Cluster Server Agent for EMC RecoverPoint Installation and Configuration Guide

Linux, Solaris

6.2

VERITAS™

Cluster Server Agent for EMC RecoverPoint Installation and Configuration Guide

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

Agent Version: 6.2

Document version: 6.2 Rev 2

Legal Notice

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization

- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

- Upgrade assurance that delivers software upgrades

- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

- Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/
forums/storage-and-clustering-documentation

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Contents

Introducing the agent for EMC RecoverPoint

This chapter includes the following topics:

- About the agent for EMC RecoverPoint

- Supported software

- Supported hardware

- Typical EMC RecoverPoint setup in a VCS cluster

- EMC RecoverPoint agent functions

About the agent for EMC RecoverPoint

The Cluster Server agent for EMC RecoverPoint provides support for application failover and recovery in environments that use the EMC RecoverPoint to replicate data across heterogeneous storage in two different Storage Area Networks (SANs). The agent provides this support for applications that are protected by the EMC RecoverPoint continuous remote replication (CRR) as well as concurrent local and remote replication (CLR) in VCS Global Cluster Option (GCO).

The agent supports synchronous, asynchronous, and dynamic synchronous modes of replication. The agent supports host-based and array-based splitters.

Note: The agent does not support Continuous Data Protection (CDP).

Note: In CLR configuration, when fail-over to LOCAL COPY or to REMOTE COPY occurs, data gets replicated between the new source (Non Production) and PRODUCTION COPY. However, the replication to the copy other than PRODUCTION COPY is stopped. As a result, if `TryMigrate` is set to 1, fail-over cannot happen between new source (Non Production) and REMOTE COPY other than PRODUCTION COPY. If you set `TryMigrate` to 0, and tries fail-over between new source (Non Production) and REMOTE COPY other than PRODUCTION COPY, fail-over happens; however, it does not contain latest data. After failback to the PRODUCTION COPY, replication to LOCAL COPY and REMOTE COPY from production is resumed.

Supported software

For information on the software versions that the agent for EMC RecoverPoint supports, see the Symantec Operations Readiness Tools (SORT) site: https://sort.symantec.com/agents.

Supported hardware

The EMC RecoverPoint agent supports all the storage arrays supported by the RecoverPoint appliance.

Typical EMC RecoverPoint setup in a VCS cluster

Figure 1-1 displays a typical setup in a EMC RecoverPoint environment.

Figure 1-1   Typical EMC RecoverPoint setup in a VCS cluster
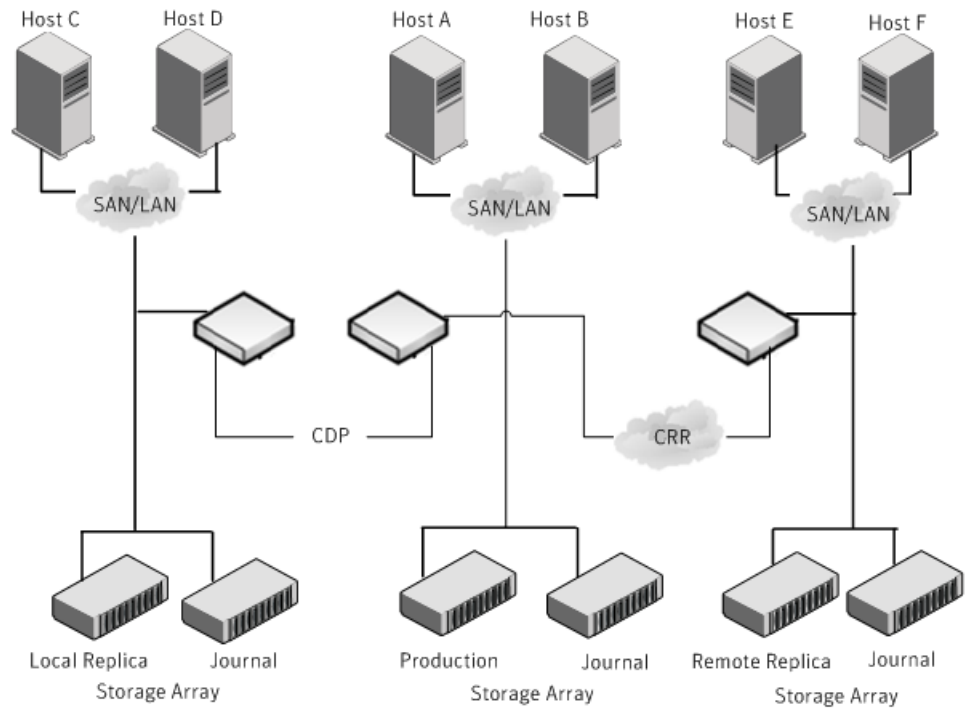
EMC RecoverPoint agent functions

The agent performs the following functions:

Table 1-1   Agent Functions

Function	Description
online	If the role of the specified Consistency Group Copy is PRODUCTION, the agent creates a lock file on the local host to indicate that the resource is online.
	If the role of the specified Consistency Group Copy is LOCAL COPY, REMOTE COPY OR REMOTE REPLICA, the agent invokes RecoverPoint commands to perform a failover.
	See "About the EMC RecoverPoint agent's online function " on page 12.

Table 1-1   Agent Functions *(continued)*

Function	Description
offline	The agent removes the lock file that the online function had created for the resource.
monitor	The agent verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.
clean	The agent removes the lock file from the local host.
open	The agent removes the online lock file on the host where the function is called. This function prevents potential concurrency violation if the service group fails over to another node while this host is down. The agent does not remove the lock file if the agent was started after running the `hastop -force` command.
info	The agent populates the value of the data transfer state of the Consistency Group in the ResourceInfo attribute, for example, ACTIVE, PAUSED.
action/failover	The agent invokes the RecoverPoint `failover` command on the specified Consistency Group and Copy. This action permanently fails over to the Replica Copy in a CRR configuration. **Note:** For this action entry point to work, you need to enable `enable_image_access` on the node where failover is taking place.
action/enable_image_access	The agent invokes the RecoverPoint `enable_image_acess` command on the specified Consistency Group and Copy. This action enables logged access to the image specified in the FailoverImage attribute.
action/disable_image_access	The agent invokes the RecoverPoint `disable_image_acess` command on the specified Consistency Group and Copy. This action disables image access to the currently accessed image. The service group will only fail-over when this action is disabled on the node where failover is taking place.

Table 1-1   Agent Functions *(continued)*

Function	Description
action/get_group_state	The agent invokes the RecoverPoint `get_group_state` command on the specified Consistency Group.
	This action displays the state of the specified Consistency Group.
action/start_transfer	The agent invokes the RecoverPoint `start_transfer` command on the specified Consistency Group.
	This action starts data transfer between the copies in the Consistency Group.
action/pause_transfer	The agent invokes the RecoverPoint `pause_transfer` command on the specified Consistency Group.
	This action pauses data transfer between the copies in the Consistency Group.
action/set_retry_migrate_flag	The agent sets a flag that indicates that the previous migration attempt must be retried. See "Exceptional Conditions and Recovery" on page 13.
action/unset_retry_migrate_flag	The agent clears a flag that indicates that the previous migration attempt must be retried. See "Exceptional Conditions and Recovery" on page 13.
action\GetCurrentRPO	Fetches the current point in time RPO. The agent performs this action function on the disaster recovery (DR) system. The RPO is computed in seconds in terms of Time, Data Lag, and Writes Lag.
	Note: The agent does not compute the RPO when the group is frozen.
	The agent does not store the computed RPO; make a note of the RPO for future reference.

About the EMC RecoverPoint agent's online function

The agent checks the role of the Copy specified in the CopyName attribute.

If the role is PRODUCTION, the agent goes online directly.

If the role is LOCAL COPY, REMOTE COPY or REMOTE REPLICA, the agent validates that no image is being accessed currently. After this validation, the agent checks the value of the TryMigrate attribute.

If the TryMigrate attribute is set to 1, the agent tries to perform a migration, that is, ensure that the latest data has arrived on the secondary site. To achieve this, the agent applies a bookmark on the current image on the primary site, and waits for the bookmark to arrive on the secondary site. Once the bookmark arrives, the agent enables logged image access to the image specified in the FailoverImage attribute. After image access is completed, the agent fails over to the image.

If the TryMigrate attribute is set to 0, the agent directly enables logged image access to the image specified in the FailoverImage attribute, and after image access has completed, the agent fails over to the image. In an asynchronous mode of replication, depending on the currently available data on the secondary site, there may potentially be a loss of data with this setting.

Exceptional Conditions and Recovery

If the value of the TryMigrate attribute is set to 1, and the latest bookmark does not arrive before the OnlineTimeout expires (with some buffer time allocated by the agent), the agent flushes and freezes the group indicating that the bookmark has not arrived.

You can either invoke a fresh online attempt or you can perform the following steps to retry the migration attempt that failed:

1   Unfreeze the group.

2   Run the set_retry_migrate_flag action. This action sets a flag that indicates that the agent must retry the migration attempt that had previously failed.

3   Bring the RecoverPoint resource online.

    The flag that indicates that the agent must retry the previously-failed migration attempt is automatically cleared.

4   If you do not bring the RecoverPoint resource online, you must run the unset_retry_migrate_flag action to clear the flag.

If the enabling of image access takes a long time to complete, and does not complete before OnlineTimeout expires, then the agent flushes and freezes the group indicating enabling image access has not completed. To recover from this condition, perform the following steps:

1   Unfreeze the group.

2   Invoke the failover action.

3   Bring the RecoverPoint resource online.

Installing and removing the agent for EMC RecoverPoint

This chapter includes the following topics:

- Before you install the agent for EMC RecoverPoint

- Installing the agent for EMC RecoverPoint

- Removing the agent for EMC RecoverPoint

Before you install the agent for EMC RecoverPoint

Before you install the Cluster Server agent for EMC RecoverPoint, ensure that you install and configure the VCS on all nodes in the cluster.

Set up replication and the required hardware infrastructure. For information about setting up Oracle RAC environment, refer to the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*.

Installing the agent for EMC RecoverPoint

You must install the EMC RecoverPoint agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

To install the agent in a VCS environment

1 Download the Agent Pack from the Symantec Operations Readiness Tools (SORT) site: https://sort.symantec.com/agents.

You can download the complete Agent Pack tar file or the individual agent tar file.

2 Uncompress the file to a temporary location, say /tmp.

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

Linux	`cd1/linux/generic/vcs/replication/recoverpoint_agent/` `agent_version/rpms/`

Solaris	`cd1/solaris/dist_arch/vcs/replication/recoverpoint_agent/` `agent_version/pkgs/`

If you downloaded the individual agent tar file, navigate to the pkgs directory (for AIX, and Solaris), or the rpms directory (for Linux).

4 Log in as a superuser.

5 Install the package.

Linux	`# rpm -ihv \` `VRTSvcsrecpt-AgentVersion-Linux_GENERIC.noarch.rpm`

Solaris	`# pkgadd -d . VRTSvcsrecpt`

Removing the agent for EMC RecoverPoint

Before you attempt to remove the agent, make sure the application service group is not online.

You must remove the HTC agent from each node in the cluster.

To remove the agent, type the following command on each node. Answer prompts accordingly:

Linux	`# rpm -e VRTSvcsrecpt`
Solaris	`# pkgrm VRTSvcsrecpt`

Configuring the agent for EMC RecoverPoint

This chapter includes the following topics:

- Configuration concepts for the EMC RecoverPoint agent
- Before you configure the agent for EMC RecoverPoint
- Configuring the agent for EMC RecoverPoint

Configuration concepts for the EMC RecoverPoint agent

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

Attribute definitions for the EMC RecoverPoint agent

The description of the agent attributes are as follows:

Table 3-1   Attributes for the EMC RecoverPoint agent

Attribute	Description
RPAAddr	Specifies the current site RPA name or IP address.
	Type-dimension: string-scalar
	Example: 10.182.200.100
	Example: RPA_US_MV

Table 3-1   Attributes for the EMC RecoverPoint agent *(continued)*

Attribute	Description
RPAUser	Specifies the user name used to connect to the current site RPA. Type-dimension: string-scalar Example: admin
SSHBinary	Contains the absolute path to the SSH binary. SSH is the mode of communication with the RPA. Type-dimension: string-scalar Default : /usr/bin/ssh
SSHPathToIDFile	Contains the absolute path to the identity file used for authenticating the host with the RPA. The corresponding public key must be added on the RPA so that the RPA can correctly authenticate the host. Type-dimension: string-scalar Example: /rpa/rpa_rsa
ConsistencyGroupName	Specifies the consistency group name. Type-dimension: string-scalar Example: SQL2008_CG
CopyName	Specifies the consistency group current site copy name. Type-dimension: string-scalar Example: copyA

Table 3-1 Attributes for the EMC RecoverPoint agent *(continued)*

Attribute	Description
FailoverImage	Specifies the image to be used for failover.
	Valid values are:
	■ IMAGE=latest
	■ IMAGE=[bookmark_name]
	■ TIME=Timestamp of the image to be accessed in the format HH:MM:[SS[:MICROS]] [DD/MM/YYYY]].
	You can specify only one image in this attribute.
	Type-dimension: string-assoc
	Examples:
	■ IMAGE=latest
	■ IMAGE=vss_exch
	■ TIME=15:27:33:123456 22/01/2010
	Note: When the FailoverImage attribute is set to a bookmark, for example, FailoverImage = { IMAGE=bookmark_name } and the TryMigrate attribute is set to 1, then the bookmark gets preference. In such a case, the bookmark data is available after the remote switchover operation is successfully completed.
TryMigrate	Specifies whether a migration must be attempted.
	Valid values are 1 and 0.
	If this attribute is set to 1, the agent tries to ensure that the latest data on the primary site has reached the secondary site.
	If this attribute is set to 0, then the agent fails over using whatever is the data currently available on the secondary site, which may not be the latest, and therefore data-loss may occur.
	Default: 1
	Type-dimension: integer-scalar
StartTransfer	Specifies whether data transfer is started from the new source after failover.
	Valid values are 1 and 0.
	If this attribute is set to 1, data transfer is started after failover.
	If this attribute is set to 0, data transfer is not started after failover.
	Default: 1
	Type-dimension: integer-scalar

Table 3-1   Attributes for the EMC RecoverPoint agent *(continued)*

Attribute	Description
AgentOpts	Lists the options for controlling agent behavior. This attribute is for future use. Type-dimension: string-vector

# Resource type definition for the EMC RecoverPoint agent

Following is the resource type definition for the EMC RecoverPoint agent:

```
type RecoverPoint (
    static keylist SupportedActions = { enable_image_access,
    disable_image_access, failover, get_group_state, start_transfer,
    pause_transfer, set_retry_migrate_flag, unset_retry_migrate_flag,
    GetCurrentRPO }
    static str ArgList[] = { RPAAddr, RPAUser, SSHBinary,
    SSHPathToIDFile, ConsistencyGroupName, CopyName, FailoverImage,
    TryMigrate, StartTransfer, AgentOpts, VCSBookMark, VCSRetryMigrate }
    str RPAAddr
    str RPAUser = admin
    str SSHBinary = "/usr/bin/ssh"
    str SSHPathToIDFile
    str ConsistencyGroupName
    str CopyName
    int TryMigrate = 1
    str FailoverImage{} = { IMAGE=latest }
    int StartTransfer = 1
    str AgentOpts[]
    temp str VCSBookMark
    temp int VCSRetryMigrate
    int ComputeDRSLA = 1
)
```

Sample configuration for the EMC RecoverPoint agent

Figure 3-1 shows the dependency graph for a VCS service group with a resource of type RecoverPoint.

Figure 3-1   Sample configuration for the EMC RecoverPoint agent
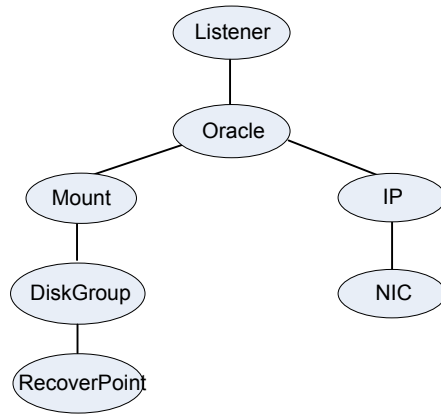
Before you configure the agent for EMC RecoverPoint

Before you configure the agent, review the following information:

- Set up passwordless ssh from all VCS hosts to the RPAs.
  See "To set up a passwordless ssh" on page 21.

- Verify that you have installed the agent on all systems in the cluster.

- Verify the hardware setup for the agent.

- Make sure that the cluster has an effective heartbeat mechanism in place.
  See "About cluster heartbeats" on page 21.

- For each RecoverPoint Consistency Group that uses Storage Foundation,
  configure the Reservations Policy for each copy to SCSI-2. Without this setting,
  DiskGroups will not function correctly in RecoverPoint environments.

To set up a passwordless ssh

To generate SSH keys on a VCS node and setup passwordless ssh to the Recover Point Appliance (RPA), perform the following steps:

1 Log on to host (VCS node).

2 Generate public and private RSA keys using ssh-keygen with no passphrase. For example:

```
bash-3.00# ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (//.ssh/id_rsa): test_rsa
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in test_rsa.
Your public key has been saved in test_rsa.pub.
The key fingerprint is:
05:c1:ce:d0:f7:a1:9f:d1:33:20:bf:c2:9b:35:c8:43 root@thor393
```

3 You can now see the generated public key to add it on the RPA.

For example: `cat //.ssh/id_rsa.pub`

4 Copy the above generated public key and add it on the RPA using the `add_ssh_key` command.

5 To cache the RPA server key on to the host for further use, run the following command from the host:

`/usr/bin/ssh -i <path to RSA private key file> -l <RPA user name> <IP address of the RPA > <dummy_rp_command>`

For example: `# /usr/bin/ssh -i /test_rsa -l admin 10.182.1.29 get_current_time`

A message appears, asking you to cache the host key with that of the RPA server.

6 Type **Y** and press **Enter**.

7 Repeat the above steps on each VCS node where the Service Group contains the RecoverPoint resource.

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.

To minimize the chances of split-brain, use the steward process.

Configuring the agent for EMC RecoverPoint

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

How the agent retrieves the RPO

In a global cluster environment, the agent for EMC RecoverPoint can retrieve the recovery point objective (RPO) from RecoverPoint. The agent uses the `get_group_statistics` command to retrieve the RPO. In a DR configuration where data is replicated asynchronously to the DR site, the DR site data might not always be as current as the primary site data.

RPO is the maximum acceptable amount of data loss in case of a disaster at the primary site. The agent computes RPO in terms of time (seconds), Data Lag, and Writes Lag.

For example:

Replication Lag = 2 sec

Data Lag = 0.00B

Writes Lag = 0

To retrieve the RPO:

1   Run the following command on a node in the DR cluster.

```
hares -action RecoverPoint_resourcename GetCurrentRPO -sys
system_name
```

The action entry point displays the RPO.

Note: The agent does not store the RPO; make a note of the RPO for future reference.

---

2   If the RPO is not reported, check the RecoverPoint configuration.

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 (RecoverPoint) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types**, and select:

/etc/VRTSvcs/conf/RecoverPointTypes.cf

3 Click **Import**.

4 Save the configuration.

5 Add a resource of type RecoverPoint at the bottom of the service group.

6 Configure the attributes of the RecoverPoint resource.

7 If the service group is not configured as a global service group, configure the service group using the Global Group Configuration Wizard.

Refer to the *Cluster Server Administrator's Guide* for more information.

8 Change the ClusterFailOverPolicy attribute from the default, if necessary. Symantec recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.

After configuration, the application service group must follow the dependency diagram. See "Sample configuration for the EMC RecoverPoint agent" on page 19.

Managing and testing clustering support for EMC RecoverPoint

This chapter includes the following topics:

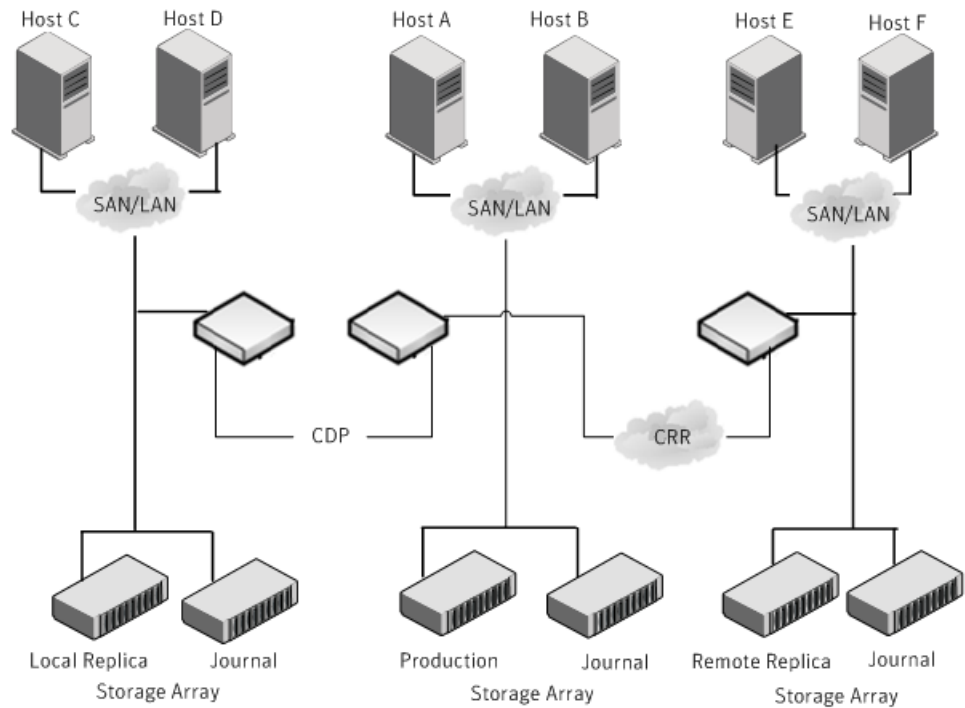- Typical test setup for the EMC RecoverPoint agent
- Testing service group migration
- Testing all hosts failure
- Performing a disaster test

Typical test setup for the EMC RecoverPoint agent

Figure 4-1 depicts a typical test environment.

Figure 4-1   Typical test setup

A typical test environment includes the following characteristics:

- A primary site RPA cluster is attached to the PRODUCTION COPY as well as LOCAL COPY, and the primary application hosts.

- A secondary site RPA cluster attached to the secondary storage and the secondary application hosts.

- Either CRR or CLR is established between the RPA clusters on the two sites.

- In CRR configuration, two hosts (Host A and Host B) are attached to the primary site RPA cluster.

- In CLR configuration, Host A is attached to the PRODUCTION COPY and Host B is attached to the LOCAL COPY. Host C and Host D are attached to REPLICA COPY and are connected to secondary site RPA cluster.

- Passwordless SSH configuration is established on the hosts (Host A and Host B) and the primary site RPA cluster so that the hosts (Host A and Host B) can communicate through SSH with the primary site RPA cluster. Similarly, passwordless SSH is established on the hosts (Host C and Host D) and the secondary site RPA cluster so that the hosts (Host C and Host D) can

communicate through SSH with the secondary site RPA cluster. The application runs on Host A, which is connected to the primary site RPA cluster.

Testing service group migration

After you configure the VCS agent for EMC RecoverPoint, verify that the global service group can migrate to hosts across the sites.

To test the global service group migration in global cluster setup using the VCS GUI

1   Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

2   In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group and click **Online** to bring the service group online on the primary cluster.

3   To switch over the service group from the primary cluster to the secondary cluster, right-click the service group and select **Remote Switch**.

   The service group comes online on the secondary cluster. Now, the secondary copy becomes the production and the original primary copy becomes the remote replica. The latest data is available after migration.

4   To switch back the service group to its original primary cluster, right-click the service group and select **Remote Switch**.

   The service group comes online on its original primary cluster, and the latest data is available.

To test the global service group migration in global cluster setup using the command line interface (CLI)

1   Switch over the global service group from the primary site to the secondary site.

   Perform the following steps:

   ■   Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

   ■   Switch the global service group from the primary site to any node in the secondary site.

       `hagrp -switch global_group -any -clus cluster_name`

       VCS brings the global service group online on a node at the secondary site.

- Verify that the EMC RecoverPoint devices at the secondary site are write-enabled, and the Copy is PRODUCTION.

2 Switch back the global service group from the secondary site to the primary site.

Perform the following steps:

- Switch the global service group from the secondary site to the primary site.

  ```
  hagrp -switch global_group -any -clus cluster_name
  ```
  VCS brings the global service group online at the primary site.

- Verify that the EMC RecoverPoint devices at the primary site are write-enabled, and the Copy is PRODUCTION.

To test the local service group migration in global cluster setup using the VCS GUI

1 Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

2 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group and click **Online** to bring the service group online on the node in primary cluster.

3 To switch over the service group from the primary node to the secondary node in the same cluster, right-click the service group and select **Switch To** and then node name. The service group comes online on the secondary node. Now, the secondary copy becomes the production and the original primary copy becomes the local replica. The latest data is available after migration.

4 To switch back the service group to its original primary node in the same cluster, right-click the service group and select **Switch To** and then node name. The service group comes online on its original primary node, and the latest data is available.

To test the local service group migration in global cluster setup using the command line interface (CLI)

1 Switch over the global service group from the primary node to the secondary node in the same cluster.

Perform the following steps:

- Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

- Switch the global service group from the primary node to the secondary node in the same cluster.

  ```
  hagrp -switch global_group -any -clus cluster_name
  ```
  or

```
hagrp -switch <group> -to <system> [-clus <cluster> |
-localclus [-nopre]]
```

VCS brings the global service group online on the secondary node in the same cluster.

■ Verify that the EMC RecoverPoint devices at the secondary node are write-enabled, and the Copy is PRODUCTION.

2 Switch back the global service group from the secondary node to primary node.

Perform the following steps:

■ Switch the global service group from the secondary node to the primary node in the same cluster.

```
hagrp -switch global_group -any -clus cluster_name
```
or
```
hagrp -switch <group> -to <system> [-clus <cluster> |
-localclus [-nopre]]
```

VCS brings the global service group online at the primary node.

■ Verify that the EMC RecoverPoint devices at the primary node are write-enabled, and the Copy is PRODUCTION.

Testing all hosts failure

Perform the following procedure to test how VCS recovers after all hosts at the primary site fail.

To test disaster recovery for all hosts failure in global cluster setup using the VCS GUI

1 Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

2 Halt the hosts at the primary site.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the VCS failover behavior.

■ Auto—VCS brings the faulted global service group online at the secondary site.

■ Manual or Connected—You must bring the global service group online at the secondary site.

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

3 Verify that the EMC RecoverPoint devices at the secondary site are write-enabled, and the Copy is PRODUCTION.

Verify that the latest data is available.

To test disaster recovery for all hosts failure in global cluster setup using the Command Line Interface (CLI)

1 Ensure that the TryMigrate attribute is set to 1 (TryMigrate = 1), which is the default setting.

2 Halt the hosts at the primary site.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the VCS failover behavior.

- Auto—VCS brings the faulted global service group online at the secondary site.

- Manual or Connected—You must bring the global service group online at the secondary site.
  On a node in the secondary site, run the following command:
  ```
  hagrp -online -force global_group -any
  ```

3 Run the following command to verify that the global service group is online at the secondary site:

```
hagrp -state global_group
```

4 Verify that the EMC RecoverPoint devices at the secondary site are write-enabled, and the Copy is PRODUCTION.

Verify that the latest data is available.

Performing a disaster test

Halt the primary hosts and the RPAs on the primary site and fail the CRR replication link between the two sites.

The primary cluster goes into the FAULTED state.

The value of the ClusterFailOverPolicy attribute for the faulted service group determines the VCS failover behavior.

- Manual—You must bring the service group online at the secondary site. Set the value of the TryMigrate attribute to 0. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group and click **Online**. After

the service group is online, Symantec recommends that you set the value of the TryMigrate attribute to 1.

- Auto—If the TryMigrate attribute is set to 1, VCS attempts to bring the faulted global service group online at the secondary site. However, the service group fails to come online, because the latest data cannot be obtained from the primary site. To successfully bring the service group online in such a situation, first set the value of the TryMigrate attribute to 0 and then bring the service group online manually at the secondary site. In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group and click **Online**. After the service group is online, Symantec recommends that you set the value of the TryMigrate attribute to 1.

Note: Disable automatic failback to the original cluster until the replication link is restored.

Index