

InfoScale™ 9.0 Cluster Server Agent for EMC SRDF Configuration Guide - Windows

Last updated: 2025-04-13

Legal Notice

Copyright ©2025 Arctera US LLC. All rights reserved.

Arctera and the Arctera Logo are trademarks or registered trademarks of Arctera US 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 Arctera is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Arctera product or available at:

<https://www.arctera.io/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and de-compilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Arctera US 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. ARCTERA US 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 Arctera 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.

Arctera US LLC | www.arctera.io

Documentation

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

<https://sort.veritas.com/arctera/documents>

Services and Operations Readiness Tools (SORT)

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

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

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.arctera.io/support

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

Contents

Technical Support	4	
Chapter 1	Introducing the agent for EMC SRDF	7
	About the agent for EMC SRDF	7
	Supported hardware for EMC SRDF	8
	Supported software	8
	Typical EMC SRDF setup in a VCS cluster	8
	EMC SRDF agent functions	9
	About the EMC SRDF agent's online function	11
	About dynamic swap support for the EMC SRDF agent	12
	Installing the agent for EMC SRDF	13
Chapter 2	Configuring the agent for EMC SRDF	14
	Configuration concepts for the EMC SRDF agent	14
	Resource type definition for the EMC SRDF agent	14
	Attribute definitions for the SRDF agent	15
	Sample configuration for the EMC SRDF agent	19
	Before you configure the agent for EMC SRDF	20
	About cluster heartbeats	21
	About configuring system zones in replicated data clusters	22
	About preventing split-brain	23
	Configuring the agent for EMC SRDF	23
	Configuring the agent manually in a global cluster	24
	Configuring the agent manually in a replicated data cluster	25
	Setting the OnlineTimeout attribute for the SRDF resource	26
	Additional configuration considerations for the SRDF agent	27
Chapter 3	Testing VCS disaster recovery support with EMC SRDF	28
	How VCS recovers from various disasters in an HA/DR setup with EMC SRDF	28
	Failure scenarios in global clusters	29
	Failure scenarios in replicated data clusters	33

	Testing the global service group migration	38
	Testing disaster recovery after host failure	40
	Testing disaster recovery after site failure	41
	Performing failback after a node failure or an application failure	43
	Performing failback after a site failure	44
Chapter 4	Setting up fire drill	46
	About fire drills	46
	Supported fire drill configurations	47
	About the SRDFSnap agent	48
	SRDFSnap agent functions	48
	Resource type definition for the SRDFSnap agent	49
	Attribute definitions for the SRDFSnap agent	50
	About the Snapshot attributes	52
	Sample configuration for a fire drill service group	52
	Additional considerations for running a fire drill	52
	Before you configure the fire drill service group	53
	Configuring the fire drill service group	54
	About the Fire Drill wizard	54
	Verifying a successful fire drill	54

Introducing the agent for EMC SRDF

This chapter includes the following topics:

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

About the agent for EMC SRDF

The High Availability agent for EMC Symmetrix Remote Data Facility (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).

Supported hardware for EMC SRDF

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.

Supported software


For information on the software versions that the InfoScale (VCS) agent for SRDF supports, see the Arctera Services and Operations Readiness Tools (SORT) site:

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

Typical EMC SRDF setup in a VCS cluster

The following figure 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 21.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.
In a global cluster environment, you must attach all hosts in a cluster to the same 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.

Table 1-1 Agent functions

Function	Description
online	<p>This operation makes the devices writable for the application.</p> <p>If one or more devices are in the write-disabled (WD) state, the agent runs the <code>symrdf</code> command to enable read-write access to the devices.</p> <p>See “About the EMC SRDF agent’s online function” on page 11.</p> <p>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 the Synchronized state, the agent creates a lock file on the local host. The lock file indicates that the resource is online.</p> <p>The agent 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.</p> <p>See “About dynamic swap support for the EMC SRDF agent” on page 12.</p> <p>If the <code>SymcliDbFile</code> attribute is defined and if a valid <code>symapi</code> database has been created with the same path and name, the agent performs further SRDF operations accordingly. If this attribute is not defined or if its value is null, the agent uses the default <code>symapi</code> database to perform the SRDF commands.</p>
offline	<p>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.</p>
monitor	<p>Verifies that the lock file exists. If the lock file exists, the monitor entry point reports the status of the resource as online. If the lock file does not exist, the monitor entry point reports the status of the resource as offline.</p>
open	<p>Removes the lock file on the host where the entry point is called. This operation prevents potential concurrency violation if the service group fails over to another node.</p> <p>Note: The agent does not remove the lock file if the agent was started after running the following command:</p> <pre>hastop<-all -local> -force</pre>
clean	<p>Determines if it is safe to fault the resource if the online entry point fails or times out.</p>

Table 1-1 Agent functions (*continued*)

Function	Description
info	Reports the device state to the VCS interface. This entry point can be used to verify the device state and to monitor dirty track trends.
action/update	Performs a <code>symrdf update</code> action from the R2 side to merge any dirty tracks from the R2 to the R1.
action/PreSwitch	<p>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 <code>hagrp -switch</code> command.</p> <p>For this, the PreSwitch Action attribute must be set to 1. The option <code>-nopre</code> indicates that the VCS engine must switch the servicegroup regardless of the value of the PreSwitch service group attribute.</p> <p>If running the PreSwitch action fails, the failover should not occur. This minimizes the application downtime and data loss.</p> <p>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>.</p>
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.

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 26.

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: `symrdf failover`. The command enables read-write on the R2 device.

If the SwapRoles attribute is set to 1, 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.

If the SwapRoles attribute is set to 2, the agent does not check for any criteria and directly initiates the role swap between the dynamic devices.

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.

Installing the agent for EMC SRDF

During the installation of the VCS server components, the High Availability hardware replication agents are also installed.

The High Availability agent for EMC SRDF is also available in the form of an agent pack, which is released on a quarterly basis. The agent pack includes support for new applications as well as fixes and enhancements to existing agents. You can install the agent pack on an existing VCS installation.

You can download and install the latest agent pack from the Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.com/agents>.

Refer to the *Veritas InfoScale Installation and Upgrade Guide* for instructions on installing and upgrading the VCS.

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 VCS agent for EMC SRDF.

```
type SRDF (
    static keylist RegList = { GrpName }
    static int ActionTimeout = 180
    static keylist SupportedActions = { update, PreSwitch }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static int OfflineMonitorInterval = 0
    static int OnlineWaitLimit = 0
    static int RestartLimit = 1
    static i18nstr ArgList[] = { SymHome, GrpName, DevFOTime,
        AutoTakeover, SplitTakeover, Mode, IsCompositeGroup, SwapRoles,
        LinkMonitor, ConfValidate, AdvancedOpts, SymcliDbFile }
    str SymHome = "C:\\Program Files\\EMC\\SYMCLI\\bin"
```

```

    str GrpName
    int DevFOTime = 2
    boolean AutoTakeover
    boolean SplitTakeover
    str Mode
    boolean IsCompositeGroup
    int SwapRoles = 1
    temp str VCSResLock
    int LinkMonitor = 0
    int ConfValidate = 0
    str AdvancedOpts{} = { ExtendMonitor = null }
    str SymcliDbFile
)
    
```

Attribute definitions for the SRDF agent

Review the description of the agent attributes.

Required attribute for SRDF

You must assign values to required attributes.

Table 2-1 Required attribute

Attribute	Description
Name: GrpName Type-dimension: string-scalar	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 <code>IsCompositeGroup</code> to 1.

Optional attributes for SRDF

Configuring these attributes is optional.

Table 2-2 Optional attributes

Attribute	Description
<p>Name: SwapRoles</p> <p>Type-dimension: integer-scalar</p>	<p>Dynamic devices only</p> <p>Specifies whether the roles of the dynamic devices must be swapped at the time of failover or not:</p> <ul style="list-style-type: none"> ■ If set to 0, the agent does not swap the role. ■ If set to 1, the agent checks whether the devices support role swapping. If the devices have the capability, the agent swaps the roles between the RDF1 and RDF2 devices. ■ If set to 2, the agent does not check whether the devices support role swapping, and directly initiates the role swap between the RDF1 and RDF2 devices. <p>Default: 1</p>
<p>Name: IsCompositeGroup</p> <p>Type-dimension: integer-scalar</p>	<p>Specifies whether the SRDF group is a composite group:</p> <ul style="list-style-type: none"> ■ If set to 0, the agent treats it as device group. ■ If set to 1, the agent treats it as composite group. <p>Default: 0</p>
<p>Name: SymHome</p> <p>Type-dimension: string-scalar</p>	<p>Path to the bin directory that contains the Symmetrix command line interface.</p> <p>Default: C:\Program Files\EMC\SMYCLI\bin</p>
<p>Name: DevFOTime</p> <p>Type-dimension: integer-scalar</p>	<p>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 the other devices or the composite groups to fail over. If the online operation cannot be completed in the remaining time, the failover does not proceed.</p> <p>Default: 2 (seconds per device)</p>

Table 2-2 Optional attributes (*continued*)

Attribute	Description
<p>Name: AutoTakeover</p> <p>Type-dimension: integer-scalar</p>	<p>Determines whether the agent performs a <code>symrdf rw_enable</code> operation on the partitioned devices at the secondary site:</p> <ul style="list-style-type: none"> ■ If set to 1, the agent fails over the service group at the DR site even when the replication is in the Partitioned state. ■ If set to 0, the agent waits for consent from an administrator to proceed with the failover. <p>Default: 0</p> <p>Note: A Partitioned state occurs when the replication link is broken, which means that the secondary devices are not in sync with the primary devices, or that 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.</p> <p>For details, refer to the Dell EMC SRDF documentation.</p>
<p>Name: SplitTakeover</p> <p>Type-dimension: integer-scalar</p>	<p>Determines whether the agent permits a failover to the R2 devices in the Split state:</p> <ul style="list-style-type: none"> ■ If set to 0, the agent does not permit a failover to the R2 devices in the Split state. ■ If set to 1, the agent permits a failover to the R2 devices in the Split state if the devices are read-write enabled. <p>Recommendation: Set the attribute to 0 to minimize the risk of data loss in case of a failover to devices that may not be in sync.</p> <p>Default: 0</p> <p>Note: The attribute has no effect on the failover operations to a host to which R1 devices are attached.</p>

Table 2-2 Optional attributes (continued)

Attribute	Description
<p>Name: LinkMonitor</p>	<p>Determines whether the agent should check the status of the replication link while bringing the resource online:</p> <ul style="list-style-type: none"> ■ If set to 1, the agent checks the status of the replication link. If the replication is in the Synchronized or the Consistent state, the agent brings the resource online. Otherwise, the resource remains offline and results in a service group fault. ■ If set to 0, the agent brings the resource online regardless of the status of the replication link. ■ Other values of the attribute are reserved for future use by the agent. <p>Default: 0</p> <p>Note: 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.</p>
<p>Name: Mode Type-dimension: string-scalar</p>	<p>Determines which commands to use to fail over the service group to the other site.</p> <p>The values for this attribute can be Asynchronous or Synchronous.</p> <p>Default: No default value</p> <p>Note: If the value is not specified, the agent assumes that the mode is Synchronous. If the devices are set up to replicate in the Asynchronous mode, you must set this attribute to Asynchronous.</p>
<p>Name: AdvancedOpts Type-dimension: string-association</p>	<p>Determines whether the agent executes a custom script during the monitor cycle for the resource.</p> <p>Use the ExtendMonitor key with this attribute—set its value to the absolute path of the script that should be executed during the monitor cycle.</p> <p>To disable the execution of the custom script, set ExtendMonitor to null, or remove the key from the AdvancedOpts attribute.</p> <p>Example: <code>AdvancedOpts{} = { ExtendMonitor=null }</code></p>

Table 2-2 Optional attributes (continued)

Attribute	Description
Name: ConfValidate Type-dimension: integer-scalar	Specifies whether the agent should check the state of the EMC device group configuration before it brings the resource online: <ul style="list-style-type: none"> ■ If set to 1, the agent compares the number of standard devices on the remote site and on the local site. If the numbers match, the agent brings the resource online; otherwise, it faults the resource. ■ If set to 0, the agent brings the resource online regardless of the state of the EMC device group configuration. Default: 0 Example: 1
Name: SymcliDbFile Type-dimension: string-scalar	Absolute path to the custom <code>symapi</code> database file. Default: No default value—the path value of the standard <code>symapi</code> database file is considered to perform the <code>symrdf</code> operations. Note: Do not modify the value of this attribute if you do not plan to use a customized database for your storage devices.

Internal attributes

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

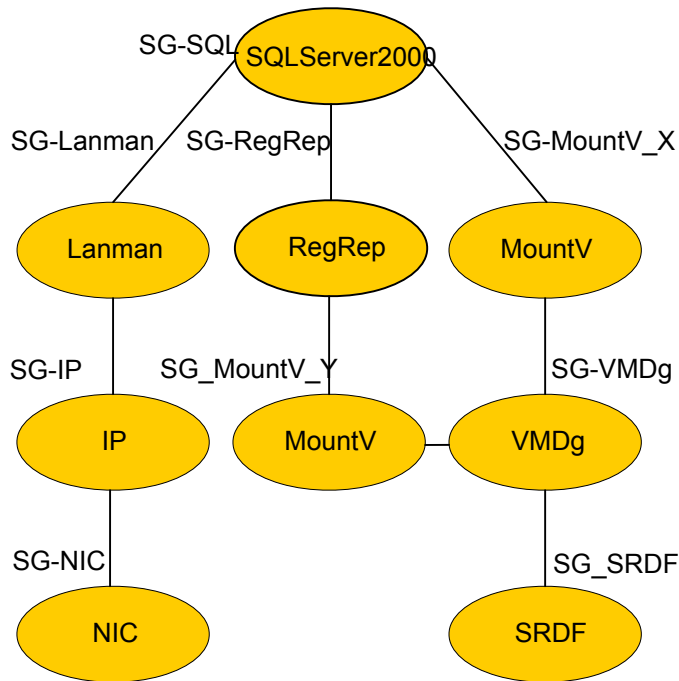
Table 2-3 Internal attributes

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

Sample configuration for the EMC SRDF agent

Figure 2-1 shows the dependency graph for a VCS service group with a resource of type SRDF. The VMDg resource depends on the SRDF resource.

Figure 2-1 Sample configuration for the SRDF agent



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:

```

SRDF SG-SRDF (
    GrpName = "SQLDg"
)
  
```

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 the cluster nodes.
- Verify the hardware setup for the agent.
 See [“Typical EMC SRDF setup in a VCS cluster”](#) on page 8.
- Make sure that the cluster has an effective heartbeat mechanism in place.
 See [“About cluster heartbeats”](#) on page 21.

See [“About preventing split-brain”](#) on page 23.

- Set up system zones in replicated data clusters.
 See [“About configuring system zones in replicated data clusters”](#) on page 22.
- 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 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 2-2 depicts a sample configuration where `hosta` and `hostb` are in one system zone, and `hostc` and `hostd` are in another system zone.

Use the `SystemZones` attribute to create these zones.

Figure 2-2 Example system zone configuration



Modify the `SystemZones` attribute using the following command:

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

The variable `grpname` represents the service group in the cluster.

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 means 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 19.

Note: You must not change the replication state of devices from primary to secondary and from secondary to primary, outside of a VCS setup. The agent for EMC 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 If the agent resource type (SRDF) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types**, and select:

```
systemdrive\Program Files\Veritas\cluster server\conf\
Sample_SRDF\SRDFTypes.cf
```

- 3 Click **Import**.
- 4 Save the configuration.
- 5 Add a resource of type SRDF at the bottom of the service group.
- 6 Configure the attributes of the SRDF resource.
- 7 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.

- 8 Change the ClusterFailOverPolicy attribute from the default, if necessary. Arctera recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
- 9 Repeat step 5 through step 8 for each service group in each cluster that uses replicated data.
- 10 Configure the Symm heartbeat on each cluster.
See [“Configuring the Symm heartbeat on each cluster”](#) on page 24.
- 11 The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Configuring the Symm heartbeat on each cluster

To configure Symm heartbeat on each cluster

- 1 From Cluster Explorer Edit menu, choose **Configure Heartbeats**.
- 2 On the Heartbeats Configuration dialog box, enter the name of the heartbeat (Symm).
- 3 Select the check box next to the name of the cluster to add it to the cluster list for the heartbeat.

- 4 Click the icon in the Configure column to open the Heartbeat Settings dialog box.
- 5 Set the value of the AYARetryLimit attribute for this heartbeat to 1 less than the value for the ICMP heartbeat.
- 6 As a first value of the Argument attribute, specify the ID of the Symmetrix array in the other cluster.
- 7 As a second value of the Argument attribute, specify the full path of SYMCLIs, using the short path name.
- 8 Click **OK**.
- 9 Symm heartbeat monitors only one array using the Symmetrix ping utility. You must configure additional heartbeats if you use devices from more than one array.

To configure additional heartbeats:

- Create a copy of <your installation directory>\cluster server\bin\hb\Symm folder using a different name under <your installation directory>\cluster server\bin\hb*, say Symm_1.
- Open the VCS Java GUI to configure Symm_1 heartbeat. The parameters are similar to Symm heartbeats.

Configuring the agent manually in a replicated data cluster

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

To configure the agent in a replicated data cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 If the agent resource type (SRDF) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types** and select:

```
systemdrive\Program Files\Veritas\Cluster Server\conf\
config\SRDFTypes.cf.
```

- 3 Click **Import**.
- 4 Save the configuration.
- 5 In each service group that uses replicated data, add a resource of type SRDF at the bottom of the service group.

- 6 Configure the attributes of the SRDF resource.
- 7 Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

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:

$$\text{OnlineTimeout} = \sum_{1}^{n_{\text{devicegroups}}} ((n_{\text{devices}} \times d_{\text{failovertime}}) + \epsilon)$$

- n_{devices} represents the number of devices in a device group.
- $d_{\text{failovertime}}$ represents the time taken to failover a device.
- $n_{\text{devicegroups}}$ 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.

```
C:\Program Files\Veritas\Cluster Server\bin\SRDF\sigma.pl
```

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.

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 26.
- 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.

Testing VCS disaster recovery support with 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 topic lists various failure scenarios and describes how VCS responds to the failures in the following DR cluster configurations.

Global clusters

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

See [“Failure scenarios in global clusters”](#) on page 29.

Replicated data clusters

When service group faults or system faults occur, the VCS failover behavior depends on the value of the `AutoFailOver` attribute of 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 33.

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

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

Table 3-1 Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes global service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the <code>ClusterFailOverPolicy</code> global service group attribute: <ul style="list-style-type: none"> ■ 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. <p>The agent write-enables the devices at the secondary site.</p> <p>For dynamic RDF devices, if the value of the <code>SwapRoles</code> attribute of the SRDF resource is 1, the agent does the following :</p> <ul style="list-style-type: none"> ■ 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. <p>See “Performing failback after a node failure or an application failure” on page 43.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the primary cluster fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ 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. <p>The agent write-enables the devices at the secondary site.</p> <p>For dynamic RDF devices, if the value of the SwapRoles attribute of the SRDF resource is 1, the agent does the following:</p> <ul style="list-style-type: none"> ■ 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. <p>See “Performing failback after a node failure or an application failure” on page 43.</p>
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>A site failure renders the devices on the array at the secondary site in the PARTITIONED state.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the cluster fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ 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. <p>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:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>symrdf failover</code> command to make the SRDF devices write-enabled. ■ 0—The agent faults the SRDF resource. <p>See “Performing failback after a site failure” on page 44.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>A replication link failure renders the SRDF devices in the PARTITIONED state. When the link is restored, the SRDF devices attain the SUSPENDED state.</p> <p>VCS response: No action.</p> <p>Agent response: No action. The agent does not monitor the replication link status and cannot detect link failures.</p> <p>After the link is restored, you must resynchronize the SRDF devices.</p> <p>To resynchronize the SRDF devices after the link is restored:</p> <ul style="list-style-type: none"> ■ 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 <code>symrdf resume</code> command. ■ After R1 and R2 devices are in sync, reestablish the mirror relationship between the BCV or target devices and the R2 devices. <p>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.</p> <p>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. Arctera recommends configuring BCV or target devices at both the sites.</p> <p>See "Typical EMC SRDF setup in a VCS cluster" on page 8.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Network failure	<p>The network connectivity and the replication link between the sites fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ VCS at each site concludes that the remote cluster has faulted. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ 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. <p>When the network (WAC and replication) connectivity is restored, you must manually resync the data.</p> <p>Note: Arctera 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.</p> <p>To resynchronize the data after the network link is restored:</p> <ul style="list-style-type: none"> ■ 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 <code>symrdf establish</code> or the <code>symrdf restore</code> command. ■ Bring the global service group online on one of the sites. <p>Agent response: Similar to the site failure.</p>

Table 3-1 Failure scenarios in a global cluster configuration with the VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Storage failure	<p>The array at the primary site fails.</p> <p>A storage failure at the primary site renders the devices on the array at the secondary site in the PARTITIONED state.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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. <p>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:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>symrdf failover</code> command to make the SRDF devices write-enabled. ■ 0—The agent faults the SRDF resource.

Failure scenarios in replicated data clusters

The following table 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 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 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. <p>The agent write enables the devices at the secondary site.</p> <p>For dynamic RDF devices, the agent does the following if the value of the SwapRoles attribute of the SRDF resource is 1:</p> <ul style="list-style-type: none"> ■ 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. <p>See “Performing failback after a node failure or an application failure” on page 43.</p>
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 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. <p>The agent write enables the devices at the secondary site.</p> <p>For dynamic RDF devices, if the value of the SwapRoles attribute of the SRDF resource is 1, the agent does the following:</p> <ul style="list-style-type: none"> ■ 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. <p>See “Performing failback after a node failure or an application failure” on page 43.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>A site failure renders the devices on the array at the secondary site in the PARTITIONED state.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 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. <p>Agent response: The agent does the following based on the AutoTakeover attribute of the SRDF resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>symrdf failover</code> command to make the SRDF devices write-enabled. ■ 0—The agent faults the SRDF resource. <p>See “Performing failback after a site failure” on page 44.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF (*continued*)

Failure	Description and VCS response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>A replication link failure renders the SRDF devices in the PARTITIONED state. When the link is restored, the SRDF devices attain the SUSPENDED state.</p> <p>VCS response: No action.</p> <p>Agent response: No action. The agent does not monitor the replication link status and cannot detect link failures.</p> <p>After the link is restored, you must resynchronize the SRDF devices.</p> <p>To resynchronize the SRDF devices after the link is restored:</p> <ol style="list-style-type: none"> 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. <p>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.</p> <p>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.</p> <p>See "Typical EMC SRDF setup in a VCS cluster" on page 8.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF (continued)

Failure	Description and VCS response
Network failure	<p>The LLT and the replication links between the sites fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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. <p>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.</p> <p>After taking the service group offline, you must manually resync the data using the <code>symrdf establish</code> or the <code>symrdf restore</code> command.</p> <p>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.</p> <p>To resynchronize the data after the network link is restored:</p> <ol style="list-style-type: none"> 1 Take the service groups offline at both the sites. 2 Manually resynchronize the data. <p>Depending on the site whose data you want to retain use the <code>symrdf establish</code> or the <code>symrdf restore</code> command.</p> 3 Bring the service group online on one of the sites. <p>Agent response: Similar to the site failure.</p>

Table 3-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF (continued)

Failure	Description and VCS response
Storage failure	<p>The array at the primary site fails.</p> <p>A storage failure at the primary site renders the devices on the array at the secondary site in the PARTITIONED state.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> ■ 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. <p>Agent response: The agent does the following based on the AutoTakeover attribute of the SRDF resource:</p> <ul style="list-style-type: none"> ■ 1—If invalid tracks do not exist, the agent issues the <code>symrdf failover</code> command to make the SRDF devices write-enabled. ■ 0—The agent does not perform failover to the secondary site.

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
- 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 29.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

Perform the procedure that is applicable to your DR configuration 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:

```
hagrps -online service_group -to sys_name
```

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

```
hagrps -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 29.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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 29.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site.

Perform the procedure that is applicable to your DR configuration.

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

- 1 Switch the global service group from the secondary site to any node at 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: <ul style="list-style-type: none">■ 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 <code>symrdf failback</code> command to resync the R1 devices and to write enable the R1 devices at the primary site. |
| For static RDF | Issues the <code>symrdf failback</code> 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 29.

See [“Failure scenarios in replicated data clusters”](#) on page 33.

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 `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.

Setting up fire drill

This chapter includes the following topics:

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

About fire drills

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

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

In clusters employing EMC 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.

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

Supported 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 52.

Gold	<p>Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.</p> <p>Arctera recommends this configuration because it does not affect production recovery.</p> <p>In the Gold configuration, VCS does the following:</p> <ul style="list-style-type: none">■ 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. <p>For non-replicated Symmetrix devices:</p> <ul style="list-style-type: none">■ You must use Veritas Volume Manager.■ You must use the Gold configuration without the option to run in the Bronze mode.
Silver	<p>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.</p> <p>If a disaster occurs while resynching data after running the fire drill, you must switch to the snapshot for recovery.</p> <p>In the Silver configuration, VCS does the following:</p> <ul style="list-style-type: none">■ 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 maintains 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.

SRDFSnap agent functions

The SRDFSnap agent performs the following functions:

online

Gold Configuration

- Takes a local snapshot of the target LUN.
- Takes the fire drill service group online by mounting the replication target LUN.
- Creates a lock file to indicate that the resource is online.

Silver Configuration

- Takes a local snapshot of the target LUN.
- Splits replication between the source and the target arrays.
- Takes the fire drill service group online by mounting the target LUN.
- Creates a lock file to indicate that the resource is online.

Bronze Configuration

- Splits replication between the source and the target arrays.
- Takes the fire drill service group online using the target array.
- Creates a lock file to indicate that the resource is online.

offline	<p>Gold Configuration</p> <ul style="list-style-type: none"> ■ Destroys the snapshot by synchronizing data between the target array and the device on which snapshot was taken. ■ Removes the lock file created by the online function. <p>Silver Configuration</p> <ul style="list-style-type: none"> ■ Resumes replication between the source and the target arrays. ■ Synchronizes data between the target array and the device on which the snapshot was taken. Destroys the snapshot of the target array after the data is synchronized. ■ Removes the lock file created by the online function. <p>Bronze Configuration</p> <ul style="list-style-type: none"> ■ Resumes the 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	<p>Monitors the change in the value of the following attributes and verifies that the new value is not invalid.</p> <ul style="list-style-type: none"> ■ CopyMode ■ SavePoolName ■ DiskGroupSnapList <p>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.</p>

Resource type definition for the SRDFSnap agent

Following is the resource type definition for the SRDFSnap agent:

```

type SRDFSnap
(
    static keylist RegList = { CopyMode, SavePoolName, DiskGroupSnapList }
    static i18nstr ArgList[] = { TargetResName, MountSnapshot, UseSnapshot,
                                RequireSnapshot, DiskGroupSnapList,
                                CopyMode, UseTgt, SavePoolName }

    i18nstr TargetResName
    i18nstr DiskGroupSnapList
  
```

```
boolean MountSnapshot = 1
boolean UseSnapshot = 1
boolean RequireSnapshot = 1
i18nstr SavePoolName
temp str Responsibility
temp i18nstr FDFFile
str CopyMode = mirror
boolean UseTgt = 0
)
```

Attribute definitions for the SRDFSnap agent

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

TargetResName	<p>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.</p> <p>For example, in a typical database setup, you might replicate data files and redo logs, but you may choose to avoid replicating temporary tablespaces. The temporary tablespace must still exist at the DR site and may be part of its own disk group.</p> <p>Type-dimension: string-scalar</p>
MountSnapshot	<p>Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
UseSnapshot	<p>Specifies whether the SRDFSnap resource takes a local snapshot of the target array. Set this attribute to 1.</p> <p>Type-Dimension: integer-scalar</p> <p>See "About the Snapshot attributes" on page 52.</p>

RequireSnapshot	<p>Specifies whether the SRDFSnap resource must take a snapshot before coming online.</p> <p>Set this attribute to 1 if you want the resource to come online only after it succeeds in taking a snapshot.</p> <p>Type-Dimension: integer-scalar</p> <p>Note: Set this attribute to 1 only if UseSnapshot is set to 1.</p>
DiskGroupSnapList	<p>This is an optional attribute that lists the original disk group names and the fire drill disk group names.</p> <p>Type-dimension: string-scalar</p>
CopyMode	<p>Indicates the array snapshot technology to be used.</p> <ul style="list-style-type: none">■ mirror indicates TimeFinder/Mirror■ clone indicates TimeFinder/Clone■ snap indicates TimeFinder/Snap <p>Type-dimension: string-scalar</p> <p>Default: 0</p>
UseTgt	<p>Indicates whether the agent should use target devices or BCVs in the device group.</p> <p>0 indicates BCV devices, 1 indicates target devices.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p>
Responsibility	<p>Do not modify. For internal use only.</p> <p>Used by the agent to keep track of resynchronizing snapshots.</p> <p>Type-Dimension: temporary string</p>
FDFile	<p>Do not modify. For internal use only.</p> <p>Used by the agent to store the absolute pathname to the file with the latest fire drill report on the local system.</p> <p>Type-Dimension: temporary string</p>

About the Snapshot attributes

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

Table 4-1 Snapshot attribute values for fire drill configurations

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

If the snapshot operation fails, set the RequireSnapshot attribute to 0 to enable a Gold or a Silver configuration to run in the Bronze mode.

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
)
```

Additional considerations for running a fire drill

Follow these guidelines for fire drills in a Windows environment:

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

- You must configure BCV or target device pairs (for SRDF) before running the wizard.

Before you configure the fire drill service group

Before you configure the fire drill service group, ensure that the following prerequisites 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.
- 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.
- 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

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

About the Fire Drill wizard

Storage Foundation and High Availability Solutions for Windows (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Volume Replicator (VVR) or that uses EMC SRDF hardware replication.

Verifying a successful fire drill

Run the fire drill routine periodically to verify that 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 where the application is not running.

If the fire drill service group comes online, it validates your DR configuration. The service group at the primary site can fail over to the secondary site in the event of an actual failure or a 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 service group offline after its functioning has been verified. Failing to take this service group offline may cause failures in your environment. For example, if the application service group fails over to the node where the fire drill service group is hosted, both the service groups may fault owing to resource conflicts.