Veritas™ Cluster Server Agent for EMC SRDF Installation and Configuration Guide

Windows Server 2003, Windows Server 2008

5.1 Service Pack 2

September 2010

symantec™

Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide

Agent Version: 5.1.SP2.0

Document version: 5.1.SP2.0.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Documentation feedback

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

sfha_docs@symantec.com

Support agreement resources

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

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

Contents

Chapter 1

Introducing the Veritas agent for EMC SRDF

This chapter includes the following topics:

About the agent for EMC SRDF

The Veritas agent for 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.

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

The agent supports SRDF device groups and consistency groups in sync and async modes. The agent also supports dynamic SRDF (role swap).

Note: The agent does not support semi-synchronous and Adaptive Copy.

Supported software for EMC SRDF

The EMC SRDF agent supports Storage Foundation and High Availability Solutions 5.0 and 5.1 for Windows.

Supported hardware for EMC SRDF

The SRDF agent supports Solutions Enabler (SE) 6.4 or later and corresponding array microcode levels. Please refer to the EMC hardware compatibility list for specific information.

Typical EMC SRDF setup in a VCS cluster

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

Figure 1-1 Typical clustering setup for the agent

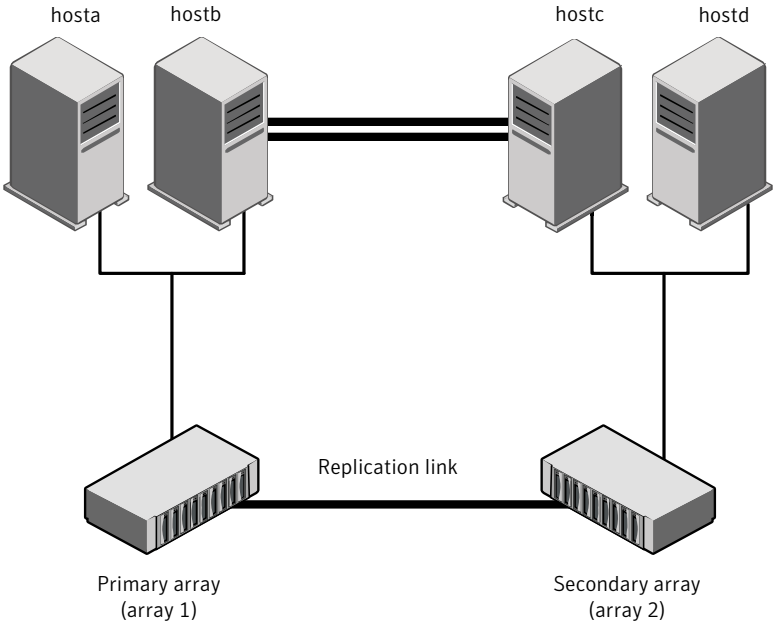

VCS clusters using SRDF for replication uses the following hardware infrastructure:

- 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 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 23.
- 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:

online	If the state of all local devices is read-write enabled (RW), the agent creates a lock file on the local host. The lock file indicates that the resource is online. This operation makes the devices writable for the application. If one or more devices are in the write-disabled (WD) state, the agent runs a `symrdf` command to enable read-write access to the devices. See "About the EMC SRDF agent's online function" on page 12.
offline	Removes the lock file on the local host. The agent does not run any SRDF commands because taking the resource offline is not indicative of the intention to give up the devices.
monitor	Verifies that the lock file exists. If the lock file exists, the monitor 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 13. 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 that 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 entry point can be used to verify the device state and to monitor dirty track trends.
action/update	Performs a `symrdf update` from the R2 side to merge any dirty tracks from the R2 to the R1.
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 is read-write enabled (RW), the agent creates a lock file on the local host to indicate that the resource is online.

If one or more devices are in the write-disabled (WD) state, the agent runs a `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 a `symrdf` command only after the devices transition to the R1 UPDATED state.
- For R2 devices in the SYNCINPROG state, the agent runs a `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 28.

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.

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 2

Installing and removing the agent for EMC SRDF

This chapter includes the following topics:

Before you install the agent for EMC SRDF

Set up your cluster. For information about installing and configuring VCS, see the *Veritas Cluster Server Installation Guide*.

Set up replication and the required hardware infrastructure.

See "Typical EMC SRDF setup in a VCS cluster" on page 10.

Installing the agent for 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.

To install the VCS agent for

1. Log on to any node in the cluster.

   Ensure that the logged on user has the domain administrative privileges.

2. Download the complete agent pack tarball from FileConnect site:

   https://fileconnect.symantec.com/

   Alternatively,

   Download the individual agent tarball from the Symantec Veritas Operations Services (VOS) site:

   https://vos.symantec.com/home

3. Uncompress the file to a temporary location.

4. If you downloaded the complete Agent Pack tarball, navigate to the directory containing the package for the platform running in your environment.

5. Double-click **vrtsvcsagentname.msi.**

   Follow the instructions that the install program provides, to complete the installation of Veritas Cluster Server agent.

Removing the agent for SRDF

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

To remove the agent SRDF

1. Open the Windows Control Panel and click **Add or Remove Programs.**
2. Click the VSFW HA Server Components entry and click **Remove.**
3. Review the Welcome page and click **Next.**
4. In the Option Selection dialog box, select the SRDF agent and click **Next.**
5. The installer validates the system for uninstallation.

   If a system is rejected, the Comments column displays the cause of rejection. Highlight the system to view detailed information about the failure in the Details box. Resolve the error, highlight the node in the selected systems list, and click **Validate Again.**

   After all the systems are accepted, click **Next.**
6. Review the summary of your selections and click **Uninstall.**

7 The installer displays the status of uninstallation.

8 After the uninstallation is complete, review the report and click **Next.**

9 Click **Finish.**

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

Chapter 3

Configuring the agent for EMC SRDF

This chapter includes the following topics:

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 str ArgList[] = { SymHome, GrpName, DevFOTime,
    AutoTakeover, SplitTakeover }
        static int NumThreads = 1
        static int ActionTimeout = 180
        static int OfflineMonitorInterval = 0
        static int MonitorInterval = 300
        static int RestartLimit = 1
        int SwapRoles = 1
        static keylist SupportedActions = { update }
        NameRule = resource.GrpName
        str SymHome = "C:\\Program Files\\EMC\\SYMCLI\\bin"
        str GrpName
```

```
    int DevFOTime = 2
    int AutoTakeover = 1
    int SplitTakeover = 0
    temp str VCSResLock
)
```

Attribute definitions for the SRDF agent

Review the description of the agent attributes.

Required attributes

You must assign values to required attributes.

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.

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-scaler

  Default: 1

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-scaler

  Default: 0

SymHome
: Path to the bin directory that contains the Symmetrix command line interface.

  Type-dimension: string-scalar

  Default is C:\Program Files\EMC\SMYCLI\bin.

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 is 2 seconds per device.

AutoTakeover — A flag that determines whether the agent performs a `symrdf rw_enable` operation on the partitioned devices at the secondary site.

Type-dimension: integer-scalar

Default is 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 is 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.

Internal attributes

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

VCSResLock — The agent uses the VCSResLock attribute to guarantee serialized management in case of a parallel application.

Type-dimension: temporary string

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 VMDg resource depends on the SRDF resource.

Figure 3-1 Sample configuration for the SRDF agent

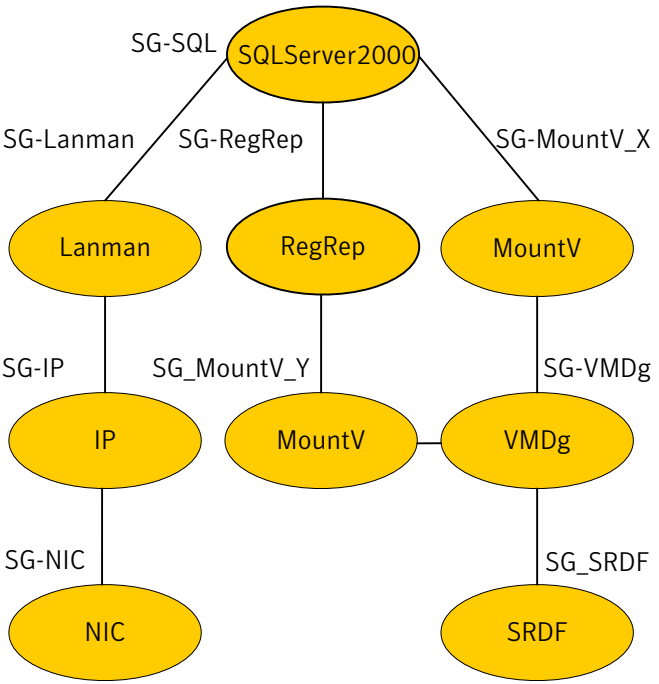

A resource of type SRDF may be configured as follows in main.cf:

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

Before you configure the agent for EMC SRDF

Before you configure the agent, review the following information:

- Review the configuration concepts, which describe the agent's type definition and attributes.
  See "Configuration concepts for the EMC SRDF agent" on page 19.
- 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 10.
- Make sure that the cluster has an effective heartbeat mechanism in place.
  See "About cluster heartbeats" on page 23.
  See "About preventing split-brain" on page 25.
- Set up system zones in replicated data clusters.
  See "About configuring system zones in replicated data clusters" on page 24.
- 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, see the *Veritas 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 3-2 depicts a sample configuration where hosta and hostb are in one system zone, and hostc and hostd are in another system zone.

Use the SystemZones attribute to create these zones.

Figure 3-2 Example system zone configuration

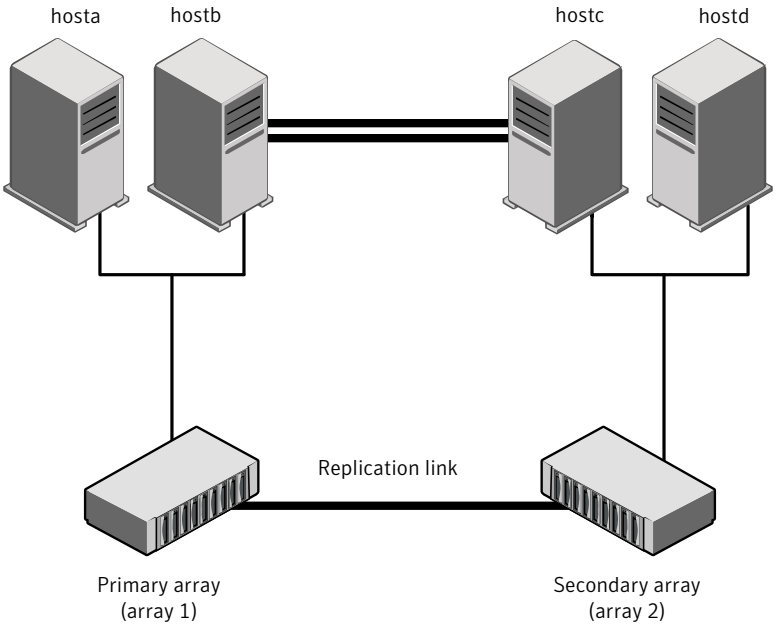

Modify the SystemZones attribute using the following command:

```
hagrp -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 vice-versa. In this case, the application faults because its underlying volumes become write-disabled, causing the service group to fault. VCS tries to fail it over to another host, causing the same consequence in the reverse direction. This phenomenon continues until the group comes online on the final node. You can avoid this situation by setting up your infrastructure such that loss of heartbeat links also mean the loss of replication links.

Configuring the agent for EMC SRDF

You can adapt most clustered applications to 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 22.

Note: You must not change the replication state of devices primary to secondary and viceversa, 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.

Configuring the agent manually in a global cluster

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

To configure the agent in a global cluster

1. Start Cluster Manager and log on to the cluster.
2. If the agent resource type (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.

   See the *Veritas Cluster Server Administartor's Guide* for more information.
8. Change the ClusterFailOverPolicy attribute from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
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.
    - From Cluster Explorer Edit menu, choose **Configure Heartbeats.**
    - On the Heartbeats Configuration dialog box, enter the name of the heartbeat (Symm).
    - Select the check box next to the name of the cluster to add it to the cluster list for the heartbeat.
    - Click the icon in the Configure column to open the Heartbeat Settings dialog box.
    - Specify the Symmetrix ID of the array in the other cluster as the first parameter of the Arguments attribute. Specify SymHome as the second argument.
    - Set the value of the AYARetryLimit attribute for this heartbeat to 1 less than the value for the ICMP heartbeat.
    - Click **OK**.

- 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. Follow 10 for more information on configuring Symm heartbeats in order to add values.

Note: The Disaster Recovery wizard configures the required settings for the SRDF resource in the VCS application service group. Optional settings are left in the default state. The wizard creates a complete disaster recovery setup using theSRDF replication and validates the replication setup. For information on using the Disaster Recovery wizard, see the Solutions guides chapters on disaster recovery.

Configuring the agent manually in a replicated data cluster

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

To configure the agent in a replicated data cluster

1. Start Cluster Manager and log on to the cluster.
2. 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:

   ```
   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. Note that some attributes must be localized to reflect values for the hosts that are attached to different arrays.
7. Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

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:

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

   - $n_{devices}$ 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.

  ```
  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 28.
- 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.

Chapter 4

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

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 VCS agent for EMC SRDF ensures safe and exclusive access to the configured EMC SRDF devices.

See "Failure scenarios in global clusters" on page 32.

Replicated data clusters	When service group or system faults occur, VCS failover behavior 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 36.

See the for more information on the DR configurations and the global service group attributes.

Failure scenarios in global clusters

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.

See the for more information on the DR configurations and the global service group attributes.

Table 4-1 Failure scenarios in a global 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 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:
	■ 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 44.

Table 4-1 Failure scenarios in a global cluster configuration with VCS 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:
	■ 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 44.

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 45.

Table 4-1 Failure scenarios in a global 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. Agent response: No action. The VCS agent for EMC SRDF does not monitor the replication link status and cannot detect link failures. 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 off the BCV 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 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 devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Symantec recommends configuring BCV devices at both the sites. See "Typical EMC SRDF setup in a VCS cluster" on page 10.

Table 4-1 Failure scenarios in a global cluster configuration with VCS 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:** Symantec 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 resync the data. Depending on the site whose data you want to retain use the `symrdf establish` or the `symrdf restore` commands. ■ Bring the global service group online on on one of the sites. Agent response: Similar to the site failure
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.

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 4-2 Failure 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 global service group:
	  ■ 1—VCS automatically brings the faulted global service group online at the secondary site.
	  ■ 2—You must bring the global service group online at the secondary site.
	Agent response:
	■ 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 44.

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:
	■ 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 44.

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for EMC SRDF *(continued)*

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 45.

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. Agent response: No action. The VCS agent for EMC SRDF does not monitor the replication link status and cannot detect link failures. 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 off the BCV 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 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 devices and if disaster occurs when resync is in progress, then the data at the secondary site becomes inconsistent. Symantec recommends configuring BCV devices at both the sites. See "Typical EMC SRDF setup in a VCS cluster" on page 10.

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` commands. **Note:** Symantec 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 resync the data. Depending on the site whose data you want to retain use the `symrdf establish` or the `symrdf restore` commands. **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 faults the SRDF resource.

Testing the global service group migration

After you configure the VCS 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 agent for EMC SRDF behave in response to the failure.

See "Failure scenarios in global clusters" on page 32.

See "Failure scenarios in replicated data clusters" on page 36.

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 agent for EMC SRDF behave in response to the failure.

See "Failure scenarios in global clusters" on page 32.

See "Failure scenarios in replicated data clusters" on page 36.

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 32.

See "Failure scenarios in replicated data clusters" on page 36.

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 32.

See "Failure scenarios in replicated data clusters" on page 36.

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.

Chapter 5

Setting up fire drill

This chapter includes the following topics:

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 drills for storage devices that are managed using Veritas Volume Manager, which is a component of Veritas Storage Foundation.

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

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.

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.
- 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

- Suspends 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	Gold Configuration ■ 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. Silver Configuration ■ 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. Bronze Configuration ■ 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.
action	For internal use.

Resource type definition for the SRDFSnap agent

Following is the resource type definition for the SRDFSnap agent:

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

Attribute definitions for the SRDFSnap agent

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

TargetResName	Name of the resource managing the LUNs that you want to take snapshot of. Set this attribute to the name of the 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. Type-dimension: string-scalar
UseSnapshot	Specifies whether the SRDFSnap resource takes a local snapshot of the target array. Set this attribute to 1 Type-Dimension: integer-scalar See "About the Snapshot attributes" on page 51.
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: integer-scalar **Note:** Set this attribute to 1 only if UseSnapshot is set to 1.
MountSnapshot	Specifies whether the resource uses the snapshot to bring the service group online. Set this attribute to 1. Type-Dimension: integer-scalar **Note:** Set this attribute to 1 only if UseSnapshot is set to 1.
Responsibility	Do not modify. For internal use only. Used by the agent to keep track of resynchonizing snapshots. Type-Dimension: temporary string

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
DiskGroupSnapList	This is an optional attribute that lists the original disk group names and the fire drill disk group names. Type-dimension: string-scaler

About the Snapshot attributes

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

Table 5-1 lists the snapshot attribute values for fire drill configurations:

Table 5-1 Snapshot attribute values for fire drill configurations

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

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

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 oradg_fd {
    TargetResName = "oradf_rdf"
    UseSnapshot = 1
    RequireSnapshot = 0
    MountSnapshot = 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 Veritas Storage Foundation HA 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 pairs (for SRDF) before running the wizard.

Before you configure the fire drill service group

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

- Make sure the application service group is configured with a SRDF resource.
- Make sure the infrastructure to take snapshots is properly configured between the source and target arrays. This process involves associating BCVs and synchronizing them with the source.
- 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, BCV's must be associated with the RDF2 device group and fully established with the devices.
- When you take snapshots of non-replicated devices, create a EMC 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 have the same BCVs 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 firedrill 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 firedrill using space-optimized snapshots.

- Make sure that the SAVE pool as specified by SavePoolName attribute exists prior to running firedrill using space-optimized snapshots.
- Make sure that the copy sessions are not created for the device or composite group prior to running firedrill with space-optimized snapshots.
- Make sure that the SRDF mode of replication is set to synchronous prior to running firedrill using space-optimized snapshots. This is because EMC does not support creation of TimeFinder/Snap copy sessions for RDF2 device, if the SRDF mode of replication is set to asynchronous.

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

Veritas Storage Foundation High Availability 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 Veritas Volume Replicator (VVR) or that uses EMC SRDF hardware replication.

For more information on using the Fire Drill wizard, see the Solutions guides chapters on testing fault readiness and running a fire drill.

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 that does not have the application running. Verify that the fire drill service group comes online.

   This action validates your disaster recovery configuration. The production service group can fail over to the secondary site in the event of an actual failure (disaster) at the primary site.

2. If the fire drill service group does not come online, review the VCS engine log for more information.

3. Take the fire drill offline after its functioning has been validated.

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

Index

R

S

T

U

V