

# Veritas™ Cluster Server Hardware Replication Agent for IBM Metro Mirror Configuration Guide

Windows 2000, Windows Server 2003

5.0

# Veritas Cluster Server Agent for IBM Metro Mirror Metro Mirror Agent Guide

Copyright © 2007 Symantec Corporation. All rights reserved.

Veritas Cluster Server 5.0

Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation 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.

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.

THIS 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” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation  
20330 Stevens Creek Blvd.  
Cupertino, CA 95014  
[www.symantec.com](http://www.symantec.com)

## Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Windows is a registered trademark of Microsoft Corporation.

## Licensing and registration

Veritas Cluster Server is a licensed product.

## Technical support

For technical assistance, visit <http://entsupport.symantec.com> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.



# Contents

Chapter 1	Introducing the Veritas agent for IBM Metro Mirror	
	About the agent for IBM Metro Mirror .....	8
	Supported software and hardware .....	8
	Typical IBM Metro Mirror setup in a VCS cluster .....	9
	IBM Metro Mirror agent functions .....	10
Chapter 2	Installing and removing the agent	
	Before you install the agent for Metro Mirror .....	12
	Installing the agent for Metro Mirror .....	12
	Removing the agent for Metro Mirror .....	13
Chapter 3	Configuring the agent for IBM Metro Mirror	
	Configuration concepts .....	16
	MetroMirror resource type definition .....	16
	Attribute definitions for the Metro Mirror agent .....	17
	Sample configuration .....	19
	Before you configure the agent for Metro Mirror .....	20
	About cluster heartbeats .....	20
	About configuring system zones in replicated data clusters .....	21
	Configuring the agent for Metro Mirror .....	22
	Configuring the agent manually in a global cluster .....	22
	Configuring the agent manually in a replicated data cluster .....	23
Chapter 4	Managing and testing clustering support for IBM Metro Mirror	
	Typical test setup .....	26
	Testing service group migration .....	27
	Testing host failure .....	27
	Performing a disaster test .....	28
	Performing the failback test .....	28
Index		31



# Introducing the Veritas agent for IBM Metro Mirror

This chapter contains the following topics:

- [About the agent for IBM Metro Mirror](#)
- [Supported software and hardware](#)
- [Typical IBM Metro Mirror setup in a VCS cluster](#)
- [IBM Metro Mirror agent functions](#)

## About the agent for IBM Metro Mirror

The Veritas agent for IBM Metro Mirror provides failover support and recovery in environments employing Metro Mirror to replicate data between IBM DS8000 arrays.

The Veritas agent for IBM Metro Mirror monitors and manages the state of replicated DS6000 and DS8000 volumes that are attached to VCS nodes. The agent ensures that the system on which the Metro Mirror resource is online has safe exclusive access to the configured volumes.

The agent can be used in single replicated data clusters and multi-cluster environments that are set up using the VCS Global Cluster Option. The agent also supports parallel applications.

The agent supports Metro Mirror (i.e. synchronous replication) only; the agent does not support Global Copy nor Global Mirror (i.e. asynchronous replication).

## Supported software and hardware


The IBM Metro Mirror agent supports VCS 4.3MP1 and later.

The agent supports all versions of IBM DSCLI. The agent supports Metro Mirror on all microcode levels on all DS6000 and DS8000 arrays, provided the host/HBA/array combination is in IBM's hardware compatibility list.



## Typical IBM Metro Mirror setup in a VCS cluster

Clustering in an IBM Metro Mirror environment typically consists of the following hardware infrastructure:



- The primary array, comprising one or more hosts directly attached by SCSI or Fibre Channel to a DS6000 or DS8000 array containing Metro Mirror source volumes.
- The secondary array, comprising one or more secondary hosts directly attached by SCSI or Fibre Channel to a DS6000 or DS8000 array containing Metro Mirror target volumes. The target volumes are paired with the source volumes in the Primary array. The secondary hosts and arrays must be at a significant distance apart from the primary to survive a disaster that may occur at the primary side.
- Network heartbeats between the two data centers to determine their health.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with dual dedicated networks that support LLT.
- In a global cluster environment, you must attach all hosts in a cluster to the same DS6000 or DS8000 array.
- In parallel applications, all hosts that are attached to the same array must be part of the same GAB membership.

## IBM Metro Mirror agent functions

The Veritas agent for IBM Metro Mirror monitors and manages the state of replicated DS6000 or DS8000 devices that are attached to VCS nodes.

The agent performs the following functions:

online	<p>If the state of all local devices is read-write enabled, the agent creates a lock file on the local host. The lock file indicates that the resource is online. This operation makes the devices writable for the application.</p> <p>If all local devices are in the write-disabled state, the agent runs a <code>failoverpprc</code> command to enable read-write access to the devices.</p> <p>For target volumes in the TARGET FULL DUPLEX state, the agent runs the <code>failoverpprc</code> command to make the volumes writable.</p> <p>If the original primary volumes are still accessible, the agent runs the <code>failbackpprc</code> command to reverse the direction of replication.</p>
offline	<p>Removes the lock file from the host. The agent does not run any Metro Mirror commands because taking the resource offline is not indicative of the intention to give up the devices.</p>
monitor	<p>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.</p>
open	<p>Removes the lock file on the host where the entry point is called. This operation prevents potential concurrency violation if the service group fails over to another node.</p> <p><b>Note:</b> The agent does not remove the lock file if the agent was started after running the <code>hastop -force</code> command.</p>
clean	<p>Determines if it is safe to fault the resource if the online entry point fails or times out. The main consideration is whether a management operation was in progress when the online thread timed out and was killed, potentially leaving the devices in an unusable state.</p>
action	<p>Performs a <code>failbackpprc</code> from the original secondary side to merge any changed tracks from the original secondary to the original primary.</p>

# Installing and removing the agent

This chapter contains the following topics:

- [Before you install the agent for Metro Mirror](#)
- [Installing the agent for Metro Mirror](#)
- [Removing the agent for Metro Mirror](#)

## Before you install the agent for Metro Mirror

Set up your cluster. For information about installing and configuring VCS, see the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide*.

Set up replication and the required hardware infrastructure.

See “[Typical IBM Metro Mirror setup in a VCS cluster](#)” on page 9.

## Installing the agent for Metro Mirror

If you did not install the IBM Metro Mirror agent when you installed Veritas Storage Foundation for Windows High Availability (SFW HA), follow these instructions to install the agent.

You must install the agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster. These instructions assume that you have already installed SFW HA.

### To install the agent

- 1 Open the Windows Control Panel and click **Add or Remove Programs**.
- 2 Click the SFW HA Server Components entry and click **Change**.
- 3 On the installer screen, click **Add or Remove** and click **Next**.
- 4 In the Option Selection dialog box, select the agent and click **Next**.
- 5 The installer validates the system for installation.  
If a system is rejected, the Comments column displays the cause of rejection. Highlight the system to view detailed information about the failure in the Details box. Resolve the error, highlight the node in the selected systems list, and click **Validate Again**.  
After all the systems are accepted, click **Next**.
- 6 An informational message appears if you selected the DMP option. Review the information and click **OK** to continue.
- 7 Review the summary of your selections and click **Next**.
- 8 Click **Update** to start the installation.
- 9 The installer displays the status of installation. After the installation is complete, review the installation report and click **Next**.
- 10 Click **Finish**.

## Removing the agent for Metro Mirror

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

- 1 Open the Windows Control Panel and click **Add or Remove Programs**.
- 2 Click the SFW HA Server Components entry and click **Remove**.
- 3 Click **Next**.
- 4 In the Option Selection dialog box, select the Metro Mirror agent and click **Next**.
- 5 The installer validates the system for uninstallation.  
If a system is rejected, the Comments column displays the cause of rejection. Highlight the system to view detailed information about the failure in the Details box. Resolve the error, highlight the node in the selected systems list, and click **Validate Again**.  
After all the systems are accepted, click **Next**.
- 6 Review the summary of your selections and click **Uninstall**.
- 7 The installer displays the status of uninstallation.
- 8 After the uninstallation is complete, review the report and click **Next**.
- 9 Click **Finish**.

---

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

---



# Configuring the agent for IBM Metro Mirror

This chapter contains the following topics:

- [Configuration concepts](#)
- [Before you configure the agent for Metro Mirror](#)
- [Configuring the agent for Metro Mirror](#)

## Configuration concepts

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

### MetroMirror resource type definition

The IBM Metro Mirror agent is represented by the MetroMirror resource type in VCS.

```
type MetroMirror (  
    static keylist SupportedActions = { failback }  
    static int MonitorInterval = 300  
    static int NumThreads = 1  
    static str ArgList[] = { DSCliHome, HMC1, HMC2, User,  
        PasswdFile, LocalStorageImageID, RemoteStorageImageID,  
        VolIds }  
    str DSCliHome = "C:\Program Files\ibm\dscli"  
    str HMC1  
    str HMC2  
    str User  
    str LocalStorageImageID  
    str RemoteStorageImageID  
    str VolIds{}  
    temp str VCSResLock  
)
```



## Attribute definitions for the Metro Mirror agent

Review the description of the agent attributes.

### Required attributes

You must assign values to required attributes.

DSCLIHome	Path to the DS6000 or DS8000 command line interface. Type-dimension: string-scalar Default is C:\Program Files\ibm\dscli.
HMC1	IP address or host name of the primary management console. Type-dimension: string-scalar
User	User name for issuing DS CLI commands from the command line. Type-dimension: string-scalar
PasswdFile	Specifies the password file that contains your password. See the <code>managepasswd</code> DS CLI command for information on how to generate a password file. Type-dimension: string-scalar
LocalStorageImageID	The image ID of the local storage, which consists of manufacturer, type, and serial number. For example, IBM.2107-75FA120 Type-dimension: string-scalar
RemoteStorageImageID	The image ID of the local storage, which consists of manufacturer, type, and serial number. For example, IBM.3108-75GB248 Type-dimension: string-scalar
VolIds	IDs of local DS6000 or DS8000 Metro Mirror volumes that the agent manages. Type-dimension: string-keylist

### Optional attributes

Configuring these attributes is optional.

HMC2	IP address or host name of the secondary management console. Type-dimension: string-scalar
------	---

## Internal attributes

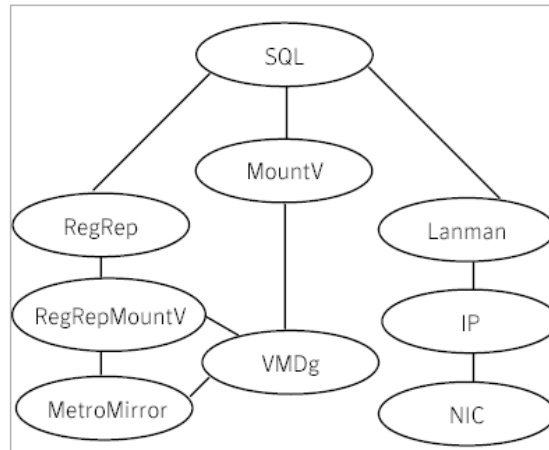
These attributes are for internal use only. Do not modify their values.

VCSResLock    The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application.  
Type-dimension: temporary string-scalar

## Sample configuration

Figure 3-1 shows the dependency graph for a VCS service group with a resource of type MetroMirror. The DiskGroup resource depends on the Metro Mirror resource.

**Figure 3-1** Sample configuration for the Metro Mirror agent



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

```

MetroMirror mir (
  DSCLiHome = "C:\Program Files\ibm\dsccli"
  HMC1 = "ds8000c.example.com"
  User = admin
  PasswdFile = "C:\Program Files\ibm\dsccli\ds_pwfile"
  LocalStorageImageID = "IBM.2107-75FA120"
  RemoteStorageImageID = "IBM.2107-75FA150"
  VolIds = { 1260, 1261 }
)
  
```

This resource manages the following objects:

- A group of two Metro Mirror volumes: 1260 and 1261 on the local array with the storage image ID IBM.2107-75FA120.
- The local array is managed by the HMC ds8000c.example.com.
- The Metro Mirror target volumes are on the remote array with the storage image ID IBM.2107-75FA150.
- The password file, created using the managepwfile DSCLI command, is located at the path C:\Program Files\ibm\dsccli\ds\_pwfile.

## Before you configure the agent for Metro Mirror

Before configuring the agent, review the following information:

- Review the configuration concepts, which describe the agent's type definition and attributes.  
See "[Configuration concepts](#)" on page 16.
- Verify that the agent is installed on all systems in the cluster.
- Verify the hardware setup for the agent.  
See "[Typical IBM Metro Mirror setup in a VCS cluster](#)" on page 9.
- Ensure that Metro Mirror paths are configured in both directions between the source and the target LSS. Metro mirror role reversal fails if paths are not configured from the current target LSS to the current source LSS.
- Make sure the cluster has an effective heartbeat mechanism in place.  
See "[About cluster heartbeats](#)" on page 20.
- Set up system zones in replicated data clusters.  
See "[About configuring system zones in replicated data clusters](#)" on page 21.
- Generate the DSCLI password file. Use the `managepwfile` DSCLI command to do so.
- Reboot the node after the DSCLI software is installed on that node. The DSCLI installation sets some system environment variables that don't take effect until after a reboot. If these environment variables are not set, the Metro Mirror agent will not function properly.

### About cluster heartbeats

In a replicated data cluster, robust heartbeating is accomplished through 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.

## About configuring system zones in replicated data clusters

In a replicated data cluster, you can prevent unnecessary failover or failback. 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



## Configuring the agent for Metro Mirror

Most clustered applications can be adapted to a disaster recovery environment by:

- Converting their devices to Metro Mirror devices
- Synchronizing the devices
- Adding the IBM Metro Mirror agent to the service group

DS6000 and DS8000 volumes are configured as resources of type MetroMirror.

## 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 If the agent resource type (MetroMirror) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select `C:\Program Files\Veritas\Cluster Server\conf\config\MetroMirrorTypes.cf`.
- 3 Click **Import**.
- 4 Save the configuration.
- 5 Add a resource of type MetroMirror at the bottom of the service group.
- 6 Configure the attributes of the MetroMirror resource.
- 7 If the service group is not configured as a global group, configure the service group using the Global Group Configuration Wizard. See the *Veritas Cluster Server Administrator's Guide* for more information.
- 8 Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.  
Repeat [step 5](#) through [step 8](#) for each service group in each cluster that uses replicated data.

## 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 If the agent resource type (MetroMirror) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select `C:\Program Files\Veritas\Cluster Server\conf\config\MetroMirrorTypes.cf`.
- 3 Click **Import**.
- 4 Save the configuration.
- 5 In each service group that uses replicated data, add a resource of type MetroMirror at the bottom of the service group.
- 6 Configure the attributes of the MetroMirror resource. Note that some attributes must be localized to reflect values for hosts that are attached to different arrays.
- 7 Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.





# Managing and testing clustering support for IBM Metro Mirror

This chapter contains the following topics:


- [Typical test setup](#)
- [Testing service group migration](#)
- [Performing a disaster test](#)
- [Performing the failback test](#)

## Typical test setup

A typical test environment includes:

- Two hosts (hosta and hostb) attached to the primary array.
- Two hosts are attached to the secondary array.
- The application is running on hosta and volumes in the local array are read-write enabled, in the FULL DUPLEX state.

The following depicts a typical test environment.



## Testing service group migration

Verify the service group can migrate to different hosts in the cluster.

### To perform the service group migration test

- 1 Migrate the service group to a host that is attached to the same array. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 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 volumes remain in the FULL DUPLEX state.
- 3 Migrate the service group to a host that is attached to a different array. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 4 Click **Switch To**, and click the system that is attached to the another array (hostc) from the menu.  
The service group comes online on hostc and the volumes there transition to the FULL DUPLEX state from the TARGET FULL DUPLEX state.
- 5 Click **Switch To**, and click the system on which the group was initially online (hosta).  
The group comes online on hosta. The devices return to the original state in step 1.

## 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 and devices are in the FULL DUPLEX state.
- 2 Halt or shut down hostb.  
In a replicated data cluster, the group fails over to hostc or hostd depending on the FailOverPolicy 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 devices transition from the TARGET FULL DUPLEX to the FULL DUPLEX state and start on the target host.

- 3 Reboot the two hosts that were shut down.
- 4 Switch the service group to its original host when VCS starts. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 5 Click **Switch To**, and click the system on which the service group was initially online (hosta).  
The service group comes online on hosta and devices swap roles again.

## 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 primary DS8000, disconnect the Metro Mirror paths and simultaneously shut down the hosts. This action mimics a disaster scenario from the point of view of the secondary side.
- 2 In a replicated data cluster, the service group fails over to hostc or hostd if:
  - All volumes were originally in the TARGET FULL DUPLEX state.
  - No copy or 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.
- 4 After the failover, the original target volumes go to the SUSPENDED state (Reason = "Host Source").

## Performing the failback test

You can set up your cluster for a failback test.

### To perform a failback test

- 1 Reconnect the replication link and reboot the original primary hosts.
- 2 Take the service group offline. If you are running this test in a replicated data cluster, run the following command from any host:

```
hagrps -offline grpname -any
```

If you are running the test in a global cluster, run the command from hostc or hostd.

- 3 Manually resynchronize the volumes using the failback action. After the resynchronization completes, the state of the original target volumes changes to FULL DUPLEX (Reason = "-"). The state of the original source volumes changes to TARGET FULL DUPLEX (Reason = "-").
- 4 Migrate the application back to the original primary side.



# Index

## A

- action entry point 10
- agent
  - configuring in a global cluster 22
  - configuring manually in global cluster 22
- Attribute definitions
  - Metro Mirror agent 17
- attribute definitions 17

## C

- clean entry point 10
- cluster heartbeats
  - about 20
- configuring system zones 21

## D

- disaster test 28
- disaster test, performing 28
- DS8000 10
- DSCliHome attribute 17

## E

- entry points
  - action 10
  - clean 10
  - monitor 10
  - offline 10
  - online 10
  - open 10

## F

- failback test 28
- failback test, performing 28

## G

- global cluster configuration 22

## H

- HMC1 attribute 17
- HMC2 attribute 17
- host failure
  - testing 27

## I

- IBM Metro Mirror agent
  - about 8
  - attribute definitions 17
  - before configuring 20
  - before installing 12
  - configuration concepts 16
  - functions 10
  - installing 12
  - removing 13
  - testing 25
  - type definition 16
  - typical setup 9
- IBM Metro Mirror agent attributes
  - DSCliHome 17
  - HMC1 17
  - HMC2 17
  - LocalStorageImageID 17
  - PasswdFile 17
  - RemoteStorageImageID 17
  - User 17
  - VCSResLock 18
  - VolIds 17

## L

- LocalStorageImageID attribute 17

## M

- migrating service group 27
- monitor entry point 10

## O

- offline entry point 10

online entry point 10  
open entry point 10

## **P**

PasswdFile attribute 17

## **R**

RemoteStorageImageID attribute 17  
resource type definition 16  
    IBM Metro Mirror agent 16

## **S**

sample configuration 19  
service group, migrating 27  
supported hardware 8  
supported software 8

## **T**

testing  
    disaster 28  
    failback 28  
    host failure 27  
type definition  
    IBM Metro Mirror agent 16

## **U**

User attribute 17

## **V**

VCSResLock attribute 18  
VolIds attribute 17