creates the user security file for the <code>navisecli</code> command.

About the agent's online function

The agent's online operation performs specific actions depending on the state of the mirrors.

If the state of all local mirrors is **MIRRORED**, the agent creates a lock file on the local host. This lock indicates that the resource is online and makes the mirrors available for the application to use.

If one or more mirrors are not in the **MIRRORED** state, the agent runs a NaviCLI command. With this command, the agent brings them into the **MIRRORED** state, which enables the application to use them.

- For secondary images in the **synchronized** state, the agent runs the `mirror -sync/-async -promoteimage/-promotegroup` command to promote the remote mirror. This command also converts the current primary to secondary.
- For secondary images in the **CONSISTENT** state, the agent waits to check if the image has transitioned to the **SYNCHRONIZED** state.
- If the images have transitioned to the **SYNCHRONIZED** state, the agent then runs the `mirror -sync/-async -promoteimage/-promotegroup` command to promote the remote mirror. This command also converts the current primary to secondary.
- If the image has not transitioned to the **SYNCHRONIZED** state, the agent checks if the remote array is accessible. If the remote array is accessible, then this condition indicates link failure—the image would be in a fractured condition.

In case of fracture:

- If the **SplitTakeover** attribute is set to 1, the agent forcibly promotes the secondary image.
- If the **SplitTakeover** attribute is set to 0, the agent does not try to promote the secondary image forcibly, and becomes the administrator's decision.

Installing and removing the agent for EMC MirrorView

This chapter includes the following topics:

- [Before you install the agent for EMC MirrorView](#)
- [Installing the agent for EMC MirrorView](#)
- [Removing the agent for EMC MirrorView](#)

Before you install the agent for EMC MirrorView

Set up your cluster. For information about installing and configuring VCS, refer to the *Symantec Cluster Server Installation Guide*.

Set up replication and the required hardware infrastructure. For information about setting up Oracle RAC environment, refer to the *Symantec Storage Foundation for Oracle RAC Installation Guide*.

See [“Typical EMC MirrorView setup in a VCS cluster”](#) on page 10.

Installing the agent for EMC MirrorView

You must install the EMC MirrorView agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

To install the VCS agent for EMC MirrorView from the Agent Pack release

- 1 Log on to any node in the cluster.
Ensure that the logged on user has the domain administrative privileges.
- 2 Download the Agent Pack from the Symantec Operations Readiness Tools (SORT) site: <https://sort.symantec.com/agents>.
You can download the complete Agent Pack zip file or the individual agent zip file.
- 3 Uncompress the file to a temporary location.
- 4 If you downloaded the complete Agent Pack zip file, navigate to the directory containing the package for the platform running in your environment.

```
Windows 2008 x64 cd1\windows\w2k8x64\vcs\replication\
mirrorview_agent\agentversion\pkgs
```

```
Windows 2008 IA64 cd1\windows\w2k8IA64\vcs\replication\
mirrorview_agent\agentversion\pkgs
```

```
Windows 2012 x64 cd1\windows\w2k12x64\vcs\replication\
mirrorview_agent\agentversion\pkgs
```

- 5 Double-click **vrtsvcsmirrorview.msi**.
Follow the instructions that the install program provides, to complete the installation of the agent.

Removing the agent for EMC MirrorView

This section describes steps for uninstalling the agent. Do not attempt to remove the agent if service groups accessing the shared storage are online.

To remove the agent for MirrorView that was installed from an Agent Pack release

- 1 Open the Windows Control Panel and click **Add or Remove Programs** or **Uninstall a Program**.
- 2 Select the VCS Agent for EMC MirrorView.
- 3 Click the **Remove** or **Uninstall** button.
- 4 Follow the instructions the installer provides to complete the uninstallation.

Note: For Win IA64 and Win x64 architectures, you must manually delete the agent directory if it is not removed after the uninstallation.

Configuring the agent for EMC MirrorView

This chapter includes the following topics:

- [Configuration concepts for the EMC MirrorView agent](#)
- [Before you configure the agent for EMC MirrorView](#)
- [Configuring the agent for EMC MirrorView](#)

Configuration concepts for the EMC MirrorView agent

Review the resource type definition and the attribute definitions for the agent.

Resource type definition for the MirrorView agent

The resource type definition defines the agent in VCS.

```
type MirrorView (
    static keylist SupportedActions = { addArrayUser }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static int OfflineMonitorInterval = 0
    static int RestartLimit = 1
    static i18nstr ArgList[] = { NaviCliHome, LocalArraySPNames,
    RemoteArraySPNames, Mode, GrpName, MirNames, SplitTakeover,
    LocalArraySecurityFilePath, RemoteArraySecurityFilePath }
    str NaviCliHome = "C:\\Program Files\\EMC\\Navisphere CLI"
    str LocalArraySPNames[]
    str RemoteArraySPNames[]
    str Mode
```

```

str GrpName
str MirNames[]
boolean SplitTakeover
temp str VCSResLock
str LocalArraySecurityFilePath
str RemoteArraySecurityFilePath
)

```

Attribute definitions for the MirrorView agent

Review the description of the agent attributes.

Required attributes

You must assign values to the following required attributes:

Table 3-1 Required attributes

Attribute	Description
NaviCliHome	NaviCLI installation directory "C:\\Program Files\\EMC\\Navisphere CLI" Type-Dimension: string-scalar
LocalArraySPNames	The list of storage processors within the array to which the local hosts are connected. Can be names or IP addresses. Type-Dimension: string-vector
RemoteArraySPNames	The list of storage processors within the array to which the remote hosts are connected. Can be names or IP addresses. Type-Dimension: string-vector
Mode	The replication mode, which is either: <i>sync</i> or <i>async</i> . Type-Dimension: string-scalar
GrpName	The name of the consistency group to which the mirrors belong. Type-Dimension: string-scalar
MirNames	The list of Remote Mirrors that this resource must manage. All the mirrors must have the same mode of replication, that is, synchronous or asynchronous. Type-Dimension: string-vector

Table 3-1 Required attributes (*continued*)

Attribute	Description
SplitTakeover	<p>This integer indicates whether VCS should forcefully promote a secondary to a primary.</p> <p>In case of a link-failure between the two arrays, the state of the mirror remains consistent or out of sync. Under such circumstances, if the application has to failover—due to disaster or user-driven action—mirrors are not in a SYNCHRONIZED state.</p> <p>If the value of the SplitTakeOver attribute is 1:</p> <ul style="list-style-type: none"> ■ The agent fails over when it discovers link failures ■ The agent determines that mirrors are out of sync <p>If the value of the attribute is 0, agent does not fail over and the administrator must determine what to do.</p> <p>Type-Dimension: boolean-scalar</p>

Optional attributes

Assign values to the following optional attributes.

Table 3-2 Optional attributes

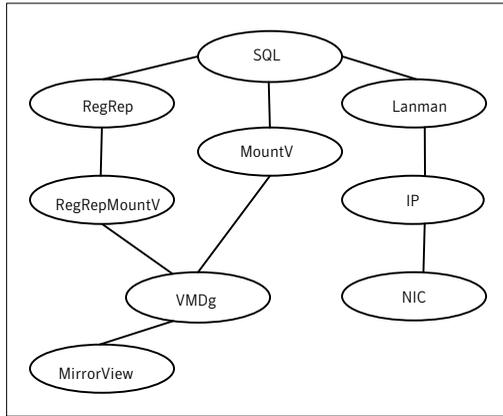
Attribute	Description
LocalArraySecurityFilePath	<p>The absolute path to the NaviSecCli security file, which contains credentials for the local array. If you do not assign a value to this attribute, the agent assumes that the security file is placed in the \$HOME of the root user.</p> <p>Type-Dimension: string-scalar</p>
RemoteArraySecurityFilePath	<p>The absolute path to the NaviSecCli security file containing credentials for the remote array. If you do not assign a value to this attribute, the agent assumes that the security file is placed in the \$HOME of the root user.</p> <p>Type-Dimension: string-scalar</p>

Sample configuration for the MirrorView agent

Figure 3-1 shows a VCS service group that has a resource of type MirrorView.

The VMDg resource depends on the MirrorView resource.

Figure 3-1 Dependency tree



You can configure a resource of type MirrorView in the `main.cf` file.

In this example, the resource is configured for asynchronous mode and consistency groups.

```

MirrorView mir (
    NaviCliHome = "C:\\Program Files\\EMC\\Navisphere CLI"
    LocalArraySPNames @sys1= { "Local_SP1_Name", "Local_SP1_IP" }
    LocalArraySPNames @sys2 = { "Local_SP2_Name", "Local_SP2_IP" }
    RemoteArraySPNames @sys1 = { "Local_SP2_IP", "Remote_SP2_Name" }
    RemoteArraySPNames @sys2= { "Local_SP1_IP", "Remote_SP1_Name" }

    Mode = async
    GrpName = consistency_grp1

    SplitTakeover = 0
)

MirrorView mir (
    NaviCliHome = "C:\\Program Files\\EMC\\Navisphere CLI"
    LocalArraySPNames = { "Local_SP1_IP", "Local_SP2_IP" }
    RemoteArraySPNames = { "Remote_SP1_IP", "Remote_SP2_IP" }
    Mode = sync
    GrpName = sync_consistency_grp1
    SplitTakeover = 0
)
  
```

If you want to configure the resource and specify the individual mirror names, configure the `MirNames` attribute, instead of the `GrpNames` attribute, as follows:

```
MirNames = { "sync_mir1", "sync_mir2" }  
GrpName = ""
```

Before you configure the agent for EMC MirrorView

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
See [“Typical EMC MirrorView setup in a VCS cluster”](#) on page 10.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 21.
See [“About preventing split-brain”](#) on page 22.
- Set up system zones in replicated data clusters.
See [“About configuring system zones in replicated data clusters”](#) on page 21.

About cluster heartbeats

In a replicated data cluster, ensure robust heartbeating by using dual, dedicated networks over which the Low Latency Transport (LLT) runs. Additionally, you can configure a low-priority heartbeat across public networks.

In a global cluster, VCS sends ICMP pings over the public network between the two sites for network heartbeating. To minimize the risk of split-brain, VCS sends ICMP pings to highly available IP addresses. VCS global clusters also notify the administrators when the sites cannot communicate.

Heartbeat loss may occur due to the failure of all hosts in the primary cluster. In such a scenario, a failover may be required even if the array is alive. In any case, a host-only crash and a complete site failure must be distinguished. In a host-only crash, only the ICMP heartbeat signals a failure by an SNMP trap. No cluster failure notification occurs because a surviving heartbeat exists. This trap is the only notification to fail over an application.

About configuring system zones in replicated data clusters

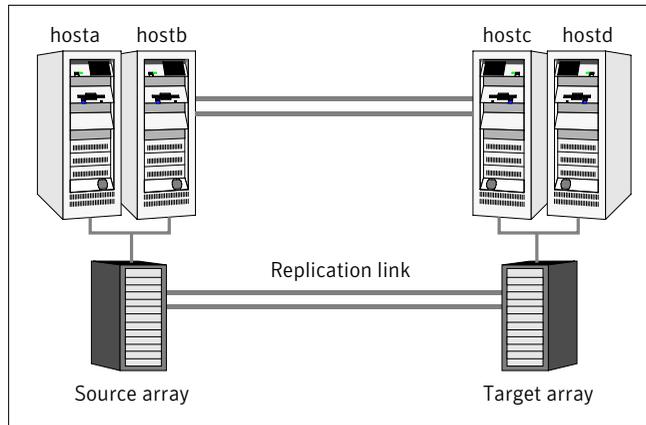
In a replicated data cluster, you can prevent unnecessary MirrorView failover or failback by creating system zones. VCS attempts to fail over applications within the same system zone before failing them over across system zones.

Configure the hosts that are attached to an array as part of the same system zone to avoid unnecessary failover.

Figure 3-2 depicts a sample configuration where hosta and hostb are in one system zone, and hostc and hostd are in another system zone.

Use the SystemZones attribute to create these zones.

Figure 3-2 Example system zone configuration



Modify the SystemZones attribute using the following command:

```
hagr -modify grpname SystemZones hosta 0 hostb 0 hostc 1 hostd 1
```

The variable grpname represents the service group in the cluster.

```
# hagr -modify <group_name> SystemZones <system_name> <zone_number>
```

Global clusters do not require system zones because failover occurs on a remote cluster if all local targets have been exhausted.

As long as a secondary image is available, MirrorView sends the writes to the secondary image immediately in synchronous mode. It does so periodically in asynchronous mode.

If the period is too long, you can perform synchronization using the resync action. The supported resync action is defined in the MirrorView resource type.

About preventing split-brain

Split-brain occurs when all heartbeat links between the primary and secondary hosts are cut. In this situation, each side mistakenly assumes that the other side is down. You can minimize the effects of split-brain by ensuring that the cluster heartbeat links pass through a similar physical infrastructure as the replication links. When you ensure that both pass through the same infrastructure, if one breaks, so does the other.

Sometimes you cannot place the heartbeats alongside the replication links. In this situation, a possibility exists that the cluster heartbeats are disabled, but the replication link is not. A failover transitions the original source to target and target to source. In this case, the application faults because its underlying volumes become write-disabled, causing the service group to fault. VCS tries to fail it over to another host, causing the same consequence in the reverse direction. This phenomenon continues until the group comes online on the final node. You can avoid this situation by setting up your infrastructure such that loss of heartbeat links also mean the loss of replication links.

Configuring the agent for EMC MirrorView

You can configure clustered applications in a disaster recovery environment by:

- Synchronizing the mirrors
- Adding the EMC MirrorView agent to the service group

After configuration, the application service group must follow the dependency diagram.

See [“Sample configuration for the MirrorView agent”](#) on page 19.

Note: You must not change the replication state of devices from primary to secondary and from secondary to primary, outside of a VCS setup. The agent for EMC MirrorView fails to detect a change in the replication state if the role reversal is done externally and RoleMonitor is disabled.

Configuring the agent manually in a global cluster

Configuring the agent manually in a global cluster involves the following tasks:

To configure the agent in a global cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 Add a resource of type mirrorview at the bottom of the service group.
- 3 Configure the attributes of the mirrorview resource.
- 4 If the service group is not configured as a global service group, configure the service group using the Global Group Configuration Wizard.

See the *Symantec Cluster Server Administrator's Guide* for more information.

- 5 Change the ClusterFailOverPolicy attribute from the default, if necessary.
Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
- 6 Repeat step 2 through step 5 for each service group in each cluster that uses replicated data.
The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Configuring the agent manually in a replicated data cluster

Configuring the agent manually in a replicated data cluster involves the following tasks:

To configure the agent in a replicated data cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 In each service group that uses replicated data, add a resource of type mirrorview at the bottom of the service group.
- 3 Configure the attributes of the mirrorview resource. Note that some attributes must be localized to reflect values for the hosts that are attached to different arrays.
- 4 Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

Executing the addArrayUser action

Executing the addArrayUser action is a critical step in configuring the EMC MirrorView agent in both global clusters and replicated data clusters.

The MirrorView agent runs under the system user context. You must execute the addArrayUser action to create the required security files, thus allowing you as a system user to run commands on both local and remote arrays.

To execute the addArrayUser action

- 1 Right click the mirrorview resource and select **Actions**.
- 2 From the list of actions, select **addArrayUser**.
- 3 Select the system on which to execute the action.

- 4 From the list of action arguments, click the plus button to add a new argument.

You can specify the credentials for the local and the remote CLARiiON arrays in the arguments as follows:

```
localuser=LocalArrayUserName  
localpass=LocalArrayPassword  
remoteuser=RemoteArrayUserName  
remotepass=RemoteArrayPassword
```

You can use this action to specify the credentials of the local array or remote array or both. You can also use the following two keys to specify the path to the security files for the arrays:

```
localsecfilepath=LocalArraySecurityFilePath  
remotesecfilepath=RemoteArraySecurityFilePath
```

If the security file path option is specified, this action also updates the appropriate resource attributes (*LocalArraySecurityFilePath* and/or *RemoteArraySecurityFilePath*) so that the subsequent `naviseccli` commands that the agent issues may succeed.

Note: Use the security file path option only if you do not want the security files to be created in the default location.

- 5 Click **OK** to execute the action.
- 6 Review the results of the action and click **OK**.
- 7 Repeat these steps for all nodes that need to access the arrays.

Managing and testing clustering support for EMC MirrorView

This chapter includes the following topics:

- [Typical test setup for the EMC MirrorView agent](#)
- [Testing service group migration](#)
- [Testing host failure](#)
- [Performing a disaster test](#)
- [Performing the failback test](#)
- [Failure scenarios for EMC MirrorView](#)

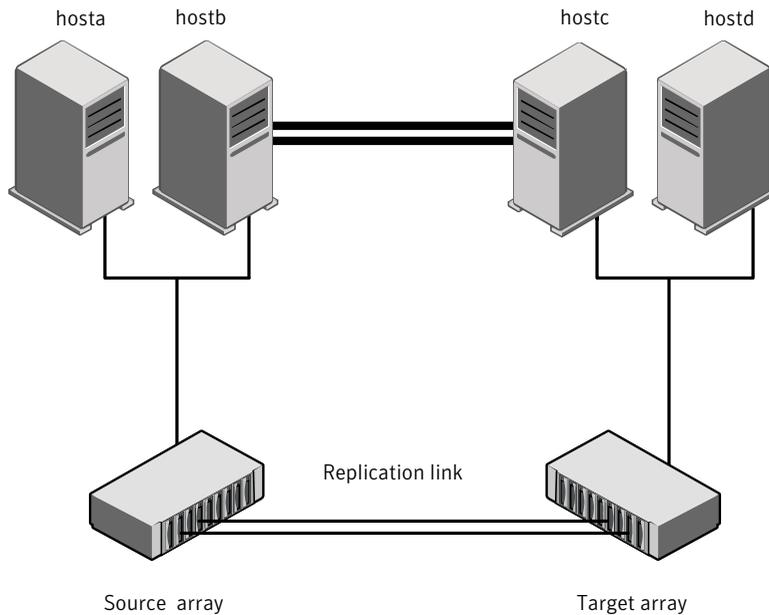
Typical test setup for the EMC MirrorView agent

A typical test environment includes the following characteristics:

- Two hosts (hosta and hostb) are attached to the source CLARiiON array.
- Two hosts (hostc and hostd) are attached to the target CLARiiON array.
- The application runs on hosta and devices in the local array are read-write enabled in the SYNCHRONIZED state.
- A replicated data cluster has two dedicated heartbeat links.
A global cluster has one network heartbeat.

[Figure 4-1](#) depicts a typical test environment.

Figure 4-1 Typical test setup



Testing service group migration

Verify that the service group can migrate to different hosts in the cluster and across clusters.

To perform the service group migration test

- 1 In the Cluster Explorer configuration tree, under the **Service Groups** tab, right-click the service group.
 Migrate the service group to a host that is attached to the same array.
- 2 Click **Switch To** and click the system that is attached to the same array (hostb) from the menu.
 The service group comes online on hostb and local image remains in the MIRRORRED state.
- 3 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
 Migrate the service group to a host that is attached to a different array.

- 4 Click **Remote Switch**, and click the remote cluster and system that are attached to the same array (hostb) from the menu.
 The service group comes online on hostc and the role of the images there transition to primary.
- 5 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
 After the devices transition to a source SYNCHRONIZED state, migrate the service group back to its original host.
- 6 Click **Remote Switch**, and click the remote cluster and system on which the group was initially online (hosta).
 The group comes online on hosta. The devices return to the MIRRORED state at the array that is attached to hosta and hostb, and then eventually transition to the SYNCHRONIZED state.

Testing host failure

In this scenario, the host where the application runs is lost. Eventually, all the hosts in the system zone or cluster are lost.

To perform the host failure test

- 1 Halt or shut down the host where the application runs (hosta).
 The service group fails over to hostb.
- 2 Halt or shut down hostb.
 In a replicated data cluster, the group fails over to hostc or hostd depending on the FailOverPolicy attribute in the cluster.
 In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.
 In both environments, the role of the devices changes from secondary to primary and starts on the target host.
- 3 Power on the two hosts that were shut down.
- 4 Switch the service group to its original host when VCS starts.
 Do the following:
 - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
 - Click **Switch To** and click the system on which the service group was initially online (hosta).

The service group comes online on hosta.

Performing a disaster test

Test how robust your cluster is in case of a disaster.

To perform a disaster test

- 1 Shut down all hosts on the source side and shut down the source array.
 If you cannot shut down the source array, change the value of the RemoteArraySPNames in the target side Mirrorview resource to non-existent names and IP addresses. This action mimics a disaster scenario from the target's point of view.
- 2 In a replicated data cluster, the service group fails over to hostc or hostd in the following conditions:
 - All devices were originally in the SYNCHRONIZED state.
 - No synchronization was in progress at the time of disaster.
- 3 In a global cluster, the administrator is notified of the failure. The administrator can then initiate the failover.

Performing the failback test

You can set up your cluster for a failback test.

The failback test verifies the application can fail back to its original host after a failover to a remote site.

To perform the failback test with Consistency groups

- 1 Remove all the mirrors from the consistency group on the old primary.
- 2 Destroy the consistency group on the old primary.
- 3 Forcefully destroy the remote mirrors on the old primary.
- 4 Remove the LUNs from the storage group on the old primary.
- 5 Remove the mirrors from the consistency group on the new primary.
- 6 Add secondary images to each of the remote mirrors on the new primary.

- 7 Add the mirrors into the consistency group on the new primary.

Between step 5 and step 7, the LUNs become vulnerable to data corruption. For example, if one of the LUNs has sustained hardware damage and failed.

During this window, the mirrors are not a part of the consistency group. The writes to other mirrors that were a part of the consistency group are not stopped. This situation could result in data corruption.

- 8 Add the LUNs, where the secondary image resides, into the appropriate storage group on the old primary.

To perform the failback test for synchronous and asynchronous mode with Individual mirrors

- 1 Forcefully destroy the remote mirrors on the old primary.
- 2 Remove the LUNs from the storage group on the old primary.
- 3 Add secondary images to each of the remote mirrors on the new primary.
- 4 Add the LUNs, where the secondary image resides, into the appropriate storage group on the old primary.

In either of the modes, the original contents of the old primary are lost.

Failure scenarios for EMC MirrorView

Review the failure scenarios and agent behavior in response to failure.

Site disaster

In a total site failure, all hosts and the array are completely disabled, either temporarily or permanently.

In a replicated data cluster, site failure is detected the same way as a total host failure, that is, the loss of all LLT heartbeats.

In a global cluster, VCS detects site failure by the loss of all configured heartbeats.

A total disaster renders the devices on the surviving array in the FRACTURED state. If the SplitTakeover attribute is set to its default value of 1, the online entry point runs the 'promote' operation. If the attribute is set to 0, no takeover occurs and the online entry point times out and faults.

The online entry point detects whether any synchronization was in progress when the source array was lost. Since the target devices are inconsistent until the synchronization completes, the agent does not write-enable the devices, but it times out and faults. You must restore consistent data from a snapshot or tape backup.

All host or all application failure

Even if both arrays are operational, the service group fails over in the following conditions:

- All hosts on the source CLARiiON side are disabled.
- The application cannot start successfully on any source host.

In replicated data cluster environments, the failover can be automatic, whereas in global cluster environments failover requires user confirmation by default.

In replicated data cluster environments, site failure is detected the same way as a total host failure, that is, the loss of all LLT heartbeats. This type of failure is communicated by the VCS engine to the other site.

In both replicated data cluster and global cluster environments, multiple service groups can fail over in parallel.

Replication link failure

Before the MirrorView takes any action, it waits for the synchronization to complete in the following situations:

- The two arrays are healthy and the link that failed is restored.
- A failover is initiated while synchronization is in progress.

After the synchronization completes, the MirrorView runs the promote operation.

If the agent times out before the synchronization completes, the resource faults.

If the SplitTakeover attribute is set to 0, the agent does not attempt a promote operation, but it times out and faults. If you write-enable the devices manually, the agent can come online after it is cleared.

Split-brain in a MirrorView environment

You must resynchronize the volumes manually by using the following command:

```
naviseccli -h array_ip mirror -async/-sync -syncimage/-syncgroup.
```

In a global cluster, you can confirm the failure before failing over the service groups. You can check with the site administrator to identify the cause of the failure. If a fail over mistakenly occurs, the situation is similar to the replicated data cluster case. However, when the heartbeat is restored, VCS does not stop HAD at either site. VCS forces you to choose which group to take offline. You must resynchronize the data manually.

Setting up fire drill

This chapter includes the following topics:

- [About fire drills](#)
- [About the MirrorViewSnap agent](#)
- [Additional considerations for running a fire drill](#)
- [Before you configure the fire drill service group](#)
- [Configuring the fire drill service group](#)
- [Verifying a successful fire drill](#)

About fire drills

A fire drill procedure verifies the fault-readiness of a disaster recovery configuration. This procedure is done without stopping the application at the primary site and disrupting user access.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group is identical to the application service group, but uses a fire drill resource in place of the replication agent resource. The fire drill service group uses a copy of the data that is used by the application service group.

In clusters employing EMC MirrorView, the MirrorViewSnap resource manages the replication relationship during a fire drill.

Bringing the fire drill service group online demonstrates the ability of the application service group to come online at the remote site when a failover occurs.

The MirrorViewSnap agent supports fire drills for storage devices that are managed using Veritas Volume Manager, which is a component of Symantec Storage Foundation.

About the MirrorViewSnap agent

The MirrorViewSnap agent is the fire drill agent for EMC MirrorView. The agent maintains the replication relationship between the source and target arrays when running a fire drill. Configure the MirrorViewSnap resource in the fire drill service group, in place of the mirrorview resource.

MirrorViewSnap agent functions

The MirrorViewSnap agent performs the following functions:

online	<p>Takes a snapshot of the LUNs in the mirror/consistency group and make it available for use.</p> <p>It performs the following steps:</p> <ul style="list-style-type: none">Retrieves the state of the mirror/consistency group. If the mirror/consistency group is not in either SYNCHRONIZED/CONSISTENT state, the fire drill cannot continue. The agent logs an error and exits.Creates the Snapview session.Takes the snapshot. In case of any error, the agent rolls back the changes for all the other mirrors. After the snapshot is successfully taken, the agent adds the snapshot LUN(s) to the specified storage group and prepares the VxVM DiskGroup for import..Creates a lock file and exits.
offline	<p>Deletes the snapshot and stops the session.. If successful, the Offline entry point removes the lock file.</p>
monitor	<p>Verifies the existence of the lock file to make sure the resource is online.</p>
clean	<p>Deletes the snapshot and stops the session.. If successful, the Clean entry point removes the lock file.</p>
open	<p>If the lock file does not exist, the agent takes no action.</p> <p>If the lock file exists:</p> <ul style="list-style-type: none">If any resources that depend on the MirrorViewSnap are online, the agent does not take any action.If any resources that depend on the MirrorViewSnap are not online, the agent removes the lock file.

Resource type definition for the MirrorViewSnap agent

Following is the resource type definition for the MirrorViewSnap agent:

```
type MirrorViewSnap (
    static int ActionTimeout = 300
    static keylist SupportedActions = { devscan }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static int OfflineMonitorInterval = 0
    static int OnlineTimeout = 600
    static int OpenTimeout = 180
    static int RestartLimit = 1
    static str ArgList[] = { StorageGrpName, TargetResName }
    str TargetResName
    str StorageGrpName
    temp str Responsibility
    temp str VCSResLock
)
```

Attribute definitions for the MirrorViewSnap agent

To customize the behavior of the MirrorViewSnap agent, configure the following attributes:

TargetResName

Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the mirrorview resource if you want to take a snapshot of replicated data. Set this attribute to the name of the DiskGroup resource if the data is not replicated.

For example, in a typical database setup, you might replicate data files and redo logs, but you may choose to avoid replicating temporary tablespaces. The temporary tablespace must still exist at the DR site and may be part of its own disk group.

Type-dimension: string-scalar

StorageGrpName

Name of the storage group that contains the snapshot. The host that runs the fire drill must be a part of this storage group. Otherwise, the fire drill fails.

Type-dimension: string-scalar

Responsibility	Do not modify. For internal use only. Used by the agent to keep track of resynchronizing snapshots. Type-Dimension: temporary string
----------------	--------------------------------------------------------------------------------------------------------------------------------------------

Sample configuration for a fire drill service group

The sample configuration of a fire drill service group is identical to an application service group with a hardware replication resource. However, in a fire drill service group, the MirrorViewSnap resource replaces the mirrorview resource.

You can configure a resource of type MirrorViewSnap in the main.cf file as follows:

```
group asyncSG_fd (
  SystemList = { VCSLX303 = 0, VCSLX319 = 1 }
)

FileOnOff fof_fd (
  PathName = "f:\\file"
)

MirrorViewSnap mvsnap (
  TargetResName = mv_async
  StorageGrpName = vcslx303_vcslx319
)

MountV mount_fd (
  MountPath = "f:"
  VolumeName = datavol
  VMDGResName = asyncDG_fd
)

VMDg asyncDG_fd (
  DiskGroupName = asyncDG_fd
)

requires group asyncSG offline local
fof_fd requires mount_fd
mount_fd requires asyncDG_fd
asyncDG_fd requires mvsnap
```

Additional considerations for running a fire drill

Follow these guidelines for fire drills in a Windows environment:

- The primary and secondary sites must be fully configured with MirrorView replication and the global cluster option. The configuration must follow the applicable instructions in the for Windows documentation for configuring disaster recovery with MirrorView.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.

Before you configure the fire drill service group

Before you configure the fire drill service group, follow the steps below:

- Make sure the application service group is configured with a mirrorview resource.
- Make sure the infrastructure to take snapshots is properly configured between the source and target arrays.
- Make sure the MirrorView relationship is established.
- Reserve sufficient unallocated LUNs in the reserved LUN pool.
- Install and enable the SnapView license.
- Install the Navisphere CLI

Configuring the fire drill service group

On the secondary site, the initial steps create a fire drill service group that closely follows the configuration of the original application service group. The fire drill service group uses a point-in-time copy of the production data. Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over and come online at the secondary site, should the need arise.

See [“Sample configuration for a fire drill service group”](#) on page 35.

You can create the fire drill service group using the Cluster Manager (Java Console).

Creating the fire drill service group using Cluster Manager (Java Console)

This section describes how to use Cluster Manager (Java Console) to create the fire drill service group.

After creating the fire drill service group, you must set the failover attribute to false so that the fire drill service group does not fail over to another node during a test.

To create the fire drill service group

- 1 Open the Cluster Manager (Java Console).
- 2 Log on to the cluster and click **OK**.
- 3 Click the **Service Group** tab in the left pane and click the **Resources** tab in the right pane.
- 4 Right-click the cluster in the left pane and click **Add Service Group**.
- 5 In the Add Service Group dialog box, provide information about the new service group.
 - In Service Group name, enter a name for the fire drill service group.
 - Select systems from the Available Systems box and click the arrows to add them to the Systems for Service Group box.
 - Click **OK**.

To disable the AutoFailOver attribute

- 1 Click the **Service Group** tab in the left pane and select the fire drill service group.
- 2 Click the **Properties** tab in the right pane.
- 3 Click the **Show all Attributes** button.
- 4 Double-click the **AutoFailOver** attribute.
- 5 In the Edit Attribute dialog box, clear the AutoFailOver check box and click **OK**.
- 6 Click the Save and Close Configuration icon in the tool bar.

Adding resources to the fire drill service group

Add resources to the new fire drill service group to recreate key aspects of the application service group.

To add resources to the service group

- 1 In Cluster Explorer, click the **Service Group** tab in the left pane, click the application service group and click the **Resources** tab in the right pane.
- 2 Right-click the resource at the top of the tree, select **Copy > Self and Child Nodes**.
- 3 In the left pane, click the fire drill service group.
- 4 Right-click the right pane, and click **Paste**.

- 5 In the Name Clashes dialog box, specify a way for the resource names to be modified, for example, insert an FD_ prefix. Click **Apply**.
- 6 Click **OK**.

Configuring resources for the fire drill service group

Edit the resources in the fire drill service group so that they work properly with the duplicated data. The attributes must be modified to reflect the configuration at the remote site. Bringing the service group online without modifying resource attributes is likely to result in a cluster fault and interruption in service.

To configure the fire drill service group

- 1 In Cluster Explorer, click the **Service Group** tab in the left pane.
- 2 Click the fire drill service group in the left pane and click the **Resources** tab in the right pane.
- 3 Right-click the Mirrorview resource and click **Delete**.
- 4 Add a resource of type MirrorViewSnap and configure its attributes.
- 5 Right-click the resource to be edited and click **View > Properties View**. If a resource to be edited does not appear in the pane, click **Show All Attributes**.
- 6 Edit attributes to reflect the configuration at the remote site. For example, change the Mount resources so that they point to the volumes that are used in the fire drill service group

Verifying a successful fire drill

Run the fire drill routine periodically to verify the application service group can fail over to the remote node.

To verify a successful fire drill

- 1** Bring the fire drill service group online on a node at the secondary site that does not have the application running.

If the fire drill service group comes online, it action validates your disaster recovery configuration. The production service group can fail over to the secondary site in the event of an actual failure (disaster) at the primary site.
- 2** If the fire drill service group does not come online, review the VCS engine log for more information.
- 3** Take the fire drill offline after its functioning has been validated.

Failing to take the fire drill offline could cause failures in your environment. For example, if the application service group fails over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Index

A

- agent functions
 - addArrayUser 12
 - clean 12
 - info 12
 - monitor 12
 - offline 12
 - online 12
 - open 12
- application failure 31
- attribute definitions 18

C

- cluster
 - heartbeats 21
- configuring
 - before 17
 - samples 19

D

- disaster test 29

F

- failback test 29
- failure scenarios
 - all application failure 31
 - all host failure 31
 - replication link failure 31
 - total site disaster 30
- fire drill
 - about 32
 - configuration wizard 36
 - MirrorViewSnap agent 33
 - running 38
 - service group for 36

H

- host failure 31

I

- installing the agent
 - Windows systems 14

M

- migrating service group 27
- MirrorViewSnap agent
 - about 33
 - attribute definitions 34
 - operations 33
 - type definition 33
- MirrorViewSnap agent attributes 34
 - Responsibility 35

O

- Online function
 - more information 12
- optional attributes
 - LocalArraySecurityFilePath 19
 - RemoteArraySecurityFilePath 19

R

- replication link failure 31
- required attributes
 - GrpName 18
 - LocalArraySPNames 18
 - MirNames 18
 - Mode 18
 - NaviCliHome 18
 - RemoteArraySPNames 18
 - SplitTakeover 18
- resource type definition
 - MirrorViewSnap agent 33
- Responsibility attribute 35

S

- sample configuration 19
- service group
 - migrating 27

- split-brain
 - handling in cluster 22
 - handling in clusters 31
- StorageGrpName attribute 34

T

- testing
 - disaster 29
 - failback 29
- testing failback
 - asynchronous with individual mirrors 30
 - synchronous with individual mirrors 30
- total site disaster 30
- type definition
 - MirrorViewSnap agent 33

U

- uninstalling the agent
 - Windows systems 15