Cluster Server Agent for EMC SRDF Storage Group Installation and Configuration Guide

AIX, Linux

7.0



Veritas InfoScale[™] Availability Agents

Last updated: 2018-08-01

Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas 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. Refer to the third-party legal notices document accompanying this Veritas product or available at:

https://www.veritas.com/about/legal/license-agreements

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 Veritas Technologies LLC 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. VERITAS TECHNOLOGIES LLC 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, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. 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.

Veritas Technologies LLC 500 E Middlefield Road Mountain View, CA 94043

http://www.veritas.com

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

https://www.veritas.com/support

You can manage your Veritas account information at the following URL:

https://my.veritas.com

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

https://sort.veritas.com/documents

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

xyz@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

http://www.veritas.com/community/

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introducing the agent for EMC SRDF Storage Group	6
	About the agent for EMC SRDF Storage Group Supported software Typical EMC SRDF Storage Group setup in a VCS cluster SRDF Storage Group agent functions About the SRDF Storage Group agent's online function About dynamic swap support for the SRDF Storage Group agent	6 7 7 8 9 . 10
Chapter 2	Installing and removing the agent for EMC SRDF Storage Group	11
	Before you install the agent for EMC SRDF Storage Group Installing the agent for EMC SRDF Storage Group Configuring LVM on AIX Upgrading the agent for EMC SRDF Storage Group Removing the agent for EMC SRDF Storage Group	11 11 12 13 14
Chapter 3	Configuring the agent for EMC SRDF Storage Group	15
	Configuration concepts for the EMC SRDF Storage Group agent Resource type definition for the EMC SRDF StorageGroup agent	15
	Attribute definitions for the SRDF StorageGroup agent Sample configuration for the SRDF Storage Group agent Additional configuration considerations for the SRDF Storage Group agent	16 16 19 20
	Before you configure the agent for EMC SRDF Storage Group About cluster heartbeats About preventing split-brain	20 20 21
	Configuring the agent for EMC SRDF Storage Group Configuring the agent manually in a global cluster	21 21

	Configuring the agent in an SF for Oracle RAC environment or Storage Foundation Cluster File System (SFCFS) environment	22
	Setting the OnlineTimeout attribute for the SRDF Storage Group resource	22
Chapter 4	Managing and testing clustering support for EMC SRDF Storage Group	24
	How VCS recovers from various disasters in an HA/DR setup with	
	EMC SRDF Storage Group	24
	Failure scenarios in global clusters	25
	Testing the global service group migration	27
	Testing disaster recovery after host failure	28
	Testing disaster recovery after site failure	29
	Performing failback after a node failure or an application failure	29
	Performing failback after a site failure	30

Chapter

Introducing the agent for EMC SRDF Storage Group

This chapter includes the following topics:

- About the agent for EMC SRDF Storage Group
- Supported software
- Typical EMC SRDF Storage Group setup in a VCS cluster
- SRDF Storage Group agent functions

About the agent for EMC SRDF Storage Group

The Cluster Server agent for EMC SRDF Storage Group (SRDF SG) provides support for application failover and recovery. The agent provides this support in environments that use SRDF storage groups to replicate data between EMC Symmetrix arrays.

The agent monitors and manages the state of replicated EMC Symmetrix devices that are attached to VCS nodes. The agent ensures that the system that has the SRDF SG resource online also has safe and exclusive access to the configured devices.

The agent for EMC SRDF Storage Group supports the following:

- Global clusters that run VCS.
- SRDF storage groups in synchronous mode only. The agent also supports dynamic SRDF (role swap).

The agent supports parallel applications, such as Storage Foundation for Oracle RAC.

Supported software

For information on the software versions that the agent for EMC SRDF Storage Group supports, see the Veritas Services and Operations Readiness Tools (SORT) site: https://sort.veritas.com/agents.

Typical EMC SRDF Storage Group setup in a VCS cluster

The following figure displays a typical cluster setup in a SRDF SG environment.



Figure 1-1 Typical clustering setup for the agent

VCS clusters using SRDF SG for replication uses the following hardware infrastructure:

- The primary array has one or more R1 devices. A Fibre Channel or SCSI directly attaches these devices to the EMC Symmetrix array that contains the SRDF SG R1 devices.
- The secondary array has one or more R2 devices. A Fibre Channel or SCSI directly attaches these devices to a EMC Symmetrix array that contains the SRDF SG R2 devices. The R2 devices are paired with the R1 devices in the R1 array. The R2 devices and arrays must be at a significant distance to survive a disaster that may occur at the R1 side.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.
 See "About cluster heartbeats" on page 20.
- In a global cluster environment, you must attach all hosts in a cluster to the same EMC Symmetrix array.

SRDF Storage Group agent functions

The VCS agent for SRDF SG monitors and manages the state of replicated Symmetrix devices that are attached to VCS nodes. The agent performs the following functions:

Function	Description		
online	This operation makes the devices writable for the application.		
	If the configured devices are in the write-disabled (WD) state, the agent runs the symrdf command to enable read-write access to the devices.		
	See "About the SRDF Storage Group agent's online function" on page 9.		
	If the configured devices are in the not-ready (NR) state, the agent faults the resource.		
	If the state of all local devices in an RDF1 type storage group is read-write enabled (RW) and the replication link is in the Synchronized or the Consistent state, the agent creates a lock file on the local host. The lock file indicates that the resource is online.		
	If the SwapRoles attribute is set, the agent swaps the roles.		
	See "About dynamic swap support for the SRDF Storage Group agent" on page 10.		

Table 1-1Agent functions

Function	Description	
offline	Removes the lock file on the local host. The agent does not run any SRDF commands because taking the resource offline is not indicative of the intention to give up the devices.	
monitor	Verifies that the lock file exists. If the lock file exists, the monitor function reports the status of the resource as online. If the lock file does not exist, the monitor function reports the status of the resource as offline.	
open	See "About dynamic swap support for the SRDF Storage Group agent" on page 10.	
	Removes the lock file on the host where the function is called. This operation 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 running the following command:	
	hastop<-all -local> -force	
clean	Determines if it is safe to fault the resource if the online function fails or times out.	

Table 1-1Agent functions (continued)

About the SRDF Storage Group agent's online function

If the state of all local devices in an RDF1 type device group is read-write enabled (RW) and the replication link is in the Synchronized or the Consistent state, the agent creates a lock file on the local host. The lock file indicates that the resource is online.

If all the local devices are in the write-disabled (WD) state, the agent runs the symrdf command to enable read-write access to the devices.

The states can be different as follows:

- For R2 devices in the SYNCHRONIZED or the CONSISTENT state, the agent runs the symrdf failover command to make the devices writable.
- For R1 devices in the FAILED OVER or R1 UPDATED state, the agent runs the symrdf failback command to make the devices writable.
- For all devices in the PARTITIONED state, if the AutoTakeover attribute is set to 1:
 - If the devices are in RW state, the agent brings the resource online.

 If the devices are in the WD state, the agent runs the symrdf command to bring the resource online.

For all devices in the PARTITIONED state, if the AutoTakeover attribute is set to 0, the agent faults the resource.

- For R1 devices in the UPDINPROG state, the agent runs the symrdf command only after the devices transition to the R1 UPDATED state.
- For R2 devices in the SYNCINPROG state, the agent runs the symrdf command only after the devices transition to the SYNCHRONIZED or the CONSISTENT state.

The agent does not run any command if there is not enough time remaining for the function to complete the command.

- All the devices should either be in the RW state or the WD state. Otherwise, the agents faults the resources.
- If the devices are in the WD state, and the replication link is SUSPENDED, the agent faults the resources.
- If the devices are in the WD state, and the replication link is TRANSIDLE, the online function waits until the timeout. If the replication does not change back to the CONSISTENT state, the agent faults the resource.

See "Setting the OnlineTimeout attribute for the SRDF Storage Group resource" on page 22.

About dynamic swap support for the SRDF Storage Group agent

The agent supports the SRDF/S dynamic swap capability. The agent performs a role swap for the healthy arrays that are configured for dynamic swap when a service group fails over between the arrays. If one array is down, a unilateral read-write enable occurs. The agent fails over the device groups that are not configured for dynamic swap using the following command: symrdf failover. The command enables read-write on the R2 device.

The agent checks the value of the SwapRoles attribute:

- If the attribute is set to 1, the agent performs the role swap.
- If the attribute is set to 0, the agent does not perform the role swap.

For SRDF/S, for R2 devices in the SYNCHRONIZED state, the agent runs the symrdf failover -establish command.

Chapter

Installing and removing the agent for EMC SRDF Storage Group

This chapter includes the following topics:

- Before you install the agent for EMC SRDF Storage Group
- Installing the agent for EMC SRDF Storage Group
- Upgrading the agent for EMC SRDF Storage Group
- Removing the agent for EMC SRDF Storage Group

Before you install the agent for EMC SRDF Storage Group

Before you install the VCS agent for EMC SRDF Storage Group, ensure that you install and configure VCS on all nodes in the cluster.

Set up replication and the required hardware infrastructure.

For information about setting up Oracle RAC environment, refer to the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*.

See "Typical EMC SRDF Storage Group setup in a VCS cluster" on page 7.

Installing the agent for EMC SRDF Storage Group

You must install the EMC SRDF Storage Group 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 VCS or SF for Oracle RAC or Storage Foundation Cluster File System (SFCFS).

To install the agent in a VCS environment

1 Download the Agent Pack from the Veritas Services and Operations Readiness Tools (SORT) site: https://sort.veritas.com/agents.

You can download the complete Agent Pack tar file or the individual agent tar file.

- **2** Uncompress the file to a temporary location, say /tmp.
- **3** If you downloaded the complete Agent Pack tar file, navigate to the directory containing the package for the platform running in your environment.

AIX	<pre>cdl/aix/vcs/replication/srdfsg_agent/ agent_version/pkgs/</pre>
Linux	<pre>cdl/linux/generic/vcs/replication/srdfsg_agent/ agent_version/rpms/</pre>

If you downloaded the individual agent tar file, navigate to the pkgs directory (for AIX, HP-UX, and Solaris), or the rpms directory (for Linux).

- 4 Log in as a superuser.
- **5** Install the package.

AIX # installp -ac -d VRTSvcssg.rte.bff VRTSvcssg.rte
Linux # rpm -ihv \
VRTSvcssg-AgentVersion-Linux GENERIC.noarch.rpm

Note: On successful installation of the agent, if VCS is running, the agent types definition is automatically added to the VCS configuration.

Configuring LVM on AIX

To support failover of the LVM volume groups to the secondary site during a disaster or normal switch, you must have the AIX ODM repository at the secondary populated with the LVM volume group entries. This must be done as part of an initial setup process before VCS starts controlling the replication.

To configure LVM on AIX

- 1 Start the replication. Wait until it is in the synchronized state. Once it is synchronized, split the replication link.
- 2 At the secondary site, run the chdev -l <diskname> -a pv=yes command for each disk inside the replicated device group lvmdg. This gets the physical volume identity (PVID) from within the disk and updates the ODM with this value. Now, these disks have the same PVIDs as their counterparts at the primary site.
- 3 Run the importvg -y <vgname> -n <diskname> command for each volume group.
- 4 Resync the replication and start VCS.

Upgrading the agent for EMC SRDF Storage Group

You must upgrade the agent on each node in the cluster.

To upgrade the agent software

1 Save the VCS configuration and stop the VCS engine.

```
# haconf -dump -makero
```

- # hastop -all -force
- 2 Remove the agent from the node.

See "Removing the agent for EMC SRDF Storage Group" on page 14.

- 3 Delete the file /etc/VRTSvcs/conf/config/SRDFSGTypes.cf.
- 4 Install the current version of the agent.

See "Installing the agent for EMC SRDF Storage Group" on page 11.

- 5 Copy the file SRDFSGTypes.cf from the directory /etc/VRTSvcs/conf/ to the /etc/VRTSvcs/conf/config directory.
- 6 Repeat step 2 through step 5 on each node.
- 7 From a node in the cluster, edit your configuration file /etc/VRTSvcs/conf/config/main.cf.

Configure the new attributes, if applicable.

8 Verify the configuration by running the following command:

hacf -verify config

- 9 Start VCS on local node first.
- **10** Start VCS on other nodes.

Removing the agent for EMC SRDF Storage Group

Before you attempt to remove the agent, make sure the application service group is not online.

You must remove the SRDF SG agent from each node in the cluster.

To remove the agent, type the following command on each node. Answer prompts accordingly:

AIX # installp -u VRTSvcssg.rte Linux # rpm -e VRTSvcssg

Chapter

Configuring the agent for EMC SRDF Storage Group

This chapter includes the following topics:

- Configuration concepts for the EMC SRDF Storage Group agent
- Before you configure the agent for EMC SRDF Storage Group
- Configuring the agent for EMC SRDF Storage Group

Configuration concepts for the EMC SRDF Storage Group agent

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

Resource type definition for the EMC SRDF StorageGroup agent

The SRDFSG resource type represents the EMC SRDF StorageGroup agent in VCS.

```
type SRDFSG (
  static int ActionTimeout = 180
  static int MonitorInterval = 300
  static int OfflineMonitorInterval = 0
  static int OpenTimeout = 180
  static int RestartLimit = 1
  static str ArgList[] = { SymHome, GrpName, DevFOTime, AutoTakeover,
```

```
SplitTakeover, Mode, SwapRoles, LinkMonitor, AdvancedOpts }
str SymHome = "/usr/symcli"
str GrpName
str SymmID
str Rdfg
int DevFOTime = 2
boolean AutoTakeover
boolean SplitTakeover
str Mode
boolean SwapRoles = 1
temp str VCSResLock
int LinkMonitor = 0
str AdvancedOpts{} = { ExtendMonitor = null }
)
```

Attribute definitions for the SRDF StorageGroup agent

Review the description of the agent attributes.

Required attributes

You must assign values to required attributes.

Attribute	Description		
GrpName	Name of the Symmetrix storage group that is failed over as part of the InfoScale service group.		
	Type-dimension: string-scalar		
	Default: ""		
Rdfg	The RDF (RA) group number that is internally used by the Symmetrix storage group, which is failed over as part of the InfoScale service group.		
	Type-dimension: string-scalar		
	Default: ""		
SplitTakeover	A flag that determines whether or not the SRDFSG resource should be brought online when the devices in the storage group are in the Split state. If SplitTakeover is set to 1, the SRDFSG resource is brought ONLINE. If SplitTakeover is set to 0, no action is taken.		
	Type-dimension: boolean-scalar		
	Default: 0		

 Table 3-1
 Required attributes

Attribute	Description
SymmID	The Symmetrix ID for the EMC storage array that is internally used by the Symmetrix storage group, which is failed over as part of the InfoScale service group.
	Type-dimension: string-scalar
	Default: ""
SwapRoles	This attribute only applies to dynamic devices. Specifies whether the roles of the dynamic devices must be swapped at the time of failover or not. If set to 1, the RDF1 dynamic devices are made RDF2, and vice-versa. If set to 0, the roles remain the same.
	Type-dimension: boolean-scalar
	Default: 1

Table 3-1Required attributes (continued)

Optional attributes

Configuring these attributes is optional.

Attribute	Description
SymHome	Path to the bin directory that contains the Symmetrix command line interface.
	Type-dimension: string-scalar
	Default: /usr/symcli.
DevFOTime	Average time in seconds that is required for each device to fail over. This value helps the agent to determine whether it has adequate time for the online operation after waiting for other device to fail over. If the online operation cannot be completed in the remaining time, the failover does not proceed.
	Modify this value if your devices take longer than 2 seconds each to fail over.
	Type-dimension: integer-scalar
	Default: 2 seconds per device

Table 3-2Optional attributes

Attribute	Description		
AutoTakeover	A flag that determines whether the agent performs a symrdf failover operation on the local devices if they are in the partitioned state.		
	See "About the SRDF Storage Group agent's online function" on page 9.		
	For more information, refer to the EMC SRDF documentation.		
	Type-dimension: integer-scalar		
	Default: 0		
LinkMonitor	A flag that determines whether the agent should check the status of the replication link while bringing the resource online.		
	This attribute is used only on the node where the resource is being brought online.		
	This attribute is of importance only at the primary site where the role is R1 and all the devices in the storage group are read-write enabled.		
	The value 1 indicates that the agent will check the status of the replication link. If replication is in the synchronized or the consistent state, then the resource comes online, otherwise, the resource remains offline and results in a service group fault.		
	The value 0 indicates that the agent will not check the status of the replication link while bringing the resource online.		
	Other values of the attribute are reserved for future use by the agent.		
	Type-dimension: integer-scalar		
	Default: 0		
AdvancedOpts	Used at the time of monitoring. This attribute enables the agent to execute custom script during the monitor cycle of the resource.		
	Use the ExtendMonitor key with this attribute. Set the value of ExtendMonitor key as the absolute path of the script that should be executed during the monitor cycle.		
	Set the value of ExtendMonitor to null or remove the key from the AdvancedOpts attribute to disable the execution of the custom script.		
	Type-dimension: string-association		
	Example: AdvancedOpts{} = { ExtendMonitor=null }		

Table 3-2Optional attributes (continued)

Internal attributes

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

Table 3-3 Internal attributes		
Attribute	Description	
VCSResLock	The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application. Type-dimension: temporary string	
Mode	The values of this attribute are reserved for future use by the agent.	

Sample configuration for the SRDF Storage Group agent

The following figure shows the dependency graph for a VCS service group with a resource of type SRDFSG. The DiskGroup resource depends on the SRDFSG resource.

Figure 3-1

Sample configuration for the SRDFSG agent



Note: In this scenario, service groups may be split as long as dependency is set to the service group that has the SRDF Storage Group agent configured.

A resource of type SRDFSG may be configured as follows in main.cf:

```
SRDFSG srdfsg res (
 GrpName = srdf sg
 SymmID = 000292601383
 Rdfg = 21
```

```
LinkMonitor = 1
```

Additional configuration considerations for the SRDF Storage Group agent

Consider the following settings for configuring the SRDF Storage Group agent:

- Set the OnlineTimeout attribute for the SRDF Storage Group resource so that its functions do not time out, or they automatically restart if they timed out.
 See "Setting the OnlineTimeout attribute for the SRDF Storage Group resource" on page 22.
- In global clusters, the value of the AYARetryLimit for the Symm heartbeat must be shorter than the ICMP retry limit. This setting allows VCS to detect an array failure first and does not confuse a site failure with an all-host failure.

Before you configure the agent for EMC SRDF Storage Group

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 SRDF Storage Group setup in a VCS cluster" on page 7.
- Make sure that the cluster has an effective heartbeat mechanism in place.
 See "About cluster heartbeats" on page 20.

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.

In global clusters, the VCS Heartbeat agent sends heartbeats directly between the Symmetrix arrays if the Symmetrix ID of each array is known. This heartbeat offers the following advantages:

- The Symmetrix heartbeat shows that the arrays are alive even if the ICMP heartbeats over the public network are lost. So, VCS does not mistakenly interpret this loss of heartbeats as a site failure.
- 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 preventing split-brain

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 R1 to R2 and R2 to R1. 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 means the loss of replication links.

To minimize the chances of split-brain, use the steward process.

Configuring the agent for EMC SRDF Storage Group

You can adapt most clustered applications to a disaster recovery environment by:

- Converting their devices to SRDF SG devices
- Synchronizing the devices
- Adding the EMC SRDF Storage Group agent to the service group

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

See "Sample configuration for the SRDF Storage Group agent" on page 19.

Configuring the agent manually in a global cluster

The following procedure describes the tasks involved in configuring the agent manually in a global cluster.

To configure the agent in a global cluster

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

Refer to the Cluster Server Administrator's Guide for more information.

- **5** Change the ClusterFailOverPolicy attribute from the default, if necessary. Veritas 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.

Configuring the agent in an SF for Oracle RAC environment or Storage Foundation Cluster File System (SFCFS) environment

To configure the agent to manage the volumes that Storage Foundation for Oracle RAC uses, do the following:

To configure the agent in a Storage Foundation for Oracle RAC or SFCFS environment:

- 1 Configure the SupportedActions attribute for the CVMVoIDg resource.
- **2** Add the following keys to the list: import, deport, and vxdctlenable.
- **3** Run the following commands sequentially to add the entry points to the CVMVoIDg resource:
 - # haconf -makerw
 # hatype -modify CVMVolDg SupportedActions import deport
 vxdctlenable
 - # haconf -dump -makero

Note that SupportedActions is a resource type attribute, and it defines a list of action tokens for the resource.

Setting the OnlineTimeout attribute for the SRDF Storage Group resource

Set the OnlineTimeout attribute for the SRDF Storage Group resource so that its entry points do not time out, or they automatically restart if they timed out.

To set the OnlineTimeout attribute

1 For each SRDF Storage Group resource in the configuration, use the following formula to calculate an appropriate value for the OnlineTimeout attribute:

OnlineTimeout =
$$\sum_{1}^{n} ((n_{devices} \times x_{fallover_time}) + \varepsilon)$$

- n_{devices} represents the number of devices that participate in the failover process of the storage group.
- d_{failovertime} represents the time taken to failover a device.
- n_{devicegroups} represents the total number of storage groups that might fail over simultaneously.
- The epsilon is for the command instantiation overhead. You can set it to any value based on your setup.

To set the Online Timeout attribute for a single storage group (typically the case for SRDF), multiply the number of devices in the storage group with the time taken to failover a device (default = 2 seconds) and add it to the value of epsilon.

For example: if you have a single storage group that consists of 5 devices and the time taken to failover a single device is 50 seconds, set the OnlineTimeout attribute to [(5*50) + 10] seconds. The value of the epsilon here is equal to 10 seconds. Thus, the OnlineTimeout attribute is equal to 260 seconds.

To set the Online Timeout attribute for multiple storage groups (currently not supported by SRDF), calculate the OnlineTimeout attribute for all storage groups and set the OnlineTimeout attribute to at least the amount of time the largest storage group takes to fail over.

2 If the resulting value seems excessive, divide it by two for every increment in the value of the RestartLimit attribute.

Chapter

Managing and testing clustering support for EMC SRDF Storage Group

This chapter includes the following topics:

- How VCS recovers from various disasters in an HA/DR setup with EMC SRDF Storage Group
- Testing the global service group migration
- Testing disaster recovery after host failure
- Testing disaster recovery after site failure
- Performing failback after a node failure or an application failure
- Performing failback after a site failure

How VCS recovers from various disasters in an HA/DR setup with EMC SRDF Storage Group

This topic lists various failure scenarios and describes how VCS responds to the failures in the following DR cluster configurations.

Global clusters

When a site-wide global service group or system fault occurs, VCS failover behavior depends on the value of the ClusterFailOverPolicy attribute for the faulted global service group. The VCS agent for EMC SRDF Storage Group ensures safe and exclusive access to the configured EMC SRDF Storage Group devices.

See "Failure scenarios in global clusters" on page 25.

Refer to the *Cluster Server Administrator's Guide* for more information on the DR configurations and the global service group attributes.

Failure scenarios in global clusters

Table 4-1

The following table lists the failure scenarios in a global cluster configuration and describes the behavior of VCS and the agent in response to the failure.

agent for EMC SRDF Storage Group

Failure scenarios in a global cluster configuration with the VCS

Failure	Description and VCS response			
Application failure	Application cannot start successfully on any hosts at the primary site.			
	VCS response at the secondary site:			
	 Causes global service group at the primary site to fault and displays an alert to indicate the fault. 			
	 Does the following based on the ClusterFailOverPolicy global service group attribute: Auto or Connected—VCS automatically brings the faulted global group online at the secondary site. 			
	 Manual—No action. You must bring the global group online at the secondary site. 			
	See "Performing failback after a node failure or an application failure" on page 29.			
Host failure	All hosts at the primary site fail.			
	VCS response at the secondary site:			
	 Displays an alert to indicate the primary cluster fault. 			
	Does the following based on the ClusterFailOverPolicy global service group attribute:			
	 Auto—VCS automatically brings the faulted global group online at the secondary site. 			
	 Manual or Connected—No action. You must bring the global group online at the secondary site. 			
	See "Performing failback after a node failure or an application failure" on page 29.			

т	able 4-1	Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF Storage Group <i>(continued)</i>	
Failure	Description	Description and VCS response	
Site failure	All hosts and t	he storage at the primary site fail.	
	VCS response	e at the secondary site:	
	 Displays an Does the for Auto—V site. Manual second See "Performing" 	n alert to indicate the cluster fault. bllowing based on the ClusterFailOverPolicy global service group attribute: VCS automatically brings the faulted global group online at the secondary or Connected—No action. You must bring the global group online at the ary site.	
Replication link failure	Poplication link between the arrays at the two sites fails		
	VCS response The DetailedM track informati	e: No action. Monitoring attribute monitors the SRDF link status and reports the dirty fon that the link is in PARTITIONED state.	
Network failure	 The network connectivity and the replication link between the sites fail. VCS response at the secondary site: VCS at each site concludes that the remote cluster has faulted. Does the following based on the ClusterFailOverPolicy global service group Manual or Connected—No action. You must confirm the cause of the refailure from the cluster administrator at the remote site and fix the issue Auto—VCS brings the global group online at the secondary site which to a site-wide split brain. This causes data divergence between the deat the primary and the secondary arrays. When the network (WAC and replication) connectivity is restored, you manually resync the data. Note: Veritas recommends that the value of the ClusterFailOverPolicy is set to Manual for all global groups to prevent unintended failovers data 		
	is set to transier	Manual for all global groups to prevent unintended failovers due to nt network failures.	

1	Table 4-1	Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF Storage Group <i>(continued)</i>	
Failure	Descriptio	Description and VCS response	
Storage failure	The array at VCS respon	The array at the primary site fails. VCS response at the secondary site:	
	 Causes to indicate to indicate to Does the Auto online Manu 	he global service group at the primary site to fault and displays an alert to he fault. following based on the ClusterFailOverPolicy global service group attribute: or Connected—VCS automatically brings the faulted global service group at the secondary site. al—No action. You must bring the global group online at the secondary site.	

Testing the global service group migration

After you configure the Cluster Server agent for EMC SRDF Storage Group, verify that the global service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global service group migration in global cluster setup

- 1 Fail over the global service group from the primary site to the secondary site. Perform the following steps:
 - Switch the global service group from the primary site to any node in the secondary site.

hagrp -switch global group -any -clus cluster name

VCS brings the global service group online on a node at the secondary site.

For static RDF FAILED OVER state

For dynamic RDF The value of the SRDF resource attribute SwapRoles determines the device state:

- 0—FAILED OVER state
- 1—SYNCHRONIZED state
- **2** Fail back the global service group from the secondary site to the primary site. Perform the following steps:
 - Switch the global service group from the secondary site to the primary site.

hagrp -switch global_group -any -clus cluster_name

VCS brings the global service group online at the primary site.

 Verify that the SRDF devices at the primary site are write-enabled and the device state is SYNCHRONIZED.

Testing disaster recovery after host failure

Review the details on host failure and how VCS and the Cluster Server agent for EMC SRDF Storage Group behave in response to the failure.

See "Failure scenarios in global clusters" on page 25.

To test disaster recovery for host failure in global cluster setup

1 Halt the hosts at the primary site.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the VCS failover behavior.

- Auto—VCS brings the faulted global service group online at the secondary site.
- Manual or Connected—You must bring the global service group online at the secondary site.

On a node in the secondary site, run the following command:

hagrp -online -force global_group -any

2 Verify that the global service group is online at the secondary site.

hagrp -state global_group

3 Verify that the SRDF devices at the secondary site are write-enabled and the device state is as follows:

 For static RDF
 FAILED OVER state

 For dynamic RDF
 The value of the SRDF resource attribute SwapRoles determines the device state:

- 0—FAILED OVER state
- 1—SYNCHRONIZED state

Testing disaster recovery after site failure

Review the details on site failure and how VCS and the Cluster Server agent for EMC SRDF Storage Group behave in response to the failure.

See "Failure scenarios in global clusters" on page 25.

To test disaster recovery for site failure in global cluster setup

1 Halt all nodes and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the failover behavior of VCS.

- Auto—VCS brings the faulted global group online at the secondary site.
- Manual or Connected—You must bring the global group online at the secondary site.

On a node in the secondary site, run the following command:

hagrp -online -force global_group -any

- 2 Verify that the SRDF devices at the secondary site are write-enabled and are in PARTITIONED state.
- **3** Verify that the global service group is online at the secondary site.

hagrp -state global_group

Performing failback after a node failure or an application failure

Review the details on node failure and application failure and how VCS and the agent for EMC SRDF Storage Group behave in response to these failures.

See "Failure scenarios in global clusters" on page 25.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site. Perform the procedure that applicable to your DR configuration.

To perform failback after a node failure or an application failure in global cluster

1 Switch the global service group from the secondary site to any node in the primary site.

hagrp -switch global_group -any -clus cluster_name

VCS brings the global service group online at the primary site.

The VCS agent for EMC SRDF Storage Group does the following based on whether the RDF pairs are static or dynamic:

- For dynamic RDF Based on the value of the SwapRoles attribute of the SRDFSG resource:
 - 1—Write enables the devices at the primary site, swaps the R1/R2 personality of each device in the storage group, and restarts replication from R1 devices on the primary site to the R2 devices at the secondary site.
 - 0—Issues the symrdf failback command to resync the R1 devices and to write enable the R1 devices at the primary site.
- For static RDF Issues the symrdf failback command to resync the R1 devices and to write enable the R1 devices at the primary site.
- 2 Verify that the SRDF devices at the primary site are write-enabled and the device state is SYNCHRONIZED.

Performing failback after a site failure

Review the details on site failure and how VCS and the agent for EMC SRDF Storage Group behave in response to the failure.

See "Failure scenarios in global clusters" on page 25.

When the hosts and the storage at the primary site are restarted and the replication link is restored, the SRDF devices attain SPLIT state at both the sites. The devices are write-enabled at both sites. You can now perform a failback of the global service group to the primary site.

To perform failback after a site failure in global cluster

1 Take the global service group offline at the secondary site. On a node at the secondary site, run the following command:

hagrp -offline global_group -any

2 Resync the devices using the symrdf restore command.

The ${\tt symrdf}$ restore command write disables the devices at both the R1 and R2 sites.

After the resync is complete, the device state is SYNCHRONIZED at both the sites.

The devices are write-enabled at the primary site and write-disabled at the secondary site.

3 Bring the global service group online at the primary site. On a node in the primary site, run the following command:

hagrp -online global group -any

This again swaps the role of R1 and R2.