Cluster Server Agent for EMC SRDF Installation and Configuration Guide

AIX, Linux, Solaris

7.0



Cluster Server Agent for EMC SRDF Installation and 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: 7.0

Document version: 7.0 Rev 0

Legal Notice

Copyright © 2016 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. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this product for more information on the Third Party Programs.

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. 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 the company to answer your questions in a timely fashion.

Our 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 our support offerings, you can visit our website at the following URL:

www.veritas.com/support

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.veritas.com/support

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 Technical Support
 - Recent software configuration changes and network changes

Licensing and registration

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

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/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
- Advice about technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Contents

Technical Su	ipport	4
Chapter 1	Introducing the agent for EMC SRDF	10
	About the agent for EMC SRDF	10
	Features of the EMC SRDF agent	11
	Supported software	11
	Supported hardware for EMC SRDF	11
	Typical EMC SRDF setup in a VCS cluster	12
	EMC SRDF agent functions	13
	About the EIVIC SRDF agent's online function	15
	About dynamic swap support for the EMC SRDF agent	10
Chapter 2	Installing and removing the agent for EMC	
·	SRDF	17
	Before you install the agent for EMC SRDF	17
	Installing the agent for EMC SRDF	17
	Installing the agent IPS package on Oracle Solaris 11	
	systems	19
	Installing agent packages on Solaris brand non-global zones	19
	Upgrading the agent for EMC SRDF	20
	Removing the agent for EMC SRDF	21
Chapter 3	Configuring the agent for EMC SRDF	22
	Configuration concepts for the EMC SRDF agent	22
	Resource type definition for the EMC SRDF agent	22
	Attribute definitions for the SRDF agent	23
	Sample configuration for the EMC SRDF agent	29
	Additional configuration considerations for the SRDF agent	30
	Considerations for using the agent in concurrent or cascaded	
	SRDF configurations	30
	Before you configure the agent for EMC SRDF	32
	About configuring system zones in replicated data elusters	33 21
	About preventing split-brain	34 35

	Configuring the agent for EMC SRDF Configuring the agent manually in a global cluster Configuring the agent in an SF for Oracle RAC environment or Storage Foundation Cluster File System (SECES)	35 35
	environment	36
	Configuring the agent manually in a replicated data cluster	36
	Configuring monitoring of SRDF link	37
	Configuring the agent to compute RPO	37
	Setting the OnlineTimeout attribute for the SRDF resource	38
	Configuring LVM on AIX	40
	Configuring LVM on HP-UX	40
Chapter 4	Managing and testing clustering support for EMC SRDF	41
	How VCS recovers from various disasters in an HA/DR setup with	
	EMC SRDF	41
	Failure scenarios in global clusters	42
	Failure scenarios in replicated data clusters	46
	Testing the global service group migration	51
	Testing disaster recovery after host failure	53
	Testing disaster recovery after site failure	54
	Performing failback after a node failure or an application failure	56
	Performing failback after a site failure	57
Chapter 5	Setting up a fire drill	59
	About fire drills	59
	Fire drill configurations	60
	About the SRDFSnap agent	61
	Support for Multi-pathing	61
	SRDFSnap agent functions	61
	Resource type definition for the SRDFSnap agent	63
	Attribute definitions for the SRDFSnap agent	63
	About the Snapshot attributes	66
	Before you configure the fire drill service group	66
	Configuring the fire drill service group	68
	Creating the fire drill service group using Cluster Manager (Java	
	Console)	68
	Creating the fire drill service group using the Fire Drill SetUp	
	Wizard	70
	Verifying a successful fire drill	71
	Sample configuration for a fire drill service group	72

Contents | 9

ex

Chapter

Introducing the agent for EMC SRDF

This chapter includes the following topics:

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

About the agent for EMC SRDF

The Cluster Server agent for EMC Symmetrix Remote Data Facility (EMC SRDF) provides support for application failover and recovery. The agent provides this support in environments that use SRDF 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 resource online also has safe and exclusive access to the configured devices.

The agent for EMC SRDF supports the following:

- Replicated data clusters and global clusters that run VCS.
- SRDF device groups and consistency groups in synchronous and asynchronous modes only. The agent also supports dynamic SRDF (role swap).

Features of the EMC SRDF agent

The features of the EMC SRDF agent are:

- Support for VCS replicated data clusters (RDCs) and global clusters The EMC SRDF agent supports RDCs and global clusters that run VCS.
- Support for concurrent and cascaded SRDF configurations. The agent always makes sure that the devices local to the host on which the resource is online are in Primary (R1) role. The agent switches the SRDF configuration from concurrent to cascaded and vice versa. The agent automatically detects the underlying configuration and does not need any additional user input.
- Ability to compute recovery point objective (RPO) in a global cluster environment.

Note: The SRDF agent calculates RPO in a global cluster environment in an SFHA configuration with VxVM DiskGroups. SFRAC environments are not supported.

- Support for multiple remote SYMAPI servers.
- Support for N_Port ID Virtualization (NPIV) configuration on AIX.
- Support for fire drills.
- Support for detailed monitoring of replication link.

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

Note: Fire drill for SF Oracle RAC is supported only on VCS and SF Oracle RAC 5.0 MP3 versions onwards.

Supported software

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

Supported hardware for EMC SRDF

In environments using Storage Foundation for Oracle RAC, the arrays must support SCSI-3 persistent reservations.

The SRDF agent supports Solutions Enabler (SE) V6.4 up to V8.0.1. The SRDF agent also supports corresponding array microcode levels. The supported array models include EMC Symmetrix DMX and EMC Symmetrix V-Max family of arrays. Refer to the EMC hardware compatibility list for specific information.

Typical EMC SRDF setup in a VCS cluster

hosta hostb hostc hostd Replication link Primary array Secondary array (array 1) (array 2)

Figure 1-1 displays a typical cluster setup in a SRDF environment.

Figure 1-1 Typical clustering setup for the agent

VCS clusters using SRDF 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 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 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.



- The arrays at both the primary and secondary sites also have the BCV or target devices configured and associated with the corresponding replication devices at each site.
- 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 33.
- 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 EMC Symmetrix array.

EMC SRDF agent functions

The VCS agent for SRDF 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 one or more 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 EMC SRDF agent's online function" on page 15.
	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 Consistent or Synchronized state, the agent creates a lock file on the local host. The lock file indicates that the resource is online.
	It checks the dynamic swap capability of the array and individual devices. It also creates the swap lock file if the device group is capable of role swap. See "About dynamic swap support for the EMC SRDF agent" on page 16.
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.

Table 1-1Agent functions

Function	Description
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	Checks the dynamic swap capability of the array and individual devices. Creates the swap lock file if the device group is capable of role swap. See "About dynamic swap support for the EMC SRDF agent" on page 16.
	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.
	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 entry point fails or times out.
info	Reports the device state to the VCS interface. This function logs the state change information of SRDF link and device groups. This function also logs the dirty track or pending writes information.
action/update	Performs a $\tt symrdf$ <code>update</code> action from the R2 side to merge any dirty tracks from the R2 to the R1.
action/PreSwitch	Ensures that the remote site cluster can come online during a planned failover within a GCO configuration. The VCS engine on the remote cluster invokes the PreSwitch action on all the resources of the remote site during a planned failover using the hagrp -switch command.
	For this, the PreSwitch attribute must be set to 1. The option -nopre indicates that the VCS engine must switch the service group regardless of the value of the PreSwitch service group attribute.
	If running the PreSwitch action fails, the failover should not occur. This minimizes the application downtime and data loss.
	For more information on the PreSwitch action and the PreSwitch feature in the VCS engine, refer to the <i>Cluster Server Administrator's Guide.</i>

 Table 1-1
 Agent functions (continued)

Function	Description
action/GetCurrentRPO	Fetches the current point in time RPO. Invoke this action function on the DR system where the ComputeDRSLA attribute is set to 1. The RPO is computed in seconds.
	Note: The agent does not store the computed RPO; make a note of the RPO for future reference.
close	Deletes the swap lock file.
Attr_changed	Monitors the changes in the attribute GrpName. If the device group name is changed, the instructions are logged for the changes to be effective.
	Also monitors the ComputeDRSLA attribute. Depending on the new value of the ComputeDRSLA attribute, this function either initiates or terminates the process of computing the RPO.

Table 1-1 Agent functions (continued)

Note: The agent uses the following internal action functions to compute the RPO: StartRPOComputation, StopRPOComputation, StartWriter, ReportRPOData.

About the EMC SRDF 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 Consistent or Synchronized 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.

Depending on SRDF/S and SRDF/A, the states can be different as follows:

- For R2 devices in the SYNCHRONIZED or 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, the agent runs the symrdf command to make the devices writable.

The agent runs the command only if the AutoTakeover attribute is set to 1 and if there are no dirty tracks on the local device. Dirty tracks indicate that an out-of-order synchronization was in progress when the devices became partitioned, rendering them inconsistent and unusable. If dirty tracks exist, the online entry point faults on timeout.

- 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 CONSISTENT state.

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

See "Setting the OnlineTimeout attribute for the SRDF resource" on page 38.

About dynamic swap support for the EMC SRDF agent

The agent supports the SRDF/S and SRDF/A 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: <code>symrdf failover</code>. The command enables read-write on the R2 device.

The agent checks the following criteria before determining if a swap occurs:

- All devices in the device group are configured as dynamic devices.
- Dynamic RDF is configured on the local Symmetrix array.
- The microcode is level 5567 or later.

The commands for online are different for SRDF/S dynamic swap and SRDF/A dynamic swap as follows:

- For SRDF/S, for R2 devices in the SYNCHRONIZED state, the agent runs the symrdf failover -establish command.
- For SRDF/A, for R2 devices in the CONSISTENT state, the agent runs the symrdf -force failover command. If consistency is enabled, the agent runs the symrdf disable command. The agent then issues the symrdf swap command to do the role-swap and the establish command to re-establish the replication, and re-enables the consistency.

Dynamic swap does not affect the ability to perform fire drills.

Chapter

Installing and removing the agent for EMC SRDF

This chapter includes the following topics:

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

Before you install the agent for EMC SRDF

Before you install the Cluster Server agent for EMC SRDF, ensure that you install and configure the 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 setup in a VCS cluster" on page 12.

Installing the agent for EMC SRDF

You must install the EMC SRDF 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	cdl/aix/vcs/replication/srdf_agent/ agent_version/pkgs/
Linux	cdl/linux/generic/vcs/replication/srdf_agent/ agent_version/rpms/
Solaris	cdl/solaris/dist_arch/vcs/replication/srdf_agent agent_version/pkgs/

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

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

	Note: To install the agent IPS package on a Solaris 11 system, see Installing the agent IPS package on Oracle Solaris 11 systems.
Solaris	# pkgadd -d . VRTSvcse
Linux	<pre># rpm -ihv \ VRTSvcse-AgentVersion-Linux_GENERIC.noarch.rpm</pre>
AIX	<pre># installp -ac -d VRTSvcse.rte.bff VRTSvcse.rte</pre>

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

Installing the agent IPS package on Oracle Solaris 11 systems

To install the agent IPS package on an Oracle Solaris 11 system

- 1 Copy the VRTSvcse.p5p package from the pkgs directory to the system in the /tmp/install directory.
- 2 Disable the publishers that are not reachable as package install may fail if any of the already added repositories are unreachable.

pkg set-publisher --disable <publisher name>

where the publisher name is obtained using the pkg publisher command.

3 Add a file-based repository in the system.

pkg set-publisher -g /tmp/install/VRTSvcse.p5p Symantec

4 Install the package.

pkg install --accept VRTSvcse

5 Remove the publisher from the system.

pkg unset-publisher Symantec

6 Enable the publishers that were disabled earlier.

pkg set-publisher --enable <publisher name>

Installing agent packages on Solaris brand non-global zones

With Oracle Solaris 11, you must install the agent package inside non-global zones. The native non-global zones are called Solaris brand zones.

To install the agent package on Solaris brand non-global zones

1 Ensure that the SMF service

```
svc:/application/pkg/system-repository:default and
svc:/application/pkg/zones-proxyd:default are online on the global
zone.
```

- # svcs svc:/application/pkg/system-repository:default
- # svcs svc:/application/pkg/zones-proxyd:default
- 2 Log on to the non-global zone as a superuser.

3 Ensure that the SMF service

svc:/application/pkg/zones-proxy-client:default is online inside
non-global zone:

svcs svc:/application/pkg/zones-proxy-client:default

- 4 Copy the VRTSvcse.p5p package from the pkgs directory to the non-global zone (for example, at the /tmp/install directory).
- **5** Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

pkg set-publisher --disable <publisher name>

6 Add a file-based repository in the non-global zone.

```
# pkg set-publisher -g/tmp/install/VRTSvcse.p5p Symantec
```

- 7 Install the package.
 - # pkg install --accept VRTSvcse
- 8 Remove the publisher on the non-global zone.

pkg unset-publisher Symantec

9 Clear the state of the SMF service, as setting the file-based repository causes the SMF service svc:/application/pkg/system-repository:default to go into the maintenance state.

svcadm clear svc:/application/pkg/system-repository:default

10 Enable the publishers that were disabled earlier.

pkg set-publisher --enable <publisher>

Note: Perform steps 2 through 10 on each non-global zone.

Upgrading the agent for EMC SRDF

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" on page 21.

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

See "Installing the agent for EMC SRDF" on page 17.

- **5** Copy the file SRDFTypes.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

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

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

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

AIX	# installp -u VRTSvcse.rte
Linux	# rpm -e VRTSvcse
Solaris	# pkgrm VRTSvcse
	Note: To uninstall the agent IPS package on a Solaris 11 system:
	# pkg uninstall VRTSvcse

Chapter

Configuring the agent for EMC SRDF

This chapter includes the following topics:

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

Configuration concepts for the EMC SRDF agent

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

Resource type definition for the EMC SRDF agent

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

```
type SRDF (
    static int ActionTimeout = 180
    static keylist RegList = { ComputeDRSLA }
    static keylist SupportedActions = { update, ld2pd, PreSwitch,
    StartRPOComputation, StopRPOComputation, ReportRPOData,
    StartWriter, GetCurrentRPO }
    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, IsCompositeGroup, SwapRoles,
    DetailedMonitoring, SymapiServers, ComputeDRSLA, LinkMonitor,
```

```
AdvancedOpts }
str SymHome = "/usr/symcli"
str GrpName
int DevFOTime = 2
boolean AutoTakeover
boolean SplitTakeover
str Mode
boolean IsCompositeGroup
boolean SwapRoles = 1
int DetailedMonitoring = 0
str SymapiServers[]
int ComputeDRSLA
boolean Tagging = 0
temp str VCSResLock
temp int Counter = 0
temp str PrevState = "normal"
int LinkMonitor = 0
str AdvancedOpts{} = { ExtendMonitor=null }
```

Attribute definitions for the SRDF agent

)

Review the description of the agent attributes.

Required attributes

You must assign values to required attributes.

 Table 3-1
 Required attributes

Attribute	Description
GrpName	Name of the Symmetrix device group or composite group that the agent manages. Specify the name of a device group or composite group.
	Note: If this is a composite group, ensure that you set the value of IsCompositeGroup to 1.
	Type-dimension: string-scalar

Optional attributes

Configuring these attributes is optional.

Attribute	Description
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: integer-scalar
IsCompositeGroup	Specifies whether the SRDF group is a composite group or not. If set to 0, VCS treats it as device group. If set to 1, VCS treats it as composite group.
	Type-dimension: integer-scalar
	Default: 0
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 or composite group to fail over. This value helps the agent to determine whether it has adequate time for the online operation after waiting for other device or composite groups to fail over. If the online operation cannot be completed in the remaining time, the failover does not proceed.
	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 rw_enable operation on the partitioned devices at the secondary site.
	The default value of the AutoTakeover value is set to 0.
	If the AutoTakeover attribute is set to 1, it allows the SRDF agent to failover the service group in the DR site even when the replication is in the "Partitioned" state. The Partitioned state means that the replication link is broken out. This means that the secondary devices are not in sync with the primary devices or the secondary devices may have invalid data. Hence, the default value of the AutoTakover attribute set to 0, so that the failover can proceed only with the admin consent.
	For more information, refer to the EMC SRDF documentation.
	Default: 0
SplitTakeover	A flag that determines whether the agent permits a failover to R2 devices in the Split state. The value 0 indicates that the agent does not permit a failover to R2 devices in the Split state. The value 1 indicates that the agent permits a failover to R2 devices in the Split state if the devices are read-write enabled. The attribute has no effect on failing over to a host attached to R1 devices.
	Set the attribute to 0 to minimize the risk of data loss on a failover to devices that may not be in synch.
	Type-dimension: integer-scalar
	Default: 0

 Table 3-2
 Optional attributes (continued)

Attribute	Description
LinkMonitor	A flag that determines whether the agent should check the status of the replication link while bringing the resource online.
	This attribute is of importance only at the primary site where the role of the device group is RDF1 and all the devices in the device 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 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
Mode	Used at the time of failover to decide which commands to use to failover to the other site.
	The values for this attribute can be Asynchronous or Synchronous.
	If the value is not specified, the agent assumes that the mode is Synchronous. If the devices are setup to replicate in the Asynchronous mode, you must set Mode to Asynchronous.

 Table 3-2
 Optional attributes (continued)

Attribute	Description
DetailedMonitoring	Used to enable detailed monitoring of the SRDF link. This attribute is used in conjunction with the InfoInterval attribute. The numeric value of the InfoInterval attribute specifies how often the link monitoring cycles must run. The numeric value of the DetailedMonitoring attribute specifies the number of monitoring cycles after which a detailed level of monitoring is performed. If, during a detailed monitoring cycle, the link is found to be in a partitioned, split, or suspended state, the agent logs the dirty track or pending write information in the VCS engine logs.
	For example, if you set the value of the InfoInterval attribute as 10 and the value of the DetailedMonitoring attribute to 5:
	 An SRDF link monitoring cycle is run every 10 seconds. After every 5 monitoring cycles (that is, after every 50 seconds), a detailed level of monitoring is performed. If the links are in partitioned, split, or suspended state, dirty track or pending write information is logged.
	Type-dimension: integer-scalar
	Default: 0
SymapiServers	Used to configure remote SYMAPI servers if the Symmetrix Gatekeeper devices are not mapped to the host on which the application is running.
	Add appropriate entries in the SYMCLI client configuration file /var/symapi/config/netcnfg depicting the remote SYMAPI service name, remote SYMAPI server IP address, connection mode, and so on. Then add the SYMAPI remote service names in this attribute.
	Type-dimension: string-vector
	See "Additional configuration considerations for the SRDF agent" on page 30.

 Table 3-2
 Optional attributes (continued)

Attribute	Description
ComputeDRSLA	Used to enable or disable Recovery Point Objective (RPO) computation. Set this attribute on any one node in the disaster recovery (DR) cluster.
	Setting this attribute to 1 starts the RPO computation process. Ensure that you reset this attribute to 0 after you use the GetCurrentRPO action function to check the RPO.
	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 3-2
 Optional attributes (continued)

Internal attributes

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

Attribute	Description
VCSResLock	The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application. Type-dimension: temporary string
Counter	The agent uses the Counter attribute for monitoring of SRDF link state. Type-dimension: temporary integer

Table 3-3Internal attributes

Attribute	Description
PrevState	The agent uses the PrevState attribute for monitoring of SRDF link state. Type-dimension: temporary string
Tagging	This attribute is used for maintaining the process of computing RPO. Type-dimension: temporary boolean

 Table 3-3
 Internal attributes (continued)

Sample configuration for the EMC SRDF agent

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





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

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

Additional configuration considerations for the SRDF agent

Consider the following settings for configuring the SRDF agent:

- Set the OnlineTimeout attribute for the SRDF resource so that its entry points do not time out, or they automatically restart if they timed out.
 See "Setting the OnlineTimeout attribute for the SRDF resource" on page 38.
- 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.
- The SymapiServers attribute can be used to configure multiple remote SYMAPI servers. While all connection modes for remote SYMAPI operations are supported, Veritas recommends the SECURE connection mode.
- If the SymapiServers attribute is set, the SRDF agent tries to find the first available remote SYMAPI server from the configured list. If none of the configured servers are available, the agent does not come online. Hence, make sure that at least one remote SYMAPI server is always online.
- If the SymapiServer attribute is set, the SRDF agent does not support fire drills.
- If the SymapiServers attribute is set, some of the action entry points might not work.

Considerations for using the agent in concurrent or cascaded SRDF configurations

You can configure the SRDF agent to manage a concurrent or cascaded SRDF configuration. The agent automatically detects the underlying configuration and takes decisions during failovers.

The present behavior of the agent is that if there are more than two sites involved, the agent manages failovers only on the SYNC links. The agent always makes sure that the devices that are local to the host on which the resource is online, are in a primary (R1) role. The agent switches the SRDF configuration from concurrent to cascaded and vice versa. The agent automatically detects the underlying configuration and requires no manual intervention.



Figure 3-2 A sample concurrent SRDF configuration

Figure 3-2 shows a sample concurrent SRDF configuration.

When the SRDF resource goes online in site A, where both the other arrays are connected through SYNC/ASYNC links, the SRDF agent makes sure that the site A array becomes the primary in both the RDF groups.

When the SRDF resource fails over to some other site, for example site B, the SRDF agent makes sure that the site B array is the primary in the corresponding RDF group and the old primary becomes a secondary for the new primary, while keeping the old primary a primary for the second RDF group. Figure 3-3 shows an example of this new configuration.



Figure 3-3 A sample cascaded SRDF configuration

The SRDF agent does not handle failovers over any ASYNC links that may be configured in such a configuration. Veritas recommends not including the hosts in site C in the configuration shown in Figure 3-3 in the VCS service group configuration.

Before you configure the agent for EMC SRDF

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 setup in a VCS cluster" on page 12.
- Make sure that the cluster has an effective heartbeat mechanism in place.
 See "About cluster heartbeats" on page 33.
 See "About preventing split-brain" on page 35.

- Set up system zones in replicated data clusters.
 See "About configuring system zones in replicated data clusters" on page 34.
- Verify that the clustering infrastructure is in place.
 - If you plan to configure the agent in a global cluster, make sure the global service group for the application is configured.
 For more information, refer to the *Cluster Server Administrator's Guide*.
 - If you want to configure the agent in an SF Oracle RAC environment, verify that the SF Oracle RAC global cluster infrastructure is in place.
 - If you plan to configure the agent in a replicated data cluster, make sure the required replication infrastructure is in place and that the application is configured.

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.
- The heartbeat is then managed completely by VCS. VCS reports that the site is down only when the remote array is not visible by the symrdf ping command.

About configuring system zones in replicated data clusters

In a replicated data cluster, you can prevent unnecessary SRDF 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 3-4 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 3-4 Example system zone configuration

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

When the SRDF runs on R2 devices, SRDF does not synchronize data back to the R1 automatically. You must update out-of-synch tracks manually. Monitor the number of out-of-synch tracks by viewing the ResourceInfo attribute of an online SRDF resource. If the value is too high, update tracks to the R1 using the update action. The update action is defined as a supported action in the SRDF 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 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 mean the loss of replication links.

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

Configuring the agent for EMC SRDF

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

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

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

See "Sample configuration for the EMC SRDF agent" on page 29.

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 SRDF 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 (Java Console) and log on to the cluster.
- 2 Add a resource of type SRDF at the bottom of the service group.
- **3** Configure the attributes of the SRDF 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 CVMVolDg resource.
- **2** Add the following keys to the list: import, deport, and vxdctlenable.
- 3 Run the following commands 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 defines a list of action tokens for the resource.

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 (Java Console) and log on to the cluster.
- 2 In each service group that uses replicated data, add a resource of type SRDF at the bottom of the service group.
- **3** Configure the attributes of the SRDF 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.

Configuring monitoring of SRDF link

The info agent function uses the InfoInterval attribute to monitor the SRDF link state and is disabled by default. The SRDF link state changes are logged with appropriate error and warning levels. If DetailedMonitoring is set with a specific frequency, the dirty track information is updated and logged in VCS engine logs.

To configure the DetailedMonitoring attribute:

1 Set the InfoInterval attribute of SRDF resource type to a defined interval value. If the value of the InfoInterval attribute is set to zero, it disables the monitoring. To set the value of InfoInterval, run the following command:

hatype -modify SRDF InfoInterval <Value>

where, <Value> is the numeric value in seconds

2 To get dirty track or pending writes information, set the DetailedMonitoring attribute to a defined frequency. All the information is logged in the engine_A.log file with appropriate error and severity levels. To set the value of the DetailedMonitoring attribute, run the following command:

hares -modify srdf res DetailedMonitoring <Value>

where, <*Value*> is the numeric value in seconds

Configuring the agent to compute RPO

In a global cluster environment, the agent for EMC SRDF can compute the recovery point objective (RPO), which is a disaster recovery (DR) SLA. In a DR configuration where data is replicated asynchronously to the DR site, the DR site data is not always as current as the primary site data.

RPO is the maximum acceptable amount of data loss in case of a disaster at the primary site. The agent computes RPO in terms of time, that is, in seconds.

Note: The agent calculates RPO in a global cluster environment in an SFHA configuration with VxVM disk groups. The agent does not calculate RPO in SFRAC environments.

Before you configure the agent to compute the RPO, ensure that the following pre-requisites are met:

- The service group containing the SRDF resource and the VxVM disk group resource are online at the production site.
- The disk group resource is dependent on the SRDF resource.

To configure the agent to compute the RPO:

1 In the DR cluster, on any one of the nodes where devices are asynchronously replicated and where the service group is configured, run the following command to start the RPO computation:

hares -modify SRDF resource name ComputeDRSLA 1 -sys system name.

2 Run the following command on the same node in the DR cluster:

hares -action SRDF_resource_name GetCurrentRPO -sys system_name

The action entry point displays the RPO. The agent does not store the computed RPO; make a note of the RPO for future reference.

If the RPO is not reported, it indicates that the agent needs more time to finish computing the RPO. Wait for some more time before you run the GetCurrentRPO action function again.

3 To stop RPO computation, run the following command:

hares -modify SRDF_resource_name ComputeDRSLA 0 -sys system_name

Note: The agent might not compute the RPO correctly when the SwapRoles attribute is disabled and the replication is in the Failedover state.

Setting the OnlineTimeout attribute for the SRDF resource

Set the OnlineTimeout attribute for the SRDF 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 resource in the configuration, use the following formula to calculate an appropriate value for the OnlineTimeout attribute:

 $\mathbf{OnlineTimeout} = \sum_{1}^{n} ((n_{\text{devices}} \times f_{\text{dilovertime}}) + \varepsilon)$

- n_{devices} represents the number of devices in a device group.
- d_{failovertime} represents the time taken to failover a device.
- n_{devicegroups} represents the total number of device 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 device group (typically the case for SRDF), multiply the number of devices in the device 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 device 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 device groups (currently not supported by SRDF), calculate the OnlineTimeout attribute for all device groups and set the OnlineTimeout attribute to at least the amount of time the largest device 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.
- To set the OnlineTimeout attribute using the sigma script
- Run the sigma script to get recommendations for VCS attribute values.

/opt/VRTSvcs/bin/SRDF/sigma

Run the script on a node where VCS is running and has the SRDF agent configured.

The sigma calculator adds 10 seconds to the value for each device group to compensate for the overhead of launching an appropriate symrdf command. Specify another value to the sigma script if the instantiation takes shorter or longer.

The script runs on the assumption that the VCS program manages all devices in the array. Other operations outside of VCS that hold the array lock might delay the online operation unexpectedly.

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.

Configuring LVM on HP-UX

To support failover of the LVM volume groups to the secondary site during a disaster or normal switch, create the LVM volume group on the primary site and export the volume group using the following command:

vgexport [-p] [-v] [-s] [-m]/vg04map.map vg04.

Copy the map file to the secondary site and then import the volume group on the secondary using the map file. Run the following command:

vgimport [-s] [-v] [-m] /vg04map.map vg04.

This must be done as part of an initial setup process before VCS starts controlling the replication.

To configure LVM on HP-UX

- 1 Configure the volume groups on a replicated primary lun.
- 2 Create the resources SRDF, LVMGroup, LVMVolume and mount and bring them online on the primary site.
- **3** Bring the resources offline on the primary site and online on the secondary. The resources must be successfully brought online on the secondary site.

Chapter

Managing and testing clustering support for EMC SRDF

This chapter includes the following topics:

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

This section covers the failure scenarios and how VCS responds to the failures for 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 Cluster Server agent for EMC SRDF ensures safe and exclusive access to the configured EMC SRDF devices.

See "Failure scenarios in global clusters" on page 42.

Replicated data	When service group or system faults occur, VCS failover behavior
clusters	depends on the value of the AutoFailOver attribute for the faulted service
	group. The VCS agent for EMC SRDF ensures safe and exclusive
	access to the configured EMC SRDF devices.

See "Failure scenarios in replicated data clusters" on page 46.

Failure scenarios in a global cluster configuration with the Cluster

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

Table 4-1 lists the failure scenarios in a global cluster configuration and describes the behavior of VCS and the agent in response to the failure.

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. 	
	Agent response:	
	 The agent Write enables the devices at the secondary site. 	
	For dynamic RDF devices, the agent does the following if the value of the SwapRoles attribute of the SRDF resource is 1:	
	 Swaps the R1/R2 personality of each device in the device group or the consistency group. 	
	 Restarts replication from R1 devices on the secondary site to the R2 devices at the primary site. 	
	See "Performing failback after a node failure or an application failure" on page 56.	

Server agent for EMC SRDF

Failure scenarios in a global cluster configuration with the Cluster

Server agent for EMC SRDF (continued)	
Failure	Description and VCS response
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.
	Agent response:
	 The agent write enables the devices at the secondary site.
	For dynamic RDF devices, the agent does the following if the value of the SwapRoles attribute of the SRDF resource is 1:
	 Swaps the R1/R2 personality of each device in the device group or the consistency group.
	 Restarts replication from R1 devices on the secondary site to the R2 devices at the primary site.
	See "Performing failback after a node failure or an application failure" on page 56.
Site failure	All hosts and the storage at the primary site fail.
	A site failure renders the devices on the array at the secondary site in the PARTITIONED state.
	VCS response at the secondary site:
	 Displays an alert to indicate the 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.
	Agent response: The agent does the following on the secondary site in case of a manual failover based on the value of the AutoTakeover attribute of the SRDF resource:
	 1—If invalid tracks do not exist, the agent issues the symrdf failover command to make the SRDF devices write-enabled. 0—The agent faults the SRDF resource.
	See "Performing failback after a site failure" on page 57.

Table 4-1

Failure scenarios in a global cluster configuration with the Cluster

Server agent for EMC SRDF (continued)		
Failure	Description and VCS response	
Replication link failure	Replication link between the arrays at the two sites fails.	
	A replication link failure renders the SRDF devices in the PARTITIONED state. When the link is restored, the SRDF devices attain the SUSPENDED state.	
	VCS response: No action.	
	The DetailedMonitoring attribute monitors the SRDF link status and reports the dirty track information that the link is in PARTITIONED state.	
	After the link is restored, you must resynchronize the SRDF devices.	
	To resynchronize the SRDF devices after the link is restored:	
	 Before you resync the R2 device, you must split the BCV or target device from the R2 device at the secondary site. 	
	 You must initiate resync of R2 device using the symrdf resume command. After R1 and R2 devices are in sync, reestablish the mirror relationship between the BCV or target devices and the R2 devices. 	
	If you initiate a failover to the secondary site when resync is in progress, the online function of the EMC SRDF agent waits for the resync to complete and then initiates a takeover of the R2 devices.	
	Note: If you did not configure BCV or target devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Veritas recommends configuring BCV or target devices at both the sites.	
	See "Typical EMC SRDF setup in a VCS cluster" on page 12.	

Table 4-1

	Table 4-1Failure scenarios in a global cluster configuration with the Cluster Server agent for EMC SRDF (continued)
Failure	Description and VCS response
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 attribute: Manual or Connected—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. Auto—VCS brings the global group online at the secondary site which may lead to a site-wide split brain. This causes data divergence between the devices on the primary and the secondary arrays. When the network (wac and replication) connectivity restores, you must manually resync the data.
	 Note: Veritas recommends that the value of the ClusterFailOverPolicy attribute is set to Manual for all global groups to prevent unintended failovers due to transient network failures. To resynchronize the data after the network link is restored: Take the global service group offline at both the sites. Manually resynchronize the data. Depending on the site whose data you want to retain use the symrdf establish or the symrdf restore command
	 Bring the global service group online on one of the sites. Agent response: Similar to the site failure.

Server agent for EMC SRDF (continued)	
Failure	Description and VCS response
Storage failure	The array at the primary site fails.
	A storage failure at the primary site renders the devices on the array at the secondary site in the PARTITIONED state.
	VCS response at the secondary site:
	 Causes the 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 service group online at the secondary site.
	 Manual—No action. You must bring the global group online at the secondary site.
	Agent response: The agent does the following on the secondary site in case of a manual failover based on the value of the AutoTakeover attribute of the SRDF resource:
	• 1—If invalid tracks do not exist, the agent issues the symrdf failover command to make the SRDF devices write-enabled.
	 0—The agent faults the SRDF resource.

Table 4-1Failure scenarios in a global cluster configuration with the Cluster
Server agent for EMC SRDF (continued)

Failure scenarios in replicated data clusters

Table 4-2 lists the failure scenarios in a replicated data cluster configuration, and describes the behavior of VCS and the agent in response to the failure.

	Table 4-2Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF
Failure	Description and VCS response
Application failure	 Application cannot start successfully on any hosts at the primary site. VCS response: Causes the service group at the primary site to fault. Does the following based on the AutoFailOver attribute for the faulted service group: 1—VCS automatically brings the faulted service group online at the secondary site. 2—You must bring the service group online at the secondary site. Agent response: The agent write enables the devices at the secondary site. For dynamic RDF devices, the agent does the following if the value of the SwapRoles attribute of the SRDF resource is 1: Swaps the R1/R2 personality of each device in the device group or the consistency group. Restarts replication from R1 devices on the secondary site to the R2 devices at the primary site.
Host failure	 All hosts at the primary site fail. VCS response: Causes the service group at the primary site to fault. Does the following based on the AutoFailOver attribute for the faulted service group: 1—VCS automatically brings the faulted service group online at the secondary site. 2—You must bring the service group online at the secondary site. Agent response: The agent write enables the devices at the secondary site. For dynamic RDF devices, the agent does the following if the value of the SwapRoles attribute of the SRDF resource is 1: Swaps the R1/R2 personality of each device in the device group or the consistency group. Restarts replication from R1 devices on the secondary site to the R2 devices at the primary site.

Table	e 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF <i>(continued)</i>
Failure	Description and VCS response
Site failure	All hosts and the storage at the primary site fail.
	A site failure renders the devices on the array at the secondary site in the PARTITIONED state.
	VCS response:
	 Causes the service group at the primary site to fault.
	 Does the following based on the AutoFailOver attribute for the faulted service group:
	 1—VCS automatically brings the faulted service group online at the secondary site.
	 2—You must bring the service group online at the secondary site.
	Agent response: The agent does the following based on the AutoTakeover attribute of the SRDF resource:
	 1—If invalid tracks do not exist, the agent issues the symrdf failover command to make the SRDF devices write-enabled. 0—The agent faults the SRDF resource.
	See "Performing failback after a site failure" on page 57.

Failure scenarios in a replicated data cluster configuration with

VCS agent for EMC SRDF (continued)		
Failure	Description and VCS response	
Replication link failure	Replication link between the arrays at the two sites fails.	
	A replication link failure renders the SRDF devices in the PARTITIONED state. When the link is restored, the SRDF devices attain the SUSPENDED state.	
	VCS response: No action.	
	The DetailedMonitoring attribute monitors the SRDF link status and reports the dirty track information that the link is in PARTITIONED state.	
	After the link is restored, you must resynchronize the SRDF devices.	
	To resynchronize the SRDF devices after the link is restored:	
	1 Before you resync the R2 device, you must split the BCV or target device from the R2 device at the secondary site.	
	2 You must initiate resync of R2 device using the update action entry point.	
	3 After R1 and R2 devices are in sync, reestablish the mirror relationship between the BCV or target devices and R2 devices.	
	If you initiate a failover to the secondary site when resync is in progress, the online function of the EMC SRDF agent waits for the resync to complete and then initiates a takeover of the R2 devices.	
	Note: If you did not configure BCV or target devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Veritas recommends configuring BCV or target devices at both the sites.	
	See "Typical EMC SRDF setup in a VCS cluster" on page 12.	

Table 4-2

Failure scenarios in a replicated data cluster configuration with

VCS agent for EMC SRDF (continued)	
Failure	Description and VCS response
Network failure	The LLT and the replication links between the sites fail. VCS response:
	 VCS at each site concludes that the nodes at the other site have faulted. Does the following based on the AutoFailOver attribute for the faulted service group:
	 2—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. 1—VCS brings the service group online at the secondary site which leads to a cluster-wide split brain. This causes data divergence between the devices on the arrays at the two sites.
	When the network (LLT and replication) connectivity is restored, VCS takes all the service groups offline on one of the sites and restarts itself. This action eliminates concurrency violation where in the same group is online at both the sites.
	After taking the service group offline, you must manually resync the data using the symrdf establish or the symrdf restore command.
	Note: Veritas recommends that the value of the AutoFailOver attribute is set to 2 for all service groups to prevent unintended failovers due to transient network failures.
	To resynchronize the data after the network link is restored:
	1 Take the service groups offline at both the sites.
	2 Manually resynchronize the data.
	Depending on the site whose data you want to retain use the symrdf establish or the symrdf restore command.
	3 Bring the service group online on one of the sites.
	Agent response: Similar to the site failure.

Table 4-2

Failure scenarios in a replicated data cluster configuration with

VCS agent for EMC SRDF (continued)	
Failure	Description and VCS response
Storage failure	 The array at the primary site fails. A storage failure at the primary site renders the devices on the array at the secondary site in the PARTITIONED state. VCS response: Causes the service group at the primary site to fault and displays an alert to indicate the fault. Does the following based on the AutoFailOver attribute for the faulted service group: 1—VCS automatically brings the faulted service group online at the secondary site. 2—You must bring the service group online at the secondary site. Agent response: The agent does the following based on the AutoTakeover attribute of the SRDF resource: 1—If invalid tracks do not exist, the agent issues the symrdf failover command to make the SRDF devices write-enabled. 0—The agent does not perform failover to the secondary site.

Table 4-2

Testing the global service group migration

After you configure the Cluster Server agent for EMC SRDF, 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.

 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 or CONSISTENT 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 or CONSISTENT.

To test service group migration in replicated data cluster setup

1 Fail over the service group from the primary site to the secondary site. Perform the following steps:

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

hagrp -switch service group -to sys name

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

 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

- U—FAILED OVER State
- 1—SYNCHRONIZED or CONSISTENT state
- 2 Fail back the service group from the secondary site to the primary site. Perform the following steps:
 - Switch the service group from the secondary site to any node in the primary site.

hagrp -switch service group -to sys name

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

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

Testing disaster recovery after host failure

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

See "Failure scenarios in global clusters" on page 42.

See "Failure scenarios in replicated data clusters" on page 46.

Depending on the DR configuration, perform one of the following procedures to test how VCS recovers after all hosts at the primary site fail.

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 or CONSISTENT state

To test disaster recovery for host failure in replicated data cluster setup

1 Halt the hosts at the primary site.

The value of the AutoFailOver attribute for the faulted service group determines the VCS failover behavior.

- 1—VCS brings the faulted service group online at the secondary site.
- 2—You must bring the service group online at the secondary site.
 On a node in the secondary site, run the following command:

hagrp -online service group -to sys name

2 Verify that the 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 or CONSISTENT state

Testing disaster recovery after site failure

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

See "Failure scenarios in global clusters" on page 42.

See "Failure scenarios in replicated data clusters" on page 46.

Depending on the DR configuration, perform one of the following procedures to test the disaster recovery in the event of site failure.

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

To test disaster recovery for site failure in replicated data cluster setup

1 Halt all hosts 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 AutoFailOver attribute for the faulted global service group determines the VCS failover behavior.

- 1—VCS brings the faulted global service group online at the secondary site.
- 2—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 service_group -sys sys_name

- 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 behave in response to these failures.

See "Failure scenarios in global clusters" on page 42.

See "Failure scenarios in replicated data clusters" on page 46.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site. Depending on your DR configuration, perform one of the following procedures.

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 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 SRDF resource:
 - 1—Write enables the devices at the primary site, swaps the R1/R2 personality of each device in the device group or the consistency 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 or CONSISTENT.

To perform failback after a host failure or an application failure in replicated data cluster

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

hagrp -switch service group -to sys name

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

The VCS agent for EMC SRDF 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 SRDF resource:

- 1—Write enables the devices at the primary site, swaps the R1/R2 personality of each device in the device group or the consistency 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 or CONSISTENT.

Performing failback after a site failure

After a site failure at the primary site, the hosts and the storage at the primary site are down. VCS brings the global service group online at the secondary site and the EMC SRDF agent write enables the R2 devices.

The device state is PARTITIONED.

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

See "Failure scenarios in global clusters" on page 42.

See "Failure scenarios in replicated data clusters" on page 46.

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 CONSISTENT or 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.

To perform failback after a site failure in replicated data cluster

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

hagrp -offline service group -sys sys name

2 Resync the devices using the symrdf restore command.

The symrdf restore command write disables the devices at both the R1 and R2 sites.

After the resync is complete, the device state is CONSISTENT or 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 service_group -sys sys_name

This again swaps the role of R1 and R2.

Chapter

Setting up a fire drill

This chapter includes the following topics:

- About fire drills
- Fire drill configurations
- About the SRDFSnap agent
- Before you configure the fire drill service group
- Configuring the fire drill service group
- Verifying a successful fire drill
- Sample configuration for a fire drill service group

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 SRDF, the SRDFSnap 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 SRDFSnap agent supports fire drill for storage devices that are managed using Veritas Volume Manager.

The agent supports fire drill in a Storage Foundation for Oracle RAC environment.

Note: On AIX, the SRDFSnap agent supports LVM only in the Gold configuration.

Fire drill configurations

VCS supports the Gold, Silver, and Bronze fire drill configurations for the agent.

Note: The values of the UseSnapshot, MountSnapshot, and RequireSnapshot attributes define the fire drill configuration. See "About the Snapshot attributes" on page 66.

Gold	Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.
	Veritas recommends this configuration because it does not affect production recovery.
	In the Gold configuration, VCS does the following:
	 Takes a snapshot of the replicated LUNS using the BCV/TGT/VDEV device on the target array. Modifies the disk group name in the snapshot. Brings the fire drill service group online using the snapshot data.
	For non-replicated Symmetrix devices:
	 You must use Veritas Volume Manager. Additionally, on AIX, you can also use LVM. You must use the Gold configuration without the option to run in the Bronze mode.
Silver	VCS takes a snapshot, but does not run the fire drill on the snapshot data. VCS breaks replication and runs the fire drill on the replicated target device.
	If a disaster occurs while resynching data after running the fire drill, you must switch to the snapshot for recovery.
	In the Silver configuration, VCS does the following:
	 Takes a snapshot of the replicated LUNS using the BCV/TGT/VDEV device on the target array.
	 Modifies the disk group name in the snapshot. Brings the fire drill service group online using the data on the target array; the agent does not use the snapshot data for the fire drill.

Bronze VCS breaks replication and runs the fire drill test on the replicated target devices. VCS does not take a snapshot in this configuration.

If a disaster occurs while resynching data after the test, it may result in inconsistent data as there is no snapshot data.

In the Bronze configuration, VCS does the following:

- Splits replication.
- Modifies the disk group name while importing.
- Brings the fire drill service group online using the data on the target array.

About the SRDFSnap agent

The SRDFSnap agent is the fire drill agent for EMC SRDF.

The agent manages the replication relationship between the source and target arrays when running a fire drill. Configure the SRDFSnap resource in the fire drill service group, in place of the SRDF resource.

The SRDFSnap agent supports TimeFinder/Snap technology for space-optimized snapshots using VxVM. The agent does not support TimeFinder/Snap with AIX LVM or CVMVoIDG.

The SRDFSnap agent supports only BCV in a CVMVoIDG environment.

The agent supports space-optimized snapshots only for the configuration in which the snapshot must be taken for the RDF2 device.

Note: If the SymapiServer attribute is set, the SRDFSnap agent does not support fire drills

Support for Multi-pathing

The SRDFSnap agent supports fire drills for disks that are under the control of the following multi-pathing solutions:

- EMC PowerPath
- Dynamic Multi-Pathing (DMP)

SRDFSnap agent functions

The SRDFSnap agent performs the following functions:

Function	Description
online	 For Gold and Silver configurations, takes a local snapshot of the target LUN. For Silver and Bronze configurations, splits replication between the source and the target arrays. For Gold Configurations, brings the fire drill service group online by mounting the local snapshot. For Silver and Bronze Configurations, brings the fire drill service group online by mounting the local snapshot. For Silver and Bronze Configurations, brings the fire drill service group online by mounting the replication target LUN. Creates a lock file to indicate that the resource is online. For AIX LVM, the agent runs the LVM command recreatevg to create the fire drill volume group.
offline	 Destroys the snapshot by synchronizing data between the target array and the device on which snapshot was taken. For AIX LVM, the agent runs the LVM command exportvg to remove the fire drill volume group entries from the ODM. For Silver and Bronze configurations, resumes replication between the source and the target arrays. Removes the lock file created by the online operation.
monitor	Verifies the existence of the lock file to make sure the resource is online.
clean	Restores the state of the LUNs to their original state after a failed online function.
attr_changed	 Monitors the change in the value of the following attributes and verifies that the new value is not invalid. CopyMode SavePoolName DiskGroupSnapList Verifies that if a value has been assigned to the SavePoolName attribute, then the value of the CopyMode attribute is set to 'snap'. Also verifies that the value of the DiskGroupSnapList attribute is in the correct format.

Table 5-1Agent functions

Resource type definition for the SRDFSnap agent

Following is the resource type definition for the SRDFSnap agent:

```
type SRDFSnap (
    static keylist RegList = { MountSnapshot, UseSnapshot,
                               SavePoolName, CopyMode,
                               DiskGroupSnapList }
    static int OpenTimeout = 180
    static str ArgList[] = { TargetResName, MountSnapshot,
                             UseSnapshot, RequireSnapshot,
                             DiskGroupSnapList, CopyMode,
                             UseTgt, SavePoolName }
    str TargetResName
    str DiskGroupSnapList
    boolean MountSnapshot = 1
    boolean UseSnapshot = 1
    boolean RequireSnapshot = 1
    str SavePoolName
    str CopyMode = mirror
    boolean UseTgt = 0
    temp str Responsibility
    temp str FDFile
    temp str VCSResLock
)
```

Attribute definitions for the SRDFSnap agent

Table 5-2 describes the attributes that you can configure to customize the behavior of the SRDFSnap agent.

TargetResName	Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the SRDF 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 Oracle 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.
	This is a required attribute.
	Type-Dimension: string-scalar
DiskGroupSnapList	Lists the original application disk group names and the fire drill disk group names.
	Set the value of this attribute in the format: application_diskgroup_name;firedrill_diskgroup_name.
	For example, if your application disk group name is Org_DG and you want to name the fire drill disk group New_DG, set this attribute to Org_DG;New_DG.
	If you don't assign a value to this attribute, the SRDFSnap agent appends '_fd' to the application disk group, to name the fire drill disk group.
	Type-dimension: string-scalar
MountSnapshot	Specifies whether the resource uses the snapshot to bring the service group online.
	Set this attribute to 1 if you want the agent to bring the service group online using the local snapshot.
	Type-Dimension: boolean
	Default: 1
	Note: Set this attribute to 1 only if UseSnapshot is set to 1.
UseSnapshot	Specifies whether the SRDFSnap resource takes a local snapshot of the target devices. Set this attribute to 1 if you want the agent to take a local snapshot of the target devices.
	Type-Dimension: boolean
	Default: 1
	See "About the Snapshot attributes" on page 66.

Table 5-2	SRDFSnap agent attributes
-----------	---------------------------

RequireSnapshot	Specifies whether the SRDFSnap resource must take a snapshot before coming online.
	Set this attribute to 1 if you want the resource to come online only after it succeeds in taking a snapshot.
	Type-Dimension: boolean
	Default: 1
	Note: Set this attribute to 1 only if UseSnapshot is set to 1.
SavePoolName	Specifies the name of the SAVE pool to be used for space-optimized snapshots. If the attribute is not set, then DEFAULT_POOL is used as SAVE pool name for space-optimized snapshots.
	Type-Dimension: string-scalar
	Note: Set this attribute only if the value of the CopyMode attribute is set to 'snap'.
CopyMode	Indicates the array snapshot technology that the agent must use. The accepted values are:
	 mirror: indicates TimeFinder/Mirror. This is the default value. clone: indicates TimeFinder/Clone. snap: indicates TimeFinder/Snap.
	Type-dimension: string-scalar
UseTgt	Determines if the agent must use target devices or BCVs in the device group.
	The value 0 indicates that the agent must use BCV devices, and the value 1 indicates that the agent must use target devices. By default, the value of this attribute is set to 0.
	Type-dimension: boolean
	Note: Set this attribute only if the value of the CopyMode attribute is set to 'clone'.
Responsibility	Do not modify. For internal use only.
	Used by the agent to keep track of resynchonizing snapshots.
	Type-Dimension: temporary string

 Table 5-2
 SRDFSnap agent attributes (continued)

FDFile	Do not modify. For internal use only.
	Used by the agent to store the absolute pathname to the file with the latest fire drill report on the local system.
	Type-Dimension: temporary string

 Table 5-2
 SRDFSnap agent attributes (continued)

About the Snapshot attributes

The UseSnapshot, MountSnapshot, and RequireSnapshot attributes define the fire drill configuration.

Table 5-3 lists the snapshot attribute values for fire drill configurations:

Attribute	Gold	Silver	Bronze
MountSnapshot	1	0	0
UseSnapshot	1	1	0

 Table 5-3
 Snapshot attribute values for fire drill configurations

Setting the RequireSnapshot attribute to 0 enables a Gold or Silver configuration to run in the Bronze mode if the snapshot operation fails.

Before you configure the fire drill service group

Before you configure the fire drill service group, ensure that the following pre-requisites are met:

- Make sure the application service group is configured with a SRDF resource.
- Make sure the infrastructure to take appropriate snapshots (mirror/clone/snap) is properly configured on the target arrays.
- If you plan to run a fire drill on SRDF/A devices, you must have a TimeFinder/CG license.
- When you use the Gold or Silver configuration, make sure TimeFinder for SRDF is installed and configured at the target array.
- When you take snapshots of R2 devices, consider the following:
 - For TimeFinder/Mirror, BCV devices must be associated with and attached to the RDF2 device group or composite group, and fully established and synchronised with the R2 devices.

- For TimeFinder/Clone, BCV or target devices must be associated with the RDF2 device group or composite group.
- For TimeFinder/Snap, or space-optimized snapshots, the VDEV and SAVE devices must be associated with the device group or composite group for which you want to run the fire drill using space-optimized snapshots. Ensure that the SAVE pool that is defined in the SavePoolName attribute exists and that the SAVE devices are enabled in the SAVE pool.
- When you take snapshots of non-replicated VxVM disk groups residing on Symmetrix devices, create a Symmetrix device group with the same name as the VxVM disk group. The device group must contain the same devices as in the VxVM disk group and additionally, have the same number of BCVs or target devices associated.

If you use LVM on AIX, the LVM volume group must have the same name as the device group.

- For non-replicated devices:
 - You must use the Gold configuration without the option to run in the Bronze mode. Set the RequireSnapshot attribute to 1.
- Add vxdctlenable action in the list of SupportedActions for the CVMVolDg resource in an SF for Oracle RAC environment.
- If you plan to run a fire drill using space-optimized snapshots, you must have a TimeFinder/Snap license.
- Make sure that the VDEV devices and SAVE devices are associated with the device group or composite group for which you want to run fire drill using space-optimized snapshots.
- Make sure that the SAVE pool as specified by SavePoolName attribute exists prior to running fire drill using space-optimized snapshots.
- Make sure that the copy sessions are not created for the device or composite group prior to running fire drill with space-optimized snapshots.
- Make sure that the SRDF mode of replication is set to Synchronous prior to running fire drill using space-optimized snapshots and clones. This is because EMC does not support creation of TimeFinder/Snap and TimeFinder/Clone copy sessions for RDF2 device, if the SRDF mode of replication is set to Asynchronous.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.

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 72.

You can use any of the following interfaces to create the fire drill service group:

- Cluster Manager (Java Console). See "Creating the fire drill service group using Cluster Manager (Java Console)" on page 68.
- Fire Drill SetUp Wizard. See "Creating the fire drill service group using the Fire Drill SetUp Wizard" on page 70.

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.
- 6 Add a fire drill service group dependency:
 - Click the Service Group tab in the left pane and select the fire drill service group.

- Right-click the fire drill service group and select Link.
- Select the application service group and select offline local to define the relationship.
- 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.
- 6 Click OK to close the Edit Attribute dialog box.
- 7 Click the Save and Close Configuration icon in the toolbar.

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' suffix. Click **Apply**.
- 6 Click OK.

Configuring resources for fire drill service group

Edit the resources in the fire drill service group so 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 SRDF resource and click Delete.
- **4** Add a resource of type SRDFSnap 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.

Note: If the application service group has resources with local attribute values, you must manually set these attributes after creating the resources.

Creating the fire drill service group using the Fire Drill SetUp Wizard

This section describes how to use the Fire Drill SetUp Wizard to create the fire drill service group.

To create the fire drill service group

1 Start the Fire Drill SetUp Wizard

/opt/VRTSvcs/bin/fdsetup-srdf

- 2 Enter the name of the application service group for which you want to configure a fire drill service group.
- **3** Select the supported snapshot technology (clone, mirror, snap)
- **4** If you selected clone or snap, enter the following:
 - For clone: Do you wish to use Target Devices? [y,n,q] (n)
 - For snap: Do you wish to specify the SAVE Pool name? [y,n,q] (n)
- **5** Select the supported snapshot configurations Gold, Silver, or Bronze.
- 6 If the snapshot fails with Gold or Silver configurations, choose whether to run a Bronze fire drill

If snapshot fails, should bronze be used? [y,n,q](n)

7 Press **Return** to verify the snapshot infrastructure.

8 In the Snapshot Details, the wizard informs whether all the devices in the device group on the target array has Synchronized state.

If the devices are in synchronized state, press Return.

If the devices are not in synchronized state, wait until the devices comes in Synchronized state and run the wizard again.

9 Enter **y** to create the fire drill service group.

The wizard runs various commands to create the fire drill service group.

10 Schedule fire drill for the service group by adding the following command to the crontab to be run at regular intervals.

/opt/VRTSvcs/bin/fdsched-srdf firedrill servicegroupname

11 Make fire drill highly available by adding the above command to the crontab on every node in this cluster.

Configuring local attributes in the fire drill service group

The fire drill setup wizard does not recognize localized attribute values for resources. If the application service group has resources with local (per-system) attribute values, you must manually set these attributes after running the wizard.

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.

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 SRDFSnap resource replaces the SRDF resource.

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

```
SRDFSnap SRDFSnap-res_srdf (
   TargetResName = res_srdf
   MountSnapshot = 1
   UseSnapshot = 1
   RequireSnapshot = 1
   CopyMode = clone
   UseTgt = 1
   )
```
Index

Α

action function 13 attribute definitions 23 AutoTakeover attribute 23

С

clean function 13 cluster heartbeats 33

D

DetailedMonitoring 23 DevFOTime attribute 23 disaster recovery 41

Ε

EMC SRDF agent attribute definitions 23 EMC SRDF agent attributes AutoTakeover 23 ComputeDRSLA 23 DetailedMonitoring 23 DevFOTime 23 GrpName 23 IsCompositeGroup 23 Mode 23 SplitTakeover 23 SwapRoles 23 SymapiServers 23 SymHome 23 VCSResLock 28

F

failure scenarios 41 global clusters 42 application failure 42 host failure 42 network failure 42 replication link failure 42 site failure 42 failure scenarios (continued) global clusters (continued) storage failure 42 replicated data clusters 46 application failure 46 host failure 46 network failure 46 replication link failure 46 site failure 46 storage failure 46 fire drill about 59 configuration wizard 66 running 71 service group for 66 SRDFSnap agent 61 supported configurations 60 functions action 13 clean 13 monitor 13 offline 13 online 13 open 13

G

global clusters failure scenarios 42 GrpName attribute 23

I

installing the agent AIX systems 17 Linux systems 17 Solaris systems 17 IsCompositeGroup attribute 23

Μ

Mode attribute 23 monitor function 13

0

offline function 13 online function 13 OnlineTimeout attribute setting 38 open functions 13

R

Recovery Point Objective (RPO) ComputeDRSLA attribute 28 Configuring RPO computation support 37 replicated data clusters failure scenarios 46 resource type definition SRDFSnap agent 63 RPO computation 23

S

sample configuration 29 split-brain handling in cluster 35 SplitTakeover attribute 23 SRDFSnap agent about 61 attribute definitions 63 operations 61 type definition 63 SwapRoles attribute 23 SymapiServers 23 SymHome attribute 23

Т

type definition SRDFSnap agent 63

U

uninstalling the agent AIX systems 21 Linux systems 21 Solaris systems 21

V

VCSResLock attribute 28