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

AIX, HP-UX, Linux, Solaris

5.0

March 2008

symantec™

Veritas Cluster Server Agent for EMC SRDF Installation and Configuration guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Agent version: 5.0.03.0

Legal Notice

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

http://www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder.

AIX is a registered trademark of IBM Corporation.

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.

Linux is a registered trademark of Linus Torvalds.

Solaris is a trademark of Sun Microsystems, Inc.

Technical support

For technical assistance, visit

http://www.symantec.com/enterprise/support/assistance_care.jsp

and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1

Introducing the Veritas agent for 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 also supports parallel applications, such as Veritas Storage Foundation for Oracle RAC.

The agent supports SRDF in the synchronous and asynchronous modes; the agent does not support semi-synchronous and Adaptive Copy. The agent does not require special configuration for SRDF/A support; the agent detects SRDF/A backed devices and manages their failover accordingly.

The agent also supports dynamic SRDF (role swap). If all devices in a given device group are configured for dynamic SRDF, the agent attempts a role swap.

See the following Technical Support TechNote for the latest updates or software issues for this agent:

http://seer.entsupport.symantec.com/docs/282004.htm

What's new in this release

The Veritas Cluster Server for EMC SRDF includes the following new or enhanced features:

- On AIX, the SRDF agent supports fire drill for storage devices that are managed with LVM.

Supported software and hardware

The agent for EMC SRDF 5.0.03.0 supports the following software versions:

Veritas Cluster Server	■ VCS 5.0 ■ VCS 5.0 MP1
Veritas SF for Oracle RAC	■ SF Oracle RAC 5.0 ■ SF Oracle RAC 5.0 MP1
Veritas Volume Manager	■ VxVM 5.0 ■ VxVM 5.0 MP1 On HP-UX and AIX, Symantec recommends using VxVM 5.0 MP1.

Operating systems	The agent supports the following operating systems:
	■ AIX 5.2 and 5.3 on pSeries
	■ HP-UX 11i v2 on PA and IA
	■ Linux on x86, Intel Xeon, AMD Opteron
	  ■ Red Hat Enterprise Linux 4 Update3
	  ■ SUSE Linux Enterprise Server 9 with SP3
	■ Solaris 2.8, 2.9, and 2.10 on SPARC
	Solaris 2.10 on x64 platform
	See the product's Release Notes for more details on the supported architectures and the operating systems.

The agent supports SYMCLI versions that EMC recommends for the firmware on the array.

The agent supports SRDF on all microcode levels on all EMC Symmetrix arrays.

This support only exists if the host, HBA, and array combination is in EMC's hardware compatibility list.

You must obtain an SRDF Consistency Groups license for taking consistent snapshots with the SRDF agent. You must also create a consistency group on all nodes in the cluster.

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

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

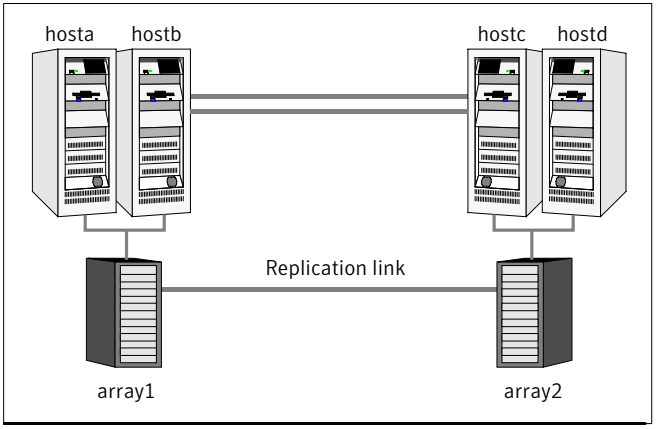

Clustering in a SRDF environment typically consists of the following hardware infrastructure:

- The primary array (array1) has one or more R1 hosts. A Fibre Channel or SCSI directly attaches these hosts to the EMC Symmetrix array that contains the SRDF R1 devices.
- The secondary array (array2) has one or more R2 hosts. A Fibre Channel or SCSI directly attaches these hosts 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 hosts and arrays must be at a significant distance to survive a disaster that may occur at the R1 side.

- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.
  See "About cluster heartbeats" on page 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.
- In parallel applications like Veritas Storage Foundation for Oracle RAC, all hosts that are attached to the same array must be part of the same GAB membership. Veritas Storage Foundation for Oracle RAC is supported with SRDF only in a global cluster environment and not in a replicated data cluster environment.

EMC SRDF agent operations

The VCS agent for EMC 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 operation" on page 11.
offline	Removes the lock file on the device. 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	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. The agent checks if a management operation was in progress when the online thread timed out. If the operation was killed, the devices are left in an unusable state.
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	Performs a `symrdf update` from the R2 side to merge any dirty tracks from the R2 to the R1.

About the EMC SRDF agent's online operation

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 30.

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 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 SYMCLI version is 5.4 or later.
- 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 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 9.

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. These instructions assume that you have already installed VCS.

To install the agent on AIX systems

1 Determine the device access name of the disc drive.

```
# cd /dev
# lsdev -C -c cdrom
```

The output resembles:

```
cd0 Available 10-60-00-4,0 16 Bit SCSI Multimedia CD-ROM Drive
```

In this example, the CD device access name is cd0.

2 Insert the disc into the system's drive.

3 Mount the disc.

```
# mkdir -p /cdrom
# mount -V cdrfs -o ro /dev/cd0 /cdrom
```

4 Navigate to the location of the agent packages:

```
# cd /cdrom/aix/replication/srdf_agent/version/pkgs
```

The variable *version* represents the version of the agent.

5 Add the filesets for the software.

```
# installp -ac -d VRTSvcse.rte.bff VRTSvcse
```

To install the agent on HP-UX systems

1 Insert the disc into the system's drive.

2 Create a mount point directory. For example, /cdrom. The directory must have read-write permissions.

3 Determine the block device file for the disc drive.

```
# ioscan -fnC disk
```

For example, the listing may indicate the block device is /dev/dsk/c1t2d0.

4 Start the Portable File System (PFS).

```
# nohup pfs_mountd &
# nohup pfsd &
```

5 Mount the disc.

```
# /usr/sbin/pfs_mount -t rrip /dev/dsk/c#t#d# /cdrom
```

The variable */c#t#d#* represents the location of the drive.

6 Install the agent software. Type one of the following commands depending on the operating system on the node.

HP-UX (PA)

```
# swinstall -s /cdrom/hpux/replication\
/srdf_agent/version/PA/depot VRTSvcse
```

HP-UX (IA)

```
# swinstall -s /cdrom/hpux/replication\
/srdf_agent/version/IA/depot VRTSvcse
```

The variable *version* represents the version of the agent.

To install the agent on Linux systems

1 Log in as superuser.

2 Insert the disc into the system's drive.

3 Mount the disc, if the disc does not automatically mount.

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

4 Navigate to the /mnt directory.

```
# cd /mnt/cdrom
```

5 Navigate to the location of the agent package.

```
# cd linux/platform/replication/srdf_agent/version/rpms/
```

The variable *platform* represents the Linux distribution and architecture.

The following are the platform values:

- redhatlinux, redhatlinuxX86_64
- suselinux, suselinuxX86_64

The variable *version* represents the version of the agent.

6 Install the agent software:

```
# rpm -ivh agentrpm
```

The variable *agentrpm* represents the agent package in the rpms directory.

To install the agent on Solaris systems

1 Insert the disc into the system's drive.

```
# cd /cdrom/cdrom0
```

2 Navigate to the location of the agent package.

```
# cd solaris/sparc/platform/replication/srdf_agent
/version/pkgs/
```

The variable *platform* represents the Solaris distribution and architecture.

The following are the platform values:

- x64
- sparc

The variable *version* represents the version of the agent.

3 Install the agent binaries.

```
# pkgadd -d . VRTSvcse
# pkgadd -d . VRTScsecw
# pkgadd -d . VRTScsfdw
```

Configuring LVM on AIX

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

To configure LVM on AIX

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

Upgrading the agent for SRDF

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

To upgrade the agent software

1. Save the VCS configuration and stop the VCS engine.

   ```
   # haconf -dump -makero
   # hastop -all -force
   ```

2. Remove the agent from the node.

   See "Removing the agent for SRDF" on page 18.

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

   See "Installing the agent for SRDF" on page 13.

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

   Configure the new attributes, if applicable.

8. Verify the configuration

   ```
   # hacf -verify config
   ```

9 Start VCS on local node first.

10 Start VCS on other nodes.

Removing the agent for SRDF

Before you attempt to remove the agent, make sure the application service group is not online. You must remove the agent from each node in the cluster.

To remove the agent from an AIX cluster

- Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# installp -u VRTSvcse.rte
```

To remove the agent from an HP-UX cluster

- Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# swremove VRTSvcse
```

To remove the agent from a Linux cluster

- Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# rpm -e VRTSvcse
```

To remove the agent from a Solaris cluster

- Type the following command on each node to remove the agent. Answer prompts accordingly:

```
# pkgrm VRTSvcse
# pkgrm VRTScsecw
# pkgrm VRTSfdw
```

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
         static keylist SupportedActions = { update }
         NameRule = resource.GrpName
         str SymHome = "/usr/symcli"
         str GrpName
         int DevFOTime = 2
```

```
        int AutoTakeover = 1
        int SplitTakeover = 1
        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 that the agent manages. Specify the name of a device group. Do not specify the name of a composite group.

Type-dimension: string-scalar

Optional attributes

Configuring these attributes is optional.

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

Type-dimension: string-scalar

Default is /usr/symcli.

DevFOTime — Average time in seconds that is required for each device in the 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 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 read-write enable on partitioned devices in the write-disabled state during a failover.

Type-dimension: integer-scalar

Default is 1, which means that the agent performs a read-write enable if devices are consistent.

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 1.

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 DiskGroup resource depends on the SRDF resource.

Figure 3-1 Sample configuration for the SRDF agent

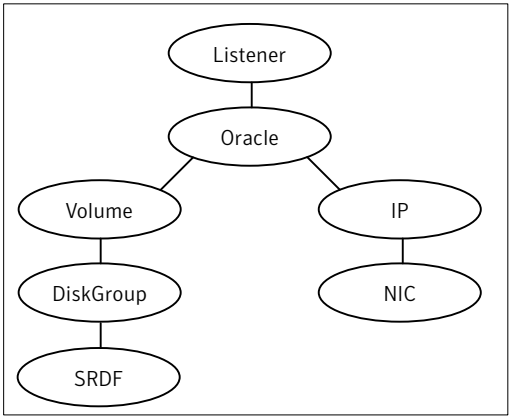

A resource of type SRDF may be configured as follows in main.cf:

```
SRDF oradf_rdf (
        GrpName = "oracle_grp")
```

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 30.
- In global clusters, the value of the AYARetryLimit for the Symm heartbeat must be shorter than the ICMP retry limit. This setting allows VCS to detect an array failure first and does not confuse a site failure with an all host failure.
- The agent supports importing and deporting a Veritas Volume Manager diskgroup when failing over Real Application Clusters across replicating arrays. Failing to do so can leave disk groups in the imported state on hosts where the storage is read-only. In this situation, any attempted write operations to the disk group are rejected, causing the disk group to be disabled. You must deport and reimport the disk group to enable it.

Before you configure the agent for 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 9.
- 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 24.
- Set up system zones in replicated data clusters.
  See "About configuring system zones in replicated data clusters" on page 23.
- 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.

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

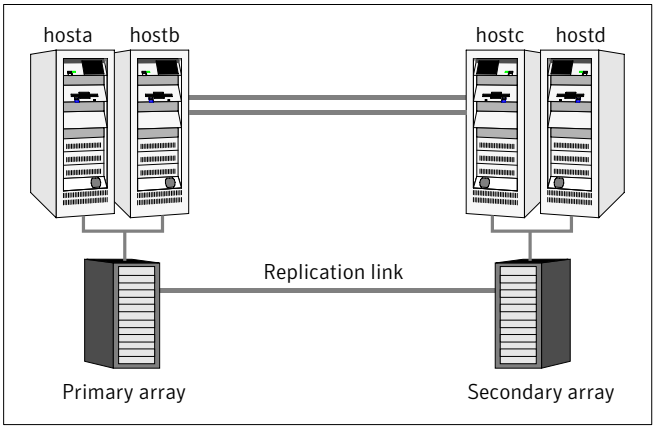

Global clusters do not require system zones because failover occurs on a remote cluster if all local targets have been exhausted.

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 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.

Configuring the agent using the wizard

Use the wizard to configure the agent in an application service group.

Note: The wizard is supported only on the Solaris operating system. The wizard does not support configuring the SplitTakeover attribute; you must configure the attribute manually.

To configure the agent using the wizard

1. Run the wizard on a system that is attached to the EMC Symmetrix array.

   Verify that the command line package (SYMCLI) for the EMC Symmetrix array is installed on the system.

2. Set the DISPLAY variable and start wizard as `root`.

   ```
   # hawizard srdf
   ```

3. Read the information on the Welcome screen and click **Next**.
4. In the Wizard Options dialog box, select the application service group to which you want to add an SRDF resource.

   The wizard displays service groups having disk group resources; it does not display service groups having SRDF resources.

5. In the Select Device Group dialog box, specify the device group from the EMC Symmetrix array where you want to add the SRDF resource.

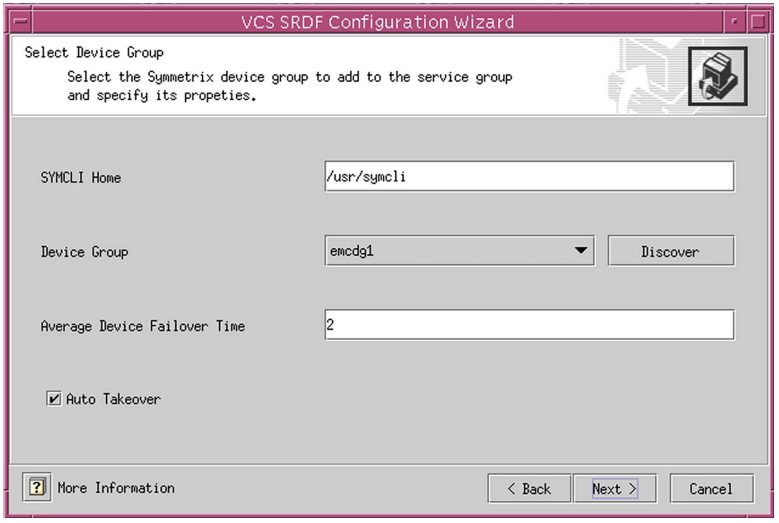

- SYMCLI Home: The path where the CLI package for the EMC Symmetrix array is installed. The default location is /usr/symcli.
- Device Group: Select the device group to be monitored. If the device group does not appear in the list, click **Discover**.
- Average Device Failover Time: Average time in seconds for each device group in the service group to fail over. Default is 2 seconds per device.
- Auto Takeover: Select if you want the SRDF resource to perform a read-write enable on partitioned devices in the write-disabled state during a failover.

6 Click **Next**.

7 Select the replicated diskgroups and click **Next**.

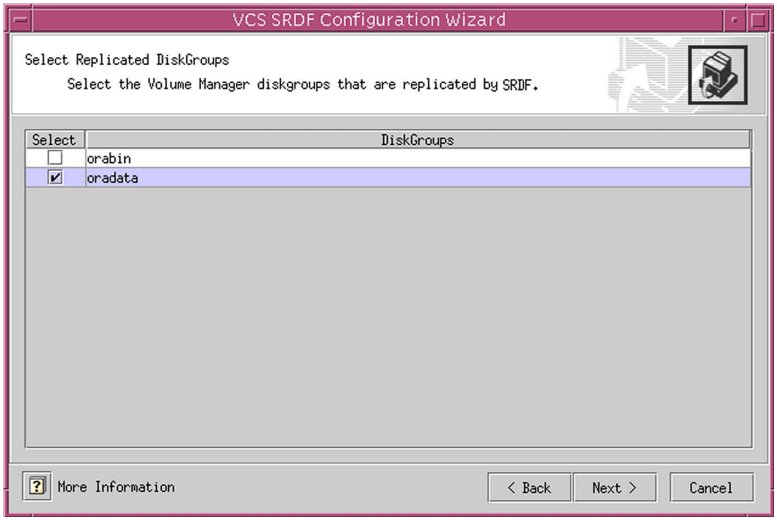

If you add the SRDF resource in a service group that is configured in a replicated data cluster, proceed to the next step. Otherwise, proceed to step 10.

8 In the System Zones Selection dialog box, specify the systems for each zone in a replicated data cluster.

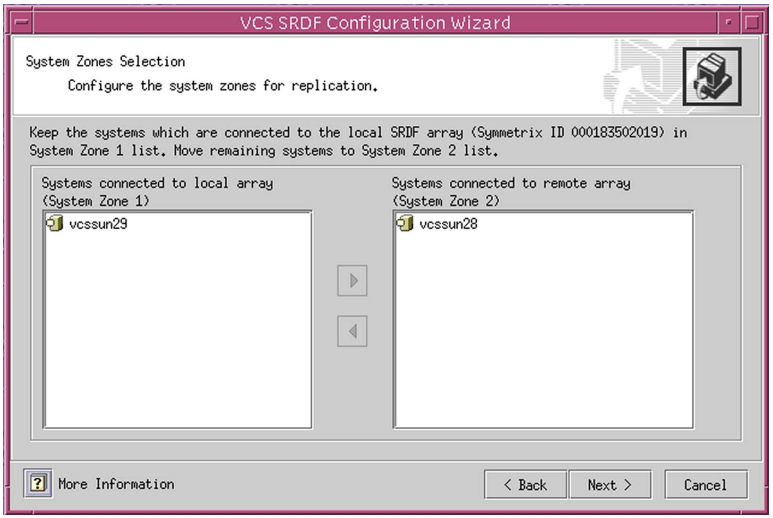

If you configured SystemZones in the application service group, verify the configuration. Use the arrows to move systems to their respective zones.

9 Click **Next**.

10 In the Service Group Summary dialog box, review the service group configuration and optionally change the name of the SRDF resource.

11 To change the name of the SRDF resource, select the resource name and either click it or press the F2 key. Press Enter after editing the resource name. To cancel editing a resource name, press Esc.

12 Click **Finish**.

The wizard starts running commands to add the SRDF resource to the service group. Various messages indicate the status of these commands.

13 In the Completing the SRDF Configuration Wizard dialog box, select the check box to bring the service group online on the local system.

14 Click **Close**.

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:

   /etc/VRTSvcs/conf/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 group, configure the service group using the Global Group Configuration Wizard.

   See the *Veritas Cluster Server User's Guide* for more information.
8. To configure the agent to manage the volumes that Veritas Storage Foundation for Oracle RAC uses:
   - Configure the SupportedActions attribute for the CVMVolDg resource
   - Add the following keys to the list: import, deport.

   Note that SupportedActions is a resource type attribute and defines a list of action tokens for the resource.
9. Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.
10. Repeat step 5 through step 9 for each service group in each cluster that uses replicated data.
11. 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 as the value of the Arguments attribute the Symmetrix ID of the array in the other cluster. Set the value of the AYARetryLimit attribute

for this heartbeat to 1 less than the value for the ICMP heartbeat. Specify SymmHome as the second argument with a value of 1.

- Click **OK**.

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:

   `/etc/VRTSvcs/conf/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:

$$\textbf{OnlineTimeout} = ((n_{devices} \times d_{failovertime}) + \varepsilon)$$

   - $n_{devices}$ 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 script

- Run the perl script to get recommendations for VCS attribute values.

  ```
  /opt/VRTSvcs/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.

Chapter 4

Managing and testing clustering support for EMC SRDF

This chapter includes the following topics:

Typical test setup for the EMC SRDF agent

A typical test environment includes the following characteristics:

- Two hosts (hosta and hostb) are attached to the R1 EMC Symmetrixarray.
- Two hosts (hostc and hostd) are attached to the R2 EMC Symmetrix array.
- The application runs on hosta and devices in the local array are read-write enabled in the SYNCHRONIZED state.
- You may add an optional SRDF link heartbeat.
- A replicated data cluster has two dedicated heartbeat links; a global cluster has one network heartbeat and an optional SRDF replication link heartbeat.

Figure 4-1 depicts a typical test environment.

Figure 4-1 Typical test setup

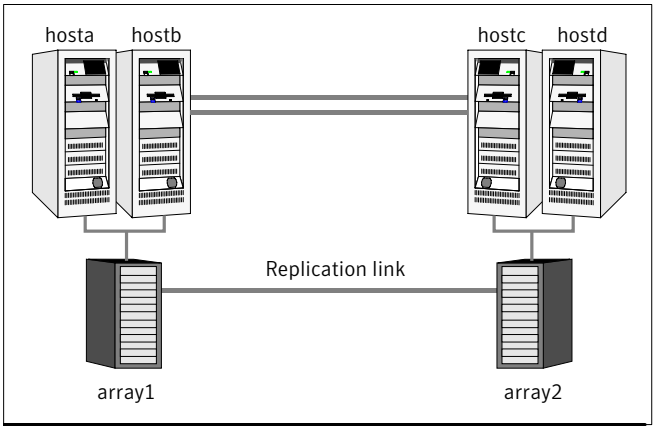

Testing service group migration

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

To perform the service group migration test

1 Migrate the service group to a host that is attached to the same array.

Do the following:

- In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- Click **Switch To**, and click the system that is attached to the same array (hostb) from the menu.
  The service group comes online on hostb and local volumes remain in the RW/SYNCHRONIZED state.

2 Migrate the service group to a host that is attached to a different array.

Do the following:

- In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- Click **Switch To**, and click the system that is attached to the another array (hostc) from the menu.
  The service group comes online on hostc and volumes there transition to the RW/FAILED OVER state.

- Accumulate dirty tracks on the R2 side and update them back on the R1:

  ```
  hares -action srdf_res_name update -sys hostc
  ```

  The variable *srdf_res_name* represents the name of the SRDF resource.

3 After the devices transition to R1 UPDATED state, migrate the service group back to its original host.

  Do the following:

  - In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
  - Click **Switch To**, and click the system on which the group was initially online (hosta).
    The group comes online on hosta. The devices return to the RW/SYNCINPROG state at the array that is attached to hosta and hostb, and then eventually transition to the SYNCHRONIZED state.

Testing host failure

In this scenario, the host where the application runs is lost. Eventually all the hosts in the system zone or cluster are lost.

To perform the host failure test

1 Halt the host where the application runs (hosta).

  The service group fails over to hostb and devices are in the RW/SYNCHRONIZED state.

2 Halt or shut down hostb.

  In a replicated data cluster, the group fails over to hostc or hostd depending on the FailOverPolicy in the cluster.

  In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.

  In both environments, the devices transition to the RW/FAILED OVER state and start on the target host.

3 Reboot the two hosts that were shut down.

4 Switch the service group to its original host when VCS starts.

  Do the following:

  - In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

- Click **Switch To**, and click the system on which the service group was initially online (hosta).

  The service group comes online on hosta and devices transition to the SYNCINPROG state and then to the SYNCHRONIZED state.

Performing a disaster test

Test how robust your cluster is in case of a disaster.

To perform a disaster test

1. Shut down all hosts on the source side and shut down the source array.

   If you can not shut down the R1 Symmetrix, disconnect the ESCON link between the two arrays and simultaneously shut down the hosts. This action mimics a disaster scenario from the point of view of the R2 side.

2. In a replicated data cluster, the service group fails over to hostc or hostd in the following conditions:

   - All devices were originally in the SYNCHRONIZED state.
   - No synchronization was in progress at the time of disaster.

3. In a global cluster, the administrator is notified of the failure. The administrator can then initiate the failover.

Performing the failback test

You can set up your cluster for a failback test.

The failback test verifies the application can fail back to its original host after a failover to a remote site.

To perform a failback test

1. Reconnect the ESCON cable and reboot the original R1 hosts.
2. Take the service group offline.

   If you run this test in a replicated data cluster, type the following command from any host:

   ```
   hagrp -offline grpname -any
   ```

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

3 After the service group goes offline, manually resynchronize the devices, which you can do only if you write-disable both sides. Type:

```
symrdf -g device_group restore
```

The variable `device_group` represents the name of the RDF device group at the R2 side. The `restore` command determines which tracks to merge between the R1 and R2 arrays and initiates the resynchronization. The operation of this command write disables both sides; use this command only when a brief downtime is acceptable.

4 Bring the service group online at the R1 side. Type:

```
hagrp -online grpname -sys hosta
```

The devices synchronize, and the environment state becomes the same as when the test began.

Failure scenarios for EMC SRDF

Review the failure scenarios and agent behavior in response to failure.

Site disaster

In a total site failure, all hosts and the array are completely disabled, either temporarily or permanently.

In a replicated data cluster, site failure is detected the same way as a total host failure, that is, the loss of all LLT heartbeats.

In a global cluster, VCS detects site failure by the loss of both the ICMP and Symm heartbeats. Make sure that a site failure is not confused with an all-host failure. Set the AYARetryLimit for the Symm heartbeat to be shorter than the ICMP retry limit. With such a setting, the failure of the Symmetrix array is detected first.

A total disaster renders the devices on the surviving array in the PARTITIONED state. If the AutoTakeover attribute is set to its default value of 1, the online entry point runs the `symrdf_rw` command. If the attribute is set to 0, no takeover occurs and the online entry point times out and faults.

The online entry point detects whether any synchronization was in progress when the source array was lost. Since the target devices are inconsistent until the synchronization completes, the agent does not write-enable the devices, but it

times out and faults. You must restore consistent data from a snapshot or tape backup.

All host or all application failure

Even if both arrays are operational, the service group fails over in the following conditions:

- All hosts on the R1 side are disabled.
- The application cannot start successfully on any R1 host.

In replicated data cluster environments, the failover can be automatic, whereas in global cluster environments failover requires user confirmation by default.

VCS serializes `symrdf` commands to ensure that SRDF does not lock out a command while another command is running.

Make sure that SRDF agent's entry points do not time out. Set the OnlineTimeout and RestartLimit attributes for the SRDF resource to restart automatically if the agent entry points are timed out.

Replication link failure

SRDF detects link failures, monitors changed tracks on devices, and resynchronizes R2 devices if the R1 was active at the time of the link failure.

Before the SRDF takes any action, it waits for the synchronization to complete in the following situations:

- The two arrays are healthy and the link that failed is restored.
- A failover is initiated while synchronization is in progress.

After the synchronization completes, the SRDF runs the `symrdf failover` command.

If the agent times out before the synchronization completes, the resource faults.

The R2 devices are rendered inconsistent and unusable in the following conditions:

- A failover is initiated due to a disaster at the R1 site, and
- A synchronization was in progress

In this case, even if the AutoTakeover attribute of the agent is set to 1, the agent does not enable read-write access to the devices. Instead, the agent faults. You must restore consistent data to these devices, either from BCV or from a tape backup. Then, you must enable read-write access to the devices manually before they can be used.

If the AutoTakeover attribute is set to 0, the agent does not attempt a `symrdf rw_enable`, but it times out and faults. If you write-enable the devices manually, the agent can come online after it is cleared.

Split-brain in a SRDF environment

When split-brain occurs in a replicated database cluster, VCS assumes a total disaster because the R1 hosts and array are unreachable. VCS attempts to start the application on the secondary site. Once the heartbeats are restored, VCS stops the applications on one side and restarts the VCS engine (HAD). This action eliminates concurrency violation of the same group being online at two places simultaneously.

You must resynchronize the volumes manually using the `symrdf merge` or `symrdf restore` commands.

In a global cluster, you can confirm the failure before failing over the service groups. You can check with the site administrator to identify the cause of the failure. If a fail over mistakenly occurs, the situation is similar to the replicated data cluster case. However, when the heartbeat is restored, VCS does not stop HAD at either site. VCS forces you to choose which group to take offline. You must resynchronize the data manually.

Chapter 5

Setting up a 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 5.0 MP1 RP2, which is a component of Veritas Storage Foundation.

Additionally, on AIX the agent also supports fire drills for storage devices that are managed with LVM.

Fire drill configurations

VCS supports the following fire drill configurations for the agent:

Gold

Runs the fire drill on a snapshot of the target array. The replicated device keeps receiving writes from the primary.

Symantec recommends this configuration because it does not affect production recovery.

In the Gold configuration, VCS does the following:

- Takes a snapshot of the BCV device on the target array.
- Modifies the disk group name in the snapshot.
- Brings the fire drill service group online using the snapshot data.

For non-replicated devices:

- You must use Veritas Volume Manager. Additionally, on AIX, you can also use LVM.
- You must use the Gold configuration without the option to run in the Bronze mode.

Silver

VCS takes a snapshot, but does not run the fire drill on the snapshot data. VCS breaks replication and runs the fire drill on the replicated target device. If a disaster occurs while resynching data after running the fire drill, you must switch to the snapshot for recovery.

In the Silver configuration, VCS does the following:

- Takes a snapshot of the BCV 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. 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: ■ Suspends replication. ■ Modifies the disk group name while importing. ■ Brings the fire drill service group online using the data on the target array.

Note: On AIX, the SRDFSnap agent supports LVM in the Gold configuration. It does not support LVM in the Silver or Bronze configurations.

About the SRDFSnap agent

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

The SRDFSnap agent performs the following operations:

Online	Gold Configuration ■ Takes a local snapshot of the target LUN. ■ Takes the fire drill service group online by mounting the snapshot. ■ For AIX LVM, the agent runs the LVM command `recreatevg` to create the firedrill volume group. ■ 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. ■ For AIX LVM, the agent runs the LVM command `exportvg` to remove the firedrill volume group entries from the ODM. ■ Removes the lock file created by the online operation. 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 operation. 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 operation.
Action	For internal use.

Resource type definition for the SRDFSnap agent

Following is the resource type definition for the SRDFSnap agent:

```
type SRDFSnap (
    static str ArgList[] = { TargetResName, MountSnapshot,
    UseSnapshot, RequireSnapshot }
    static keylist RegList = { MountSnapshot, UseSnapshot }
    static int NumThreads = 1
    str TargetResName
    int MountSnapshot
    int UseSnapshot
    int RequireSnapshot
    temp str Responsibility
    temp str FDFile
    )
```

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 for Gold and Silver configurations. For Bronze, set this attribute to 0. Type-Dimension: integer-scalar See "About the Snapshot attributes" on page 46.

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. Set this attribute to 0 if you do want the resource to come online even if it fails to take a snapshot. Setting this attribute to 0 creates the Bronze configuration. 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 for Gold configuration. For Silver and Bronze configurations, set the attribute to 0. 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 locate the latest fire drill report. Type-Dimension: temporary string

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.

The following configuration creates a Gold fire drill configuration, but allows VCS to run a Bronze fire drill if the snapshot does not complete successfully.

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
    }
```

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. If you use LVM on AIX, the LVM volume group must have the same name as the device group.
- For non-replicated devices:
  - You must use the Gold configuration without the option to run in the Bronze mode. Set the RequireSnapshot attribute to 1.

Configuring the fire drill service group

On the secondary site, the initial steps create a fire drill service group that closely follows the configuration of the original application service group. The fire drill service group uses a point-in-time copy of the production data. Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over and come online at the secondary site, should the need arise.

See "Sample configuration for a fire drill service group" on page 47.

The fire drill service group uses a point-in-time snapshot of the application data.

You can create the fire drill service group using one of the following methods:

- Cluster Manager (Java Console)
  See "Creating the fire drill service group using Cluster Manager (Java Console)" on page 48.
- Fire Drill Configuration wizard
  The wizard is supported only on Solaris systems.
  See "Configuring the fire drill service group using the wizard" on page 50.

Creating the fire drill service group using Cluster Manager (Java Console)

This section describes how to use Cluster Manager (Java Console) to create the fire drill service group. After creating the fire drill service group, you must set the failover attribute to false so that the fire drill service group does not fail over to another node during a test.

To create the fire drill service group

1. Open the Veritas Cluster Manager (Java Console).
2. Log on to the cluster and click **OK**.
3. Click the **Service Group** tab in the left pane and click the **Resources** tab in the right pane.
4. Right-click the cluster in the left pane and click **Add Service Group**.
5. In the Add Service Group dialog box, provide information about the new service group.
   - In Service Group name, enter a name for the fire drill service group
   - Select systems from the Available Systems box and click the arrows to add them to the Systems for Service Group box.
   - Click **OK.**

To disable the AutoFailOver attribute

1. Click the **Service Group** tab in the left pane and select the fire drill service group.
2. Click the **Properties** tab in the right pane.
3. Click the **Show all attributes** button.
4. Double-click the **AutoFailOver** attribute.
5. In the Edit Attribute dialog box, clear the **AutoFailOver** check box.
6. Click **OK** to close the Edit Attribute dialog box.
7. Click the **Save and Close Configuration** icon in the tool bar.

Adding resources to the fire drill service group

Add resources to the new fire drill service group to recreate key aspects of the application service group.

To add resources to the service group

1. In Cluster Explorer, click the **Service Group** tab in the left pane, click the application service group and click the **Resources** tab in the right pane.
2. Right-click the resource at the top of the tree, select **Copy > Self and Child Nodes**.
3. In the left pane, click the fire drill service group.
4. Right-click the right pane, and click **Paste**.
5. In the Name Clashes dialog box, specify a way for the resource names to be modified, for example, insert an FD_ prefix. Click **Apply**.
6. Click **OK**.

Configuring resources for fire drill service group

Edit the resources in the fire drill service group so they work properly with the duplicated data. The attributes must be modified to reflect the configuration at the remote site. Bringing the service group online without modifying resource attributes is likely to result in a cluster fault and interruption in service.

To configure the fire drill service group

1. In Cluster Explorer, click the **Service Group** tab in the left pane.
2. Click the fire drill service group in the left pane and click the **Resources** tab in the right pane.
3. Right-click the SRDF resource and click **Delete**.

4 Add a resource of type SRDFSnap and configure its attributes.

5 Right-click the resource to be edited and click **View > Properties View**. If a resource to be edited does not appear in the pane, click **Show All Attributes**.

6 Edit attributes to reflect the configuration at the remote site. For example, change the Mount resources so that they point to the volumes that are used in the fire drill service group.

Configuring the fire drill service group using the wizard

Review the procedure to create the fire drill service group using the configuration wizard. Note that the configuration wizard is supported only on Solaris systems.

To configure the fire drill service group

1 Set the DISPLAY variable and start the Fire Drill Configuration wizard as a superuser.

```
# hawizard firedrill
```

2 Read the information on the Welcome screen and click **Next**.

3 In the Wizard Options dialog box, select the application service group for which you want to configure a fire drill service group.

The wizard does not display service groups that do not have SRDF resources.

4 Verify the information in the Device Group Details dialog box and click **Next**.

5 In the Snapshot Methods dialog box, choose the **Gold**, **Silver**, or **Bronze** configuration option for the fire drill service group.

See "Fire drill configurations " on page 42.

6 Select the **Use Bronze method if snapshot fails** check box if you want the fire drill service group to come online even if the resource fails to take a snapshot. This check box is enabled only if you choose the Gold or Silver configuration.

7 Click **Next**.

8 In the Snapshot Details dialog box, the wizard informs whether the device group on the target array has synchronized BCV devices to take a snapshot. If the devices are synchronized, click **Next**.

If the devices are not synchronized, quit the wizard, synchronize data between the target array and the BCV device, and rerun the wizard.

9 In the Service Group Summary dialog box, review the service group configuration and change the resource names if wanted.

10 To edit a resource name, select the resource name and either click it or press the F2 key. Press Enter after editing each resource name. To cancel editing a resource name, press Esc.

11 Click **Finish**.

The wizard starts running commands to create the fire drill service group. Various messages indicate the status of these commands.

12 In the Completing the Fire Drill Configuration Wizard dialog box, select the check box to bring the service group online on the local system.

13 Click **Close**.

14 In Linux clusters, verify that the StartVolumes attribute for each DiskGroup type resource in the fire drill group is set to 1. If not, modify the resource to set the value to 1.

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.

You can also view the fire drill log, which is located at /tmp/fd-servicegroup.

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

S

T

U

V