

Cluster Server 7.4.2 Agent for EMC SRDF/Star Configuration Guide - Windows

Windows

Last updated: 2020-05-31

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third-party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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

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

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054
<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

infoscaledocs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the agent for EMC SRDF/Star	5
	About the agent for EMC SRDF/Star	5
	Supported software	6
	Typical EMC SRDF/Star setup using VCS clusters	6
	EMC SRDF/Star agent functions	9
	About the agent's online function	13
Chapter 2	Configuring the agent for EMC SRDF/Star	16
	Configuration concepts for the EMC SRDF/Star agent	16
	SRDF/Star resource type definition	16
	Attribute definitions for the SRDF/Star agent	17
	Sample configuration	21
	Before you configure the agent for EMC SRDF/Star	22
	About cluster heartbeats	22
	About preventing split-brain	22
	Configuring the agent for EMC SRDF/Star	23
	Configuring the agent manually in a global cluster	23
	Additional configuration considerations	24
Chapter 3	Managing and testing clustering support for EMC SRDF/Star	26
	Typical test setup for the EMC SRDF/Star agent	26
	Testing service group migration	27
	Testing host failure	28
	Performing a disaster test	29
	Performing the failback test	29
	Failure scenarios for EMC SRDF/Star	30
	Replication link failure	30
	Workload hosts and storage failure	31
	Workload hosts failure	31
	Workload storage failure	31
	Workload and sync target total site failure	32
	Workload and async target total site failure	32

Introducing the agent for EMC SRDF/Star

This chapter includes the following topics:

- [About the agent for EMC SRDF/Star](#)
- [Supported software](#)
- [Typical EMC SRDF/Star setup using VCS clusters](#)
- [EMC SRDF/Star agent functions](#)

About the agent for EMC SRDF/Star

The Cluster Server 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. The 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 that:

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

Supported software


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

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

Typical EMC SRDF/Star setup using VCS clusters

Figure 1-1 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

The workload site comprises one or more workload hosts directly attached to a EMC Symmetrix array. The array contains R1 devices that are part of the SRDF/Star configuration.

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 asynch target.

The array at the sync target

The sync target comprises one or more sync target hosts directly attached to a EMC 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 at the workload site, yet still be close enough to maintain synchronous replication.

The array at the async target

The async target comprises 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 sync target hosts and array must be a sufficient distance away to survive a regional disaster across the workload and sync target sites.

A global cluster configuration

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:

- Workload site cluster A
 ClusterList=[Cluster_A=0, Cluster_B=1, Cluster_C=2]
- Sync target site cluster B
 ClusterList=[Cluster_A, Cluster_B, Cluster_C]
- Async target site cluster C
 ClusterList=[Cluster_A, Cluster_B, Cluster_C]

Note: In a global cluster environment, you must attach all hosts in a cluster to the same EMC Symmetrix array.

A replicated data cluster configuration

The workload hosts and sync hosts belong to one cluster. The async hosts belong to a second cluster. The workload hosts belong to one system zone of replicated data cluster, and sync hosts belong to the second system zone of the replicated data cluster. All workload hosts are connected to the same Symmetrix array. All sync site hosts are connected to the same Symmetrix array.

The cluster (preference) list in the cluster configuration would look like the following for each cluster:

- Workload site cluster A
 ClusterList=[Cluster_A=0, Cluster_B=1]
 System Zones=[Workload_Host=0;
 Sync_Host1=1]

Network heartbeats

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

See [“About cluster heartbeats”](#) on page 22.

EMC SRDF/Star agent functions

The Cluster Server 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 using action entry points

[Table 1-1](#) describes the agent functions:

Table 1-1 Agent functions

Function	Description
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.</p> <p>If open detects the lock file, it waits until at least one of the resource's parent resources is probed.</p> <p>It then does either 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. <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>
online	<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.</p> <p>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 the online entry point determines that the local site is not the workload site, it tries to convert the local site to the workload site. To achieve this conversion, the online entry point 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	<p>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.</p>

Table 1-1 Agent functions (*continued*)

Function	Description
offline	Checks the value of the 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 SMTP notification (if configured to do so) whenever the value of StarState changes to a different value.

Table 1-1 Agent functions (*continued*)

Function	Description
action	<p>Initiates 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.</p> <p>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.</p> <p>The following are the action entry points:</p> <ul style="list-style-type: none"> ■ 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 <p>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.</p>

About the agent's online function

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 same as the workload site or the sync target site. The other is when the local site name is same as 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 same as the workload site or the sync target site (as specified in the options file as a part of the SRDF/Star configuration).

Under these conditions, the online entry point checks the state of Star and then takes one of the following actions:

- If STAR is in the PROTECTED state, the online entry point 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 from 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 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` nor `asyncSite` attributes is specified, the SRDFStar resource remains `OFFLINE`.

Local site name same as asynch target

In this scenario, the local site is the current workload site and the local site name is the same as `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 the `ForceOnline` attribute 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 `TripTakeover` attribute. If `TripTakeover` is set to 0, the SRDFStar resource does not go online. If the `TripTakeover` attribute 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 the 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 the UseMostRecentData attribute is set to 1, then the entry point determines which site has the most recent data using the symstar query command output and uses that output to synchronize the local site.
 - If the UseMostRecentData attribute is set to 0 and the KeepDataOn attribute is set to switch, then the 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 13.

Configuring the agent for EMC SRDF/Star

This chapter includes the following topics:

- [Configuration concepts for the EMC SRDF/Star agent](#)
- [Before you configure the agent for EMC SRDF/Star](#)
- [Configuring the agent for EMC SRDF/Star](#)

Configuration concepts for the EMC SRDF/Star agent

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

SRDF/Star resource type definition

The agent is represented by the SRDF/Star resource type in VCS.

```
type SRDFStar (
static keylist SupportedActions = { connect, protect,
    enable,disconnect, unprotect, disable, reset,
    halt_reset, halt, cleanup,switch }
static int InfoInterval = 600
static int OfflineMonitorInterval = 0
static int OnlineTimeout = 86400
static int OfflineTimeout = 1200
static int RestartLimit = 1
static str ArgList[] = { SymHome, GrpName, SyncSite,
    AsyncSite, UseMostRecentData, KeepDataOn,
```



```

        KeepDataFrom, OnlineOnState, TripTakeover,
        ForceOnline, HaltOnOffline }
str SymHome = "C:\\Program Files\\EMC\\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 StarState
)

```

Attribute definitions for the SRDF/Star agent

Review the description of the agent attributes.

Required Attributes

You must assign values to the following required attributes.

Table 2-1 Required attributes

Attribute	Description
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. Set this attribute at only the sync site; do not set it 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.

Table 2-2 Optional attributes

Attribute	Description
SymHome	<p>Specifies the home directory for the Symmetrix CLI software.</p> <p>Type-dimension: string-scalar</p>
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>

Table 2-2 Optional attributes (*continued*)

Attribute	Description
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 the 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 attribute is set to the proper value before you bring the agent online.</p>

Table 2-2 Optional attributes (*continued*)

Attribute	Description
HaltOnOffline	<p>Becomes relevant and the time of offlining the SRDFStar resource.</p> <p>Indicates if the agent must issue the halt command while going offline.</p> <p>If HaltOnOffline is set to 1, the agent issues the halt command 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.

Table 2-3

Attribute	Description
VCSResLock	<p>Used by the agent to serialize management of composite groups for parallel service groups. Do not modify the value of this attribute</p> <p>Type-dimension: temp string-scalar</p>
StarState	<p>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:</p> <pre>hares -display SRDFStarResourceName</pre> <p>Type-dimension: temp string-scalar</p>

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	Set this attribute to star_protected if the applications should be brought only if STAR is in protected state. The 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 cannot be enabled at async site.
ForceOnline	Set this attribute to 1 only if the target sites or STAR cannot attain the PROTECTED state. Otherwise, set this attribute to 0. If ForceOnline is set to 1, then the value of the OnlineOnState attribute is ignored by the agent. The ForceOnline attribute value is checked by the agent either at the time of VCS failover after switching the workload site or at the time of brining the SRDFStar resource online on the current workload site.
TripTakeover	Set this attribute to 1 only if the local site has to be converted to workload site and STAR is in the 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.
HaltOnOffline	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.

Sample configuration

The DiskGroup resource depends on the SRDFStar resource.

[Figure 2-1](#) shows the dependency graph for a VCS service group with a resource of type SRDF/Star.

Figure 2-1 Sample configuration for the agent



A resource of type SRDF/Star comes online first in the application service group.

Before you configure the agent for EMC SRDF/Star

Before you configure the agent, review the following information:

- Set up the SRDF/Star configuration.
- Verify that you have installed the agent on all the cluster nodes.
- Verify the hardware setup for the agent.
See [“Typical EMC SRDF/Star setup using VCS clusters”](#) on page 6.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 22.
See [“About preventing split-brain”](#) on page 22.

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 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.

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 for EMC SRDF/Star

You can adapt the most clustered applications to a disaster recovery environment by:

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

Symmetrix device groups are configured as resources of type SRDF/Star.

Ensure that you get the OnlineTimeout and OnlineRetryLimit attributes for the SRDF/Star resource so that its entry points do not time out, or so that they are automatically restarted if they do not time out.

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

See [“Sample configuration”](#) on page 21.

Note: You must not change the replication state of devices from primary to secondary and from secondary to primary, outside of a VCS setup. The agent for EMC SRDF/Star fails to detect a change in the replication state if the role reversal is done externally and RoleMonitor is disabled.

Configuring the agent manually in a global cluster

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

To configure the agent in a global cluster

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

```
systemdrive\Program Files\Veritas\cluster server\conf\  
Sample_SRDF/Star\SRDF/StarTypes.cf
```

- 3 Click **Import**.
- 4 Save the configuration.
- 5 The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Additional configuration considerations

The following are important implementation guidelines for the SRDF/Star agent:

- Set the OnlineTimeout, Offline Timeout, 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 timeout, 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 a 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 runs 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 -sys SysName
```

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

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


Sample configuration file for a global cluster environment

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

```
group SRDFStar (
SystemList = { sysA = 0 }
ClusterList = { vcs_A = 0, vcs_B = 2, vcs_C = 1 }
)

SRDFStar srdfstar (
SymHome = "C:\\Program Files\\EMC\\SYMCLI"
GrpName = STAR-cg
SyncSite = SYNC_TARGET
AsyncSite = ASYNC_TARGET
OnlineOnState = STAR_PROTECTED
)
```

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(
SymHome = "C:\\Program Files\\EMC\\SYMCLI"
GrpName = STAR-cg
SyncSite = WORKLOAD_TARGET
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

This chapter includes the following topics:

- [Typical test setup for the EMC SRDF/Star agent](#)
- [Testing service group migration](#)
- [Testing host failure](#)
- [Performing a disaster test](#)
- [Performing the failback test](#)
- [Failure scenarios for EMC SRDF/Star](#)


Typical test setup for the EMC SRDF/Star agent

A typical test environment includes the following characteristics:

- 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.

[Figure 3-1](#) depicts a typical test environment.

Figure 3-1 Typical test setup



Testing service group migration

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

To perform the service group migration test

- 1 In the Cluster Explorer configuration tree, under the **Service Groups** tab, right-click the service group.
 Migrate the service group to a host that is attached to the same array.
- 2 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 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.
 Migrate the service group to a host that is attached to the array at the sync target site B.
- 4 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.
 Migrate the service group back to its original host.
- 6 Click **Switch To** and 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

In this scenario, the host where the application runs is lost. Eventually, all the hosts in the cluster or system zone 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 the site B or site C, depending on the values of the ClusterList and ClusterFailoverPolicy attributes.
 In a replicated data cluster, the service group fails over to the site B or site C, depending on the values of the SystemZones, FailoverPolicy, AutoFailover, ClusterList, and ClusterFailoverPolicy attributes.
- 3 Power on 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.

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 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, 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.

In a replicated data cluster, the service group fails over to site B or site C, depending on the values of the SystemZones, FailoverPolicy, AutoFailover, ClusterList, and ClusterFailoverPolicy attributes.

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 workload site hosts.
- 2 Take the service group offline using the following command:

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

If you run the test in a global cluster, type 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 using the following command:

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

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

Failure scenarios for EMC SRDF/Star

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

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: Veritas 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.

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 [“EMC SRDF/Star agent functions”](#) on page 9.

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.