

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

AIX, HP-UX, Linux, Solaris

5.0

March 2008



Veritas Cluster Server Agent for EMC SRDF/Star Installation and Configuration Guide

Copyright © 2008 Symantec Corporation. All rights reserved.

Agent version: 5.0.02.0

Symantec, the Symantec logo, Veritas, are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Symantec Corporation and its licensors, if any.

THIS DOCUMENTATION IS PROVIDED “AS IS” AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID, SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
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.

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 VCS agent for EMC SRDF/Star	
	About the VCS agent for EMC SRDF/Star	8
	Supported software and hardware	8
	Typical EMC SRDF/Star setup using VCS clusters	9
	Agent entry points	11
	About the agent's online entry point	15
	Scenario I: The local site is the current workload site	15
	Scenario II: The local site is not the current workload site	16
Chapter 2	Installing and removing the agent	
	Before you install the agent	20
	Installing the agent	20
	Removing the agent	22
Chapter 3	Configuring the agent for EMC SRDF/Star	
	Configuration concepts	24
	SRDFStar resource type definition	24
	Attribute definitions for the SRDFStar agent	25
	Sample configuration	30
	Before you configure the agent	30
	About cluster heartbeats	31
	About preventing split-brain	31
	Configuring the agent	31
	Additional configuration considerations	32
	Configuring the agent for VCS global clusters	32
Chapter 4	Managing and testing clustering support for EMC SRDF/Star	
	Typical test setup	36
	Test scenarios	36
	Testing service group migration	36
	Testing host failure	37
	Testing disaster recovery	37
	Testing failback	38

Failure scenarios	38
Workload hosts and storage failure	39
Workload hosts failure	39
Workload storage failure	39
Workload and sync target total site failure	40
Workload and async target total site failure	40
Replication link failure	40

Index	41
-------------	----

Introducing the VCS agent for EMC SRDF/Star

- About the VCS agent for EMC SRDF/Star
- Supported software and hardware
- Typical EMC SRDF/Star setup using VCS clusters
- Agent entry points

About the VCS agent for EMC SRDF/Star

The VCS enterprise agent for EMC SRDF/Star provides failover support and recovery in environments that use the SRDF/Star solution to protect their data. SRDF/Star employs a three-point triangular topology to provide data protection and failure recovery among three geographically separated sites.

The SRDF/Star solution mirrors production data on a Symmetrix array at a local site, called the workload site, to Symmetrix arrays at two remote sites. SRDF/Star technology replicates data synchronously to one remote site, called the sync target, and asynchronously to the other remote site, called the async target. The sync target is in the same region as the workload site; the async target is at a distant location.

SRDF/Star protects against a workload site failure by performing incremental recovery and data replication among the three sites. Depending upon the failure scenario, either of the remote sites can become the new production workload site.

The agent monitors and manages data recovery and replication for a specific composite group that you configure as part of the SRDF/Star configuration.

Before bringing the agent online, ensure the following:

- That SRDF/Star is configured such that STAR can be enabled from the sites
- That the Workload site can be failed over from one site to the other

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

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

Supported software and hardware

The agent supports VCS 5.0 and Solutions Enabler version 6.4.1 and later. Symantec recommends that you use the latest Solutions Enabler version available.

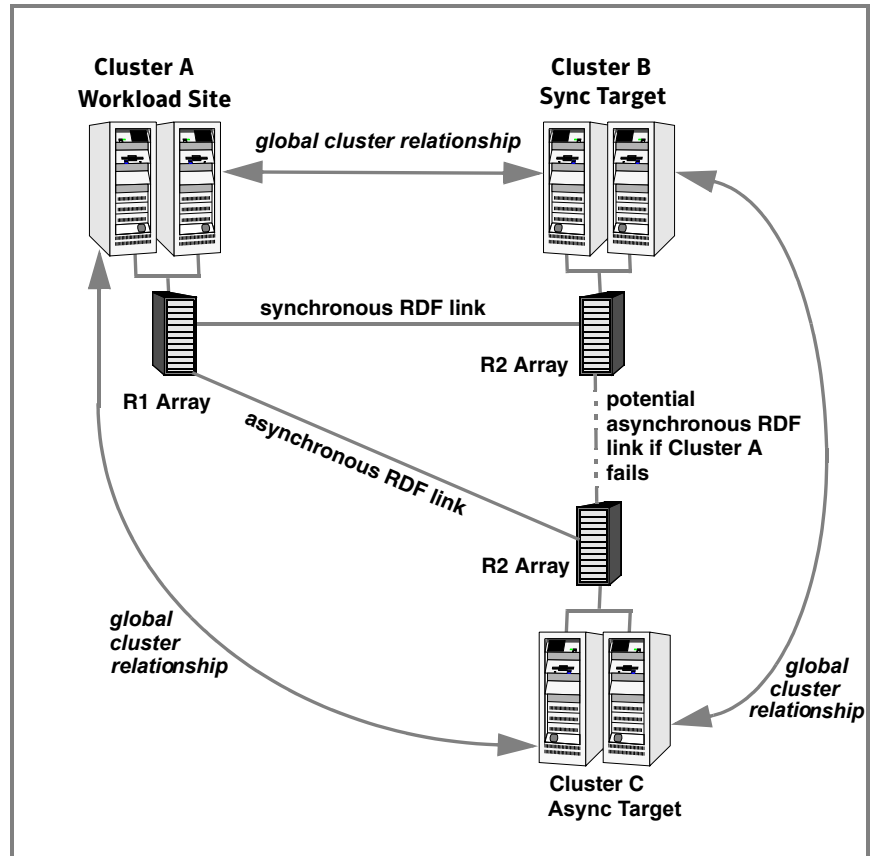
The agent supports SRDF/Star on all microcode levels on all Symmetrix arrays that support SRDF/Star, provided that the host/HBA/array combination is in the EMC hardware compatibility list.

Note: The agent supports the configuration in which the STAR composite group has multiple RDF (RA) groups configured for either link (sync or async link). The agent also supports the configuration that has meta devices.

Typical EMC SRDF/Star setup using VCS clusters

The following figure shows a typical cluster setup in an EMC SRDF/Star configuration.

Figure 1-1 Typical cluster setup in SRDF/Star



Clustering in an SRDF/Star environment typically consists of the following infrastructure:

The array at the workload site	<p>The workload site consists of one or more workload hosts directly attached to a Symmetrix array. The array contains R1 devices that are part of the SRDF/Star configuration.</p> <p>The R1 devices must also be in an SRDF/S relationship with the regional sync target and in a concurrent SRDF/A relationship with the distant async target.</p>
The array at the sync target	<p>The sync target consists of one or more sync target hosts directly attached to a Symmetrix array containing R2 devices that are part of the SRDF/Star configuration. The sync target array is paired with the workload array through a synchronous RDF link. The sync target hosts and array must be a sufficient distance away to survive a local disaster that may occur at the workload side, yet still be close enough to maintain synchronous replication.</p>
The array at the async target	<p>The async target consists of one or more async target hosts directly attached to a Symmetrix array containing R2 devices that are part of the SRDF/Star configuration. The async target array is paired with the workload array through an asynchronous RDF link. The async target hosts and array must be a sufficient distance away to survive a regional disaster that may occur across the workload and sync target sides.</p>
A global cluster configuration	<p>All three clusters are members of a VCS global cluster relationship. The sync target hosts and the async target hosts are paired with the workload hosts in a VCS global cluster relationship. In VCS parlance, each cluster has two remote clusters. The cluster (preference) list in the cluster configuration would look like the following for each cluster:</p> <ul style="list-style-type: none">■ Workload site cluster A ClusterList=[Cluster_A, Cluster_B, Cluster_C]■ Sync target site cluster B ClusterList=[Cluster_A, Cluster_B, Cluster_C]■ Async target site cluster C ClusterList=[Cluster_A, Cluster_B, Cluster_C] <p>Note: In a global cluster environment, you must attach all hosts in a cluster to the same Symmetrix array.</p>

Network heartbeats

You must configure network heartbeats among the three data centers to monitor their health.

See “[About cluster heartbeats](#)” on page 31.

Agent entry points

The VCS enterprise agent for EMC SRDF/Star (SRDFStar agent) facilitates workload site transitions.

The SRDFStar agent uses SYMCLI to do the following:

- Verify the momentary state of the target sites and STAR (PROTECTED, CONNECTED, DISCONNECTED, HALTED, TRIPPED)
- Switch the workload site from one site to the other
- Provide a mechanism for manual interventions via action entry points

The following table describes the agent entry points:

open	<p>Checks if the lock file exists in the agent directory. The existence of the lock file indicates that a resource of type SRDFStar (the SRDFStar resource) is online. If open detects the lock file, it waits until at least one of the resource's parent resources is probed and then does one of the following:</p> <ul style="list-style-type: none">■ If the parent resource is ONLINE, the agent probably called open in response to being restarted. The agent can restart whenever it or HAD has been forcibly killed. In this case, open does not remove the lock file.■ If the parent resource is OFFLINE, then the agent is being started in response to HAD being stopped gracefully and then restarted. In this case, open removes the lock file.
online	<p>For example, if the agent calls open and detects that the lock file is present and that the disk group resource above it is ONLINE, the SRDFStar resource was probably already ONLINE. In this case, open leaves the lock file in place. If the agent detects that the disk group is OFFLINE, then the host rebooted and the application failed over, allowing no chance to delete the lock file. In this case, open removes the lock file and the agent reports OFFLINE because the group is probably ONLINE elsewhere.</p> <p>Determines if the local site is the workload site and checks the state of Star. If the local site is the workload site and the state of Star is PROTECTED, the agent creates a lock file on the workload host. The lock file indicates that the SRDFStar resource is ONLINE and that the workload site is STAR-PROTECTED. If the local site is the workload site but state of STAR is not PROTECTED, then the online entry point brings the workload site into STAR-PROTECTED state and creates a lock file on the workload host. If online determines that the local site is not the workload site, it tries to convert the local site to the workload site. During this conversion, it tries to change the local site from either the sync or async target into the workload site. If successful, the online entry point creates the lock file on the new workload host.</p>

monitor	Verifies that the lock file exists. If the lock file exists, the agent reports the status of the SRDFStar resource as ONLINE. If the lock file does not exist, the agent reports the status of the resource as OFFLINE.
offline	Checks the value of HaltOnOffline attribute. If the value is 1 and the target sites are in the CONNECTED/PROTECTED state, the offline entry point issues a symstar halt command. This command halts STAR. The offline entry point also removes the lock file that is created by the online entry point.
clean	Removes the lock file that is created by the online entry point.
info	Sets the StarState attribute to the Star state as displayed in the output of the symstar query command. Reports the state of the sync and async target sites and the state of the SRDF/Star configuration. Also sends an SMTP notification (if configured to do so) whenever the value of StarState changes to a different value.

action

Note: Invoking some of the action entry points while the application is running can cause disruptions. For example, invoking the halt action while the application is running makes the underlying devices write-disabled and causes the application to crash.

Note: The timeout for action entry points is specified by the value of the ActionTimeout attribute. The action entry points extend this timeout value to 1 day before running the symstar commands. This behavior avoids the termination of any running symstar commands.

The SRDFStar agent has the following action entry points. Actions initiate symstar commands to perform various operations. For entry points that perform operations on specified sites, the site name is passed to the entry point as arguments.

The following are the action entry points:

- connect
Connects to a specified remote site
- disconnect
Disconnects from a specified remote site
- protect
Protects a specified remote site
- unprotect
Unprotects a specified remote site
- enable
Enables Star protection
- disable
Disables Star protection
- reset
Moves a specified site from the PATHFAIL state to the DISCONNECTED state
- cleanup
Performs cleanup on the specified site after a link failure
- halt
Brings the target sites into the HALTED state
- halt-reset
Brings the target sites into the DISCONNECTED state from the HALTED state

About the agent's online entry point

The online entry point determines the current workload site and then acts as described in these scenarios.

Scenario I: The local site is the current workload site

This scenario has two subcases. One is when the local site name is the same as that of either the workload site or the sync target site. The other is when the local site name is the same as that of the async target site.

The local site name same as workload or sync target

In this scenario, the local site is the current workload site and the local site name is the same as either the workload site or the sync target site (as specified in options file as part of the SRDF/Star configuration).

Under these conditions, the online entry point takes the following actions:

- The entry point checks the state of STAR. If STAR is in the PROTECTED state, it creates the lock file and the SRDFStar resource comes ONLINE. If STAR is not in the PROTECTED state, the online entry point determines the state of the remote sites as specified by the values of the syncSite and asyncSite attributes.

At this point, the remote sites can be in one of the following states:

- HALTED
- HALTFAIL
- PATHFAIL
- ISOLATED
- CONNECTED
- DISCONNECTED

The online entry point then takes one of the following actions:

- If one or both of the remote sites cannot transition to the PROTECTED state, the entry point checks the value of the ForceOnline attribute. If ForceOnline is set to 0, then the SRDFStar resource does not go ONLINE. If ForceOnline is set to 1, the entry point creates the lock file and the SRDFStar resource goes ONLINE.
- If both of the syncSite and asyncSite attributes are specified and both the remote sites transition to the PROTECTED state, the entry point checks the value of OnlineOnState attribute. If OnlineOnState is set to PROTECTED, the entry point issues the command to move Star to the PROTECTED state in the background. The

entry point then creates the lock file and the SRDFStar resource goes ONLINE.

If `OnlineOnState` is set to `STAR_PROTECTED`, the entry point issues the command to move Star into the `PROTECTED` state in the foreground. That is, the SRDFStar resource goes ONLINE only after the command has finished. If the command is not successful, the entry point checks the value of the `ForceOnline` attribute. If `ForceOnline` is set to 0, the SRDFStar resource does not go ONLINE. If `ForceOnline` is set to 1, the entry point creates the lock file and the SRDFStar resource goes ONLINE.

- If only one of the `syncSite` or `asyncSite` attributes is specified and that site transitions to the `PROTECTED` state successfully, the entry point creates the lock file and the SRDFStar resource goes ONLINE. Otherwise, the SRDFStar resource goes ONLINE only if `ForceOnline` is set to 1.
- If neither of the `syncSite` or `asyncSite` attributes is specified, the SRDFStar resource remains OFFLINE.

Local site name same as async target

In this scenario, the local site is the current workload site and the local site name is the same as the async target site (as specified in options file as part of the SRDF/Star configuration). The `syncSite` attribute is not specified.

The entry point attempts to transition only the site specified by the `asyncSite` attribute to the `PROTECTED` state. The entry point also attempts to move the second remote site to the `CONNECTED` state. The entry point does this so that, in the future, it is possible to switch from the local site by halting STAR. In such a case, the `OnlineOnState` attribute is irrelevant because both remote sites cannot be moved to the `PROTECTED` state and STAR protection cannot be enabled. If the `asyncSite` does not transition to the `PROTECTED` state, then the SRDFStar resource goes online only if `ForceOnline` is set to 1.

Scenario II: The local site is not the current workload site

In this scenario, the local site is not the current workload site. The online entry point attempts to make the local site the current workload site by doing the following:

- The entry point determines the state of STAR. If STAR is in a `TRIPPED` state, the entry point checks the value of the attribute `TripTakeover`. If `TripTakeover` is set to 0, the SRDFStar resource does not go ONLINE. If `TripTakeover` is set to 1, the entry point makes the local site the current workload site. The entry point also makes the local site the current workload site if Star is in the `HALTED` or `PROTECTED` state.

- The entry point checks the value of `KeepDataFrom`, `KeepDataOn`, and `UseMostRecentData` attributes as it makes the local site the current workload site. It checks these values to determine which site's data it uses to synchronize the local site.
 - If `UseMostRecentData` is set to 1, then the entry point determines which site has most recent data using `symstar` query command output and uses that output to synchronize the local site.
 - If `UseMostRecentData` is set to 0 and `KeepDataOn` is set to `SWITCH`, then `KeepDataFrom` attribute value is used to determine which site's data is to be used to synchronize the local site.

After making the local site the current workload site, the entry point tries to transition the remote sites and Star to the `PROTECTED` state as described in the first scenario.

See “[Scenario I: The local site is the current workload site](#)” on page 15.

Installing and removing the agent

- Before you install the agent
- Installing the agent
- Removing the agent

Before you install the agent

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/Star setup using VCS clusters](#)” on page 9.

Installing the agent

You must install the agent on each node in each cluster. The following 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/srdfstar_agent/version/pkgs
```

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

- 5 Add the filesets for the software.

```
installp -ac -d VRTSvcsst.rte.bff VRTSvcsst
```

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\  
/srdfstar_agent/version/PA/depot VRTSvcsst
```

```
HP-UX (IA) swinstall -s /cdrom/hpux/replication\  
/srdfstar_agent/version/IA/depot VRTSvcsst
```

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 Change to the `/mnt` directory.

```
cd /mnt/cdrom
```

- 5 Change to the directory that contains the agent package.

```
cd linux/replication/srdfstar_agent/version/rpms/
```

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/replication/srdfstar_agent/version/pkgsrc/
```

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

- 3 Install the agent binaries.

```
pkgadd -d . VRTSvcsst
```

Removing the agent

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 VRTSvcsst
```

To remove the agent from an HP-UX cluster

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

```
swremove VRTSvcsst
```

To remove the agent from a Solaris cluster

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

```
pkgrm VRTSvcsst
```

To remove the agent from a Linux cluster

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

```
rpm -e VRTSvcsst
```

Configuring the agent for EMC SRDF/Star

- Configuration concepts
- Before you configure the agent
- Configuring the agent

Configuration concepts

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

SRDFStar resource type definition

The agent is represented by the SRDFStar resource type in VCS.

```
type SRDFStar (  
    static keylist SupportedActions = { connect, protect, enable,  
    disconnect, unprotect, disable, reset, halt_reset, halt, cleanup  
    }  
    static int InfoInterval = 600  
    static int NumThreads = 1  
    static int OfflineMonitorInterval = 0  
    static int OnlineTimeout = 1200  
    static int OfflineTimeout = 1200  
    static int RestartLimit = 1  
    static str ArgList[] = { SymHome, GrpName, SyncSite, AsyncSite,  
    UseMostRecentData, KeepDataOn, KeepDataFrom, OnlineOn  
    State, TripTakeover, ForceOnline, HaltOnOffline  
    }  
    str SymHome = "/usr/symcli"  
    str GrpName  
    str SyncSite  
    str AsyncSite  
    int UseMostRecentData = 1  
    str KeepDataOn = SWITCH  
    str KeepDataFrom  
    str OnlineOnState = STAR_PROTECTED  
    int TripTakeover = 0  
    int ForceOnline = 0  
    int HaltOnOffline = 1  
    temp str VCSResLock  
    temp str StarState  
)
```


Attribute definitions for the SRDFStar agent

Review the description of the agent attributes.

Required Attributes

You must assign values to the following required attributes.

GrpName	Specifies the name of one Symmetrix composite group that is a part of the SRDF/Star configuration. Type-dimension: string-scalar Default: ""
SyncSite	Specifies the name of the target site for synchronous replication when the local site becomes the workload site. This attribute should not be set in the VCS configuration at the async site. Type-dimension: string-scalar Default: ""
AsyncSite	Specifies the name of the target site for asynchronous replication when the local site becomes the workload site. Type-dimension: string-scalar Default: ""

Optional attributes

You may optionally configure the following attributes.

SymHome	Specifies the home directory for the Symmetrix CLI software. Type-dimension: string-scalar Default: /usr/symcli
---------	---

UseMostRecentData	<p>Specifies if the local site is to be synchronized so that the most recently available data is used. This attribute value is used only during VCS failover and when the agent is converting the local site to the workload site.</p> <p>If UseMostRecentData is set to 0, the most recent data will not be used.</p> <p>If UseMostRecentData is set to 1, the most recent data will be used.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 1</p> <p>Possible values: 0 or 1</p>
KeepDataOn	<p>Specifies which of either a symstar switch or a symstar connect operation causes the data to be synchronized.</p> <p>Type-dimension: string-scalar</p> <p>Default: "SWITCH"</p> <p>Possible values: "SWITCH"</p>
KeepDataFrom	<p>The name of the site that would serve as a source for data at the time of a symstar switch or symstar connect operation.</p> <p>Set this attribute only if UseMostRecentData is set to 0.</p> <p>Type-dimension: string-scalar</p> <p>Default: ""</p>
OnlineOnState	<p>Indicates the state that the Star system or the target sites must attain before the SRDFStar resource is considered to be in the ONLINE state.</p> <p>Type-dimension: string-scalar</p> <p>Default: "STAR_PROTECTED"</p> <p>Possible values: "STAR_PROTECTED" or "PROTECTED"</p> <p>A value of "STAR_PROTECTED" indicates that the resource is considered to be ONLINE only when the Star system is in the PROTECTED state.</p> <p>A value of "PROTECTED" indicates that the resource is considered to be online only when both target sites are in the PROTECTED state.</p>

TripTakeover	<p>Becomes relevant at the time of switching the workload site to another site. Specifies whether or not the agent is to convert the local site to a workload site when Star is in the TRIPPED state.</p> <p>If TripTakeover is set to 0, the local site is not converted to a workload site.</p> <p>If TripTakeover is set to 1, the local site is converted to a workload site.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p> <p>Possible values: 0 or 1</p>
ForceOnline	<p>Becomes relevant at the time of enabling STAR at the current workload site. Indicates if VCS is to bring the SRDFStar resource online in the event that the target sites, the Star, or both, can not attain a PROTECTED state.</p> <p>If ForceOnline is set to 1, then OnlineOnState attribute is overridden.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p> <p>Possible values: 0 or 1</p> <p>Note: Ensure that the ForceOnline is set to the proper value before you bring the agent online.</p>
HaltOnOffline	<p>Becomes relevant at the time of offlining the SRDFStar resource. Indicates if the halt command must be issued by the agent while going offline.</p> <p>If HaltOnOffline is set to 1, the halt command is issued by the agent while going offline. Set HaltOnOffline to 1 at the async site</p> <p>Type-dimension: integer-scalar</p> <p>Default: 1</p> <p>Possible values: 0 or 1</p>

Internal attributes

The following attributes are for internal use only; do not modify their values.

VCSResLock	Used by the agent to serialize management of composite groups for parallel service groups. Do not modify the value of this attribute Type-dimension: temp string-scalar
StarState	Used by the agent to store the current state of STAR. Do not modify the value of this attribute. The value of this attribute can be displayed using the GUI or the following command: <code>hares -display SRDFStarResourceName</code> Type-dimension: temp string-scalar

Recommended values for the SRDFStar agent attributes

SyncSite	If the sync target site is down, set the SyncSite to "". Set the attribute to "" at the async site because no sync target site exists.
AsyncSite	If the async target site is down, you can set AsyncSite to ""
UseMostRecentData	If the most recent data is to be used at the time of VCS failover, set UseMostRecentData to 1. If UseMostRecentData is set to 1, then the SRDF agent ignores the values of the KeepDataFrom and KeepDataOn.
KeepDataOn	Set this attribute to SWITCH if UseMostRecentData is set to 0.
KeepDataFrom	Set this attribute to the name of the site whose data is to be used at the time of symstar switch and symstar connect. Set this attribute only if UseMostRecentData is set to 0.

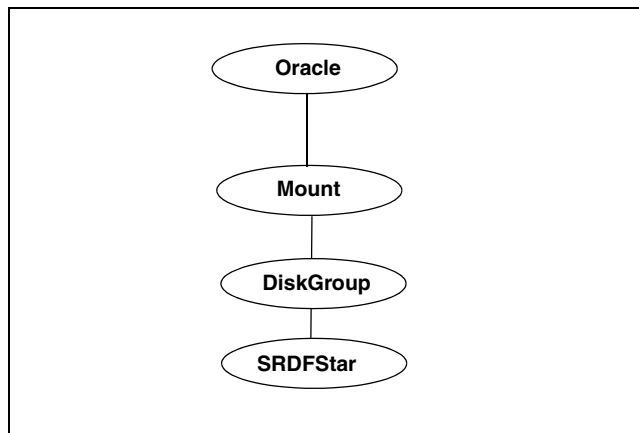
OnlineOnState	<p>Set this attribute to STAR_PROTECTED if the applications should be brought only if STAR is in PROTECTED state. OnlineOnState attribute should be set to PROTECTED if the applications can be brought up even if STAR is not PROTECTED. This attribute is irrelevant at async site. This is because STAR can not be enabled at async site.</p>
ForceOnline	<p>Set this attribute to 1 only if the target sites or STAR cannot attain PROTECTED state. Otherwise, set it to 0. If ForceOnline is set to 1 then OnlineOnState attribute is ignored by the agent. ForceOnline attribute value is checked by the agent at the time of VCS failover after switching the workload site or at the time of onlining the SRDFStar resource on the current workload site.</p>
TripTakeover	<p>Set this attribute to 1 only if the local site has to be converted to workload site and STAR is in TRIPPED state. Otherwise, set this attribute to 0. Its value is checked by the SRDFStar agent before switching the workload site at the time of VCS failover.</p>
HaltOnOffline	<p>Set this attribute to 1 if the data from the current workload site has to be flushed to the target sites at the time of offlining the SRDFStar resource or at the time of VCS failover. Otherwise it should be set to 0. Set HaltOnOffline to 1 at the async site.</p>

Sample configuration

The following figure shows the dependency graph for a VCS service group with a resource of type SRDFStar.

The DiskGroup resource depends on the SRDFStar resource.

Figure 3-2 Sample configuration for the agent



A resource of type SRDFStar comes online first in the application service group.

Before you configure the agent

Before you configure the agent, do the following:

- Set up the SRDF/Star configuration.
- Review the configuration concepts, which describe the type definition and attributes of the agent.
See “[Configuration concepts](#)” on page 24.
- Verify that the agent is installed on all systems in the cluster.
- Verify the hardware setup for the agent.
See “[Typical EMC SRDF/Star setup using VCS clusters](#)” on page 9.
- Make sure the cluster has an effective heartbeat mechanism in place.
See “[About cluster heartbeats](#)” on page 31.
See “[About preventing split-brain](#)” on page 31.

About cluster heartbeats

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.

About preventing split-brain

Split-brain occurs when all heartbeat links between the primary and secondary sites are cut but the hosts at each site remain online. In this situation, each site mistakenly assumes that the other site is down. Minimize the effects of split-brain by ensuring that the cluster heartbeat links pass through similar physical infrastructure as the replication links so that if one breaks, so does the other.

Even if the cluster heartbeat links pass through the same physical infrastructure as the replication links and all links break simultaneously, split-brain can still occur if the service group starts on a remote site. In this case, neither remote site is protecting the original workload site. To prevent this type of split-brain scenario, configure the steward utility on all three sites such that each site acts as the arbitrator for the other two sites.

Configuring the agent

Most clustered applications can be adapted to a disaster recovery environment by doing the following:

- Converting their devices to be a part of the SRDF/Star composite group
- Synchronizing the devices
- Creating a resource of type SRDFStar and adding it to the service group

Symmetrix device groups are configured as resources of type SRDFStar.

Ensure also that you set the `OnlineTimeout` and `OnlineRetryLimit` attributes for the SRDFStar resource so that its entry points do not time out, or so that they are automatically restarted if they do time out.

To configure the agent in a global cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 If the agent resource type (SRDFStar) is not added to your configuration, add it.

In Cluster Explorer, on the File menu, click **Import Types** and then select `/etc/VRTSvcs/conf/SRDFStarTypes.cf`.

- 3 Click **Import**.
- 4 Save the configuration.

Additional configuration considerations

The following are important implementation guidelines for the agent:

- Set the `OnlineTimeout`, `OfflineTimeout`, and `OnlineRetryLimit` attributes for the SRDFStar resource such that its entry points do not time out. If you must set them so that they do time out, ensure that they are automatically restarted.
- If `HaltOnOffline` is set to 1, then the SRDFStar agent takes a long time to go OFFLINE because the `symstar halt` command is invoked by the offline entry point. This means that the system reboot and execution of the `hastop -all` command also takes long time. Set `HaltOnOffline` in these cases to 0. However, in case of planned failover, before switching the workload site to another site, set `HaltOnOffline` to 1 to prevent data loss.
- The info entry point is run periodically as defined by the `InfoInterval` attribute. `InfoInterval` is set to 600 seconds by default. Set it to a different value according to your requirements.
- The info entry point can also be invoked manually to get the state of target sites and STAR using the following command:

```
hares -refreshinfo SRDFStarResName -sys SysName
```
- The value of `ResourceInfo` can be displayed using the following command:

```
hares -value SRDFStarResName ResourceInfo
```
- The value of `ResourceInfo` can be flushed using the following command:

```
hares -flushinfo SRDFStarResName -sys SysName
```

Configuring the agent for VCS global clusters

See the *Veritas Cluster Server Implementation Guide* for the details of setting up VCS global clusters.

The following is a sample `main.cf` file for the Workload site in a VCS global cluster environment:

```
include "types.cf"
include "SRDFStarTypes.cf"
cluster vcs_A (
    UserNames = { admin = aHIaHChEIdIIgQIcHF }
    Administrators = { admin }
)
remotecluster vcs_B (
    ClusterAddress = "a.b.c.d"
```



```
)
remotecluster vcs_C (
    ClusterAddress = "a.b.c.d"
)
heartbeat Icmp (
    ClusterList = { vcs_B, vcs_C }
    AYATimeout = 30
    Arguments @vcs_B = { "a.b.c.d" }
    Arguments @vcs_C = { "a.b.c.d" }
)
system sysA (
)
group ClusterService (
    SystemList = { sysA = 1 }
    AutoStartList = { sysA }
    Tag = CSG
Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
)
IP webip (
    Device = bge0
    Address = "a.b.c.d"
    NetMask = "255.255.240.0"
)
NIC csgnic (
    Device = bge0
)
wac requires webip
webip requires csgnic
group myapp (
    SystemList = { sysA = 1 }
    ClusterList = { vcs_B = 0, vcs_C = 1 }
)
Application testApp (
    StartProgram = "/testApp/start.sh"
    StopProgram = "/testApp/stop.sh"
    MonitorProgram = "/testApp/monitor.sh"
)
Mount Mount-testvoll (
    MountPoint = "/testdir/testvoll"
    BlockDevice = "/dev/vx/dsk/stardg/testvoll"
    FSType = vxfs
    FsckOpt = "-y"
)
DiskGroup VM_stardg (
    DiskGroup = stardg
)
NotifierMngr ntfr (
    SnmpConsoles = { sysA = Error }
    SntpServer = "x.y.yourcompany.com"
```

```
        Smtprcipients = { "xyz@yourcompany.com" = Error }
    )
SRDFStar srdfstar_testAppstar (
    GrpName = STAR-cg
    SyncSite = SYNC_TARGET
    AsyncSite = ASYNC_TARGET
    OnlineOnState = STAR_PROTECTED
)
Mount-testvoll requires VM_stardg
VM_stardg requires srdfstar_testAppstar
testApp requires Mount-testvoll
```

The main.cf file for the sync site is similar to the main.cf for the Workload site except for the SRDFStar resource definition.

The following is the SRDFStar resource definition in the main.cf file for the sync site:

```
SRDFStar srdfstar_testAppstar (
    GrpName = STAR-cg
    SyncSite = WORKLOAD_SITE
    AsyncSite = ASYNC_TARGET
    OnlineOnState = STAR_PROTECTED
)
```

The main.cf file for the async site is similar to main.cf file for the Workload site except for the SRDFStar resource definition.

The following is the SRDFStar resource definition in the main.cf file for the async site:

```
SRDFStar srdfstar_testAppstar (
    GrpName = STAR-cg
    AsyncSite = SYNC_TARGET
)
```

Managing and testing clustering support for EMC SRDF/Star

- Typical test setup
- Test scenarios
- Failure scenarios

Typical test setup

Ensure that you have a working SRDF/Star configuration before testing the agent.

A typical test environment includes the following:

- Two hosts (hosta and hostb) attached to the array at the workload site A.
- Two hosts are attached to the array at the remote sync target site B.
- Two hosts are attached to the array at the remote async target site C.
- The application is running on hosta and Star protection is enabled.
- Clusters in the Star system have one network heartbeat.

See “[Typical cluster setup in SRDF/Star](#)” on page 9.

Test scenarios

Perform the tests described in the following topics to ensure the proper operation of the agent:

- Testing service group migration
- Testing host failure
- Testing disaster recovery
- Testing failback

Testing service group migration

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

To perform the service group migration test

- 1 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group and migrate it to a host that is attached to the same array.
- 2 In the menu, click **Switch To**, and click the hostb system that is attached to the same array.
The service group comes online on hostb and site A remains the workload site with Star protection enabled.
- 3 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group and migrate it to a host that is attached to the array at the sync target site B.

- 4 In the menu, click **Switch To** and click the hostc system that is attached to the array at the sync target site B.
 The service group comes online on hostc and sites A and B switch roles. Site A becomes the new sync target and site B becomes the new workload site with Star protection enabled.
- 5 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group and migrate it back to its original host.
- 6 In the menu, click **Switch To** and then click the system on which the group was initially online (hosta).
 The group comes online on hosta and sites A and B switch roles again. Site A is now the workload site with Star protection enabled and site B is now the sync target.

Testing host failure

Test a recovery procedure for a scenario in which the application host is lost. Eventually, all the hosts in the cluster are lost.

To perform the host failure test

- 1 Halt or shut down the host where the application runs (hosta).
 The service group fails over to hostb.
- 2 Halt or shut down hostb.
 In a global cluster, a cluster down alert appears and the service group fails over to site B or site C, depending on the values of the ClusterList and ClusterFailoverPolicy attributes.
- 3 Reboot the two hosts that were shut down.
- 4 Switch the service group to its original host when VCS starts. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 5 Click **Switch To**, and click the system on which the service group was initially online (hosta).
 The service group comes online on hosta.

Testing disaster recovery

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

To perform a disaster test

- 1 Shut down all hosts and the Symmetrix array at the workload site.

If you cannot shut down the workload site Symmetrix array, disconnect all of the links for the Symmetrix array and simultaneously shut down the hosts. This action mimics a disaster scenario from the point of view of the remote sites.

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

Testing failback

Test the ability of the application to 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 workload site hosts.
- 2 Take the service group offline:

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

To run the test in a global cluster, run the command from `hostc` or `hostd`. Set the SRDFStar resource attribute `KeepDataFrom` before bringing the service group online. `KeepDataFrom` must point to the site from which you want to synchronize data. The default value is the site with the most recent data, which is known to the agent.

- 3 Bring the service group online at the workload site. Type:

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

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

Failure scenarios

The agent enables you to respond to many combinations of host, storage, ICMP link, and replication link failure.

The following topics describe the more common failure scenarios and their recommended responses:

- Workload hosts and storage failure
- Workload hosts failure
- Workload storage failure
- Workload and sync target total site failure
- Workload and async target total site failure
- Replication link failure

Workload hosts and storage failure

In this scenario, the workload hosts and the storage array become unavailable. The original sync and async RDF links and the workload ICMP heartbeat links are lost. Under these conditions, you must use VCS to switch the application to either the sync target site or the async target site. While switching, you have a choice of using the data at either the sync or async target site. Set the `UseMostRecentData` and `KeepDataOn` attributes to specify which data you want to use.

If you switch to the sync target site, it becomes the new workload site. The agent then establishes a new async RDF link from this new workload site to the async target site.

At this point, no synchronous replication is available. Star is not in the `PROTECTED` state.

After the original workload site is available, you can regain Star protection by making the original workload site the new sync target site. You can then use the agent entry points to transition Star to the `PROTECTED` state.

See “[Agent entry points](#)” on page 11.

If you switch to the async target site, it becomes the new workload site. The agent then establishes a new async RDF link from this new workload site to one of the remote sites as specified by the `asyncSite` attribute.

At this point, no synchronous replication link is available. Star is not in the `PROTECTED` state.

Star protection is not available until the original workload site is recovered.

Workload hosts failure

In this scenario, the workload hosts become unavailable, but the storage array is still online. The workload ICMP heartbeat links are lost, but the original synchronous and asynchronous RDF links are still available. Under these conditions, you must use VCS to switch the application over to either the sync or async target site, which becomes the sync or async site.

At this point, the original synchronous and asynchronous replication links and the original workload array are still available.

Workload storage failure

In this scenario, the workload array becomes unavailable, but the application hosts are still online. The original synchronous and asynchronous RDF links are unavailable, but the workload ICMP heartbeat links are available. However, the application cannot run at the original workload site while the storage is down. Under these conditions, you must use VCS to switch the application to either the

sync or the async target site. The result of switching to either site is the same as described in the workload hosts and storage failure scenario.

Workload and sync target total site failure

In this scenario, the hosts and storage at both the workload site and the sync target site become unavailable. Under these conditions, you must use VCS to bring the application online at the async target site.

At this point, no synchronous replication link is available.

After both the workload site and sync target sites are available, you can regain Star protection by using VCS to perform a manual switch back to the workload site.

Workload and async target total site failure

In this scenario, the hosts and storage at both the workload site and the async target become unavailable. Under these conditions, you must use VCS to switch the application to the sync target site.

At this point, no replication link is available.

After both the workload site and sync target sites are available, you can regain Star protection by using VCS to switch the application back to the workload site or you can use the agent action entry points to transition Star to the `PROTECTED` state.

Replication link failure

In this scenario, either or both of the synchronous and asynchronous replication links are lost. When any replication (RDF) link is lost, Star goes into the `TRIPPED` state.

At this point, the application continues to run on the workload site, but simultaneous synchronous and asynchronous replication is not available. Star is not in the `PROTECTED` state.

If desired, you can use the agent to switch the application to the sync or async target sites while the RDF links are unavailable. To do this, set the `ForceOnline` and `TripTakeover` attributes to 1.

Note: Symantec recommends that you set these attributes back to their default values after performing a switch of this type.

After both RDF links are available, Star remains in the `TRIPPED` state because the agent does not monitor the RDF link status. You can regain Star protection by re-enabling it outside of VCS control.

Index

A

attribute definitions 25

C

cluster heartbeats 31

D

disaster test 37

E

EMC SRDF/Star agent

about 8

attribute definitions 25

configuration concepts 24

testing 35

type definition 24

F

failback test 38

M

migrating service group 36

R

resource type definition

EMC SRDF/Star agent 24

S

sample configuration 30

service group, migrating 36

split-brain, preventing in cluster 31

supported hardware 8

supported software 8

T

testing

disaster 37

failback 38

type definition

EMC SRDF/Star agent 24

