Veritas™ Cluster Server One Agent for Oracle Data Guard Installation and Configuration Guide

Solaris

5.0 Service Pack 2

December 2010

Symantec™

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

Agent Version: 5.0SP2.0

Document version: 5.0SP2.0.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Documentation

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:

docs@symantec.com

Support agreement resources

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

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

Contents

Chapter 1

Introducing the Veritas Cluster Server One Agent for Oracle Data Guard

This chapter includes the following topics:

About the agent for Oracle Data Guard

The Veritas agent for Oracle Data Guard provides failover support and recovery in an environment that uses the Oracle Data Guard. Oracle Data Guard replicates data between Oracle databases.

The agent monitors and manages the state of replicated Oracle databases that run on VCS nodes. The Data Guard resource is online on the system with the primary database server. The agent makes sure that Oracle Data Guard replicates the database information from the primary database server to the standby database server.

You can use the Data Guard agent in global clusters that run VCS.

The Veritas agent for Oracle Data Guard Broker manages the replication in Oracle 10gR2 and 11gR1 databases in parallel applications such as Veritas Storage Foundation for Oracle RAC. This agent uses the Oracle Data Guard Broker to manage the database replication in a parallel application environment. The Data

Guard Broker agent simplifies the RAC database switch over or fail over using the Data Guard command-line interface DGMGRL.

You can use the Data Guard Broker agent in global clusters that run SF Oracle RAC.

The VCS agent for Data Guard does not support database environments under the control of Oracle Enterprise Manager.

Note: The Data Guard agent and the Data Guard Broker agent do not support replicated data clusters.

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 for Oracle Data Guard

The agent for Oracle Data Guard supports the following software versions:

Veritas Cluster Server One	■ VCS One 5.0 SP2 on Solaris SPARC Oracle 11gR2 See the product's Release Notes for more details on the supported architectures and the operating system versions.

Typical Oracle Data Guard setup in a VCS cluster

Figure 1-1 displays a typical cluster setup in a Data Guard environment.

Figure 1-1 Typical clustering setup for the agent

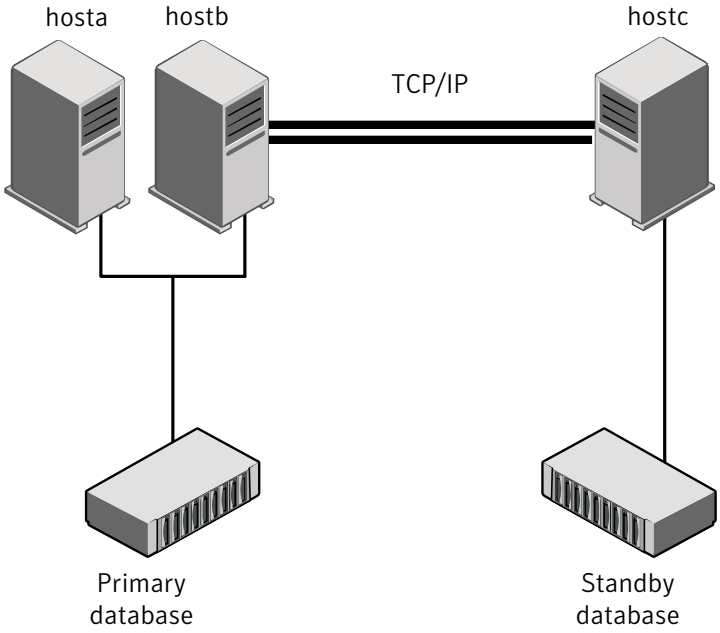

Clustering in a Data Guard environment typically consists of the following hardware infrastructure:

- The primary database instance (db1) sends redo data across a TCP/IP link to a standby database instance (db2). A local cluster protects the primary database and makes it highly available.
- The standby database instance applies the redo information to a physical copy of the primary database.
- The primary and standby sites must be connected through a single TCP/IP network connection. This link can be shared with VCS global clusters for heartbeat communication.
- Network heartbeating between the two data centers to determine their health; this network heartbeating could be LLT or TCP/IP.

Agent functions for the Data Guard agent

The Oracle Data Guard agent monitors and manages the state of replicated Oracle database that runs on VCS nodes. Agent functions bring resources online, take

them offline, and perform different monitoring actions. Agent functions are also known as entry points.

The agent also supports DataGuard role transition.

See "About DataGuard role transition" on page 14.

online	Creates a lock file on the local host to indicate that the resource is online. Depending on the role of the database, the agent performs actions to make the database accessible. See "About the Oracle DataGuard agent's online function" on page 13.
offline	Removes the lock file on the local node. Because a switch of the replication direction, promoting the standby and demoting the primary is executed on the target node. Oracle reconfiguration is not done as part of offline. In case of a complete shutdown, an Oracle resource is responsible to close the database.
monitor	Verifies that the lock file exists. ■ If the lock file does not exist, the monitor entry point reports the status of the resource as offline. ■ If the lock file exists, the agent checks if the role of the database is still PRIMARY and the open mode is WRITE.
open	Creates a lock file in the local agent directory if the role of the database is PRIMARY and the open mode is WRITE.
clean	Removes the lock file for the following resource states: ■ OFFLINE TIMEOUT ■ OFFLINE INEFFECTIVE ■ ONLINE TIMEOUT ■ UNEXPECTED OFFLINE ■ MONITOR HUNG
info	Reports the state and the role of the database.
start_stb_curlog.sql	Custom startup script for the VCS agent for Oracle. See "About the custom startup script for the Oracle agent" on page 13.
actions/DGStatus	Reports the current state and role of the database in real time.

actions/activateStandby	It enables the physical standby database to be opened in a read only mode with redo apply from a mounted state. It works only when the state of database is Mounted.
actions/deactivateStandby	Running this action entry point causes the physical standby database to be shutdown and then mounted with redo apply.

About the Oracle DataGuard agent's online function

The agent determines the role of the database and the type of open mode using the SQL commands:

```
DATABASE_ROLE from V$DATABASE
OPEN_MODE from V$DATABASE
```

If the role of the replicated database is PRIMARY and the open mode is MOUNT, the agent makes the database accessible for clients as follows:

- Alters the database to open mode READ WRITE.
- Creates a lock file on the local host to indicate that the resource is online.

If the role of the database is PHYSICAL STANDBY, the agent assumes a site fault and reconfigures the database as follows:

- The agent first tries to demote a primary database instance by executing the action `DGDemotePri` inside the remote cluster.
- Then, the agent changes the mode of the local database from PHYSICAL STANDBY to PRIMARY.
  The agent stops the reception of redo log information using the SQL command:

  ```
  ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL
  ```

  The agent changes the role of the database using the SQL command:

  ```
  ALTER DATABASE COMMIT TO SWITCHOVER TO PRIMARY
  ```

- The agent then restarts the local database instance to make the changes effective and creates a lock file in the local agent home directory.

About the custom startup script for the Oracle agent

The Oracle Data Guard agent uses a custom startup script start_stb_curlog.sql to start the Oracle agent. The Oracle database instance start has to be implemented by using a VCS resource of type Oracle with the attribute StartUpOpt set to

CUSTOM. The necessary file start_custom_<InstID>.sql can then be implemented as a symbolic link to the start_stb_curlog.sql file.

Depending on the database role, the agent does the following actions:

- If the database role is PRIMARY, the agent mounts the database.
- If the database role is PHYSICAL STANDBY, the agent mounts the database. Then, the agent executes the following SQL command to start the replication reception:

```
ALTER DATABASE RECOVER MANAGED STANDBY DATABASE USING
CURRENT LOGFILE DISCONNECT FROM SESSION
```

About DataGuard role transition

You can switch the service group in which the DataGuard resource resides using the `hagrp -switch` command.

If the agent is OFFLINE on the original primary, the agent removes the lock file.

If the agent is ONLINE on the former standby, the agent executes the following actions:

- Execute action DGDemotePri on the original primary.
- Alter database role from standby to primary.
- Restart Oracle instance on the standby.

Chapter 2

Installing and removing the agent for Oracle Data Guard

This chapter includes the following topics:

Installing the Veritas Cluster Server One agent for Data Guard on UNIX

You must install the agent for Data Guard on all the client systems of the VCS One cluster that will host the Data Guard service group. You can install the agent for Data Guard using the installagpack program or using the command line interface (CLI).

The installation of the agent packs involves the following phases:

Installing the agent packages	See "Installing the agent packages using the installer" on page 16.
Adding the agent resource type definitions	See "Adding the agent resource type definitions to the Policy Master Server on UNIX" on page 17.

Note: The installagpack program supports only the -addtypes, -rmtypes, -responsefile, and -rsh options. Symantec recommends that you do not use any of the other options from the `installagpack` command help output.

Installing the agent packages using the installer

You can install the agent packages on one or more client systems of a specific platform type.

Note: To install the VCS One client for managing VMware ESX Servers, download the tar ball for Red Hat Enterprise Linux 4 (RHEL 4) x86 (32-bit) or RHEL 5 x86_64

Perform the following steps to install the agent packages using the installer

1. Download the complete Agent Pack tarball from FileConnect site, on the Policy Master system:

   https://fileconnect.symantec.com/

   Alternatively,

   Download the individual agent tarball from the Symantec Operations Readiness Tools (SORT) site:

   https://sort.symantec.com/home

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

3. If you downloaded the complete Agent Pack tarball, navigate to the following directory containing the installer for the VCS One agents, for the platform running in your environment:

Solaris	`cd1/solaris/dist_arch/vcsone/vcsone_version` Where, *dist_arch* is 'sol_sparc' or 'sol_x64'.

4. Enter the following command to start the agent pack installation:

   ```
   # ./installagpack [-rsh]
   ```

   You can use the `-rsh` option if rsh and rcp are used for communication between systems instead of the default ssh and scp. This option requires that systems be preconfigured such that the `rsh` commands between systems execute without prompting for passwords or confirmations.

5. Enter the name of the client systems where you want to install the agents.

6. Choose whether to install all the agents or any specific agent. Follow the installer prompt to specify your option.

7. Review the output as the installation program installs the agent packages.

   You can view installation logs in the /var/VRTS/install/logs directory.

Installing the agent package using the CLI

You can install the desired agent package using the CLI, on one or more client systems of a specific platform type.

Perform the following steps to install the agent packages using CLI

1. Download the complete Agent Pack tarball from FileConnect site, on the Policy Master system:

   https://fileconnect.symantec.com/

   Alternatively,

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

   https://vos.symantec.com/home

2. Uncompress the file to a temporary location, say /tmp.
3. If you downloaded the complete Agent Pack tarball, navigate to the following directory containing the installer for the VCS One agents, for the platform running in your environment:

   Solaris `# cd1/solaris/dist_arch/vcsone/vcsone_version/pkgs`

   Where *dist_arch* is 'sol_sparc' or 'sol_x64'

4. Type the following command on each client system to install the agent. Answer the prompt accordingly:

   Solaris `# pkgadd -d .`

Adding the agent resource type definitions to the Policy Master Server on UNIX

You must add the agent resource type definitions to the Policy Master database configuration. You can perform this task from any client system in the VCS One cluster.

Note: You must add the agent resource type definitions only one time per platform type.

To add the agent resource types to the policy master database configuration

1. Set up RSH or SSH communications between the client system and the policy master system.

   For information on configuring SSH for remote communication, refer to the *Veritas Cluster Server One Installation Guide*.

2. Make sure that the PM daemon is running.

   ```
   # /opt/VRTSvcsone/bin/haclus -display
   ```

   The output should show ClusterState is RUNNING.

3. Access the temporary location where you downloaded the tar ball and depending on the platform type, navigate to the directory containing the agent installer:

Solaris	`cd1/solaris/`*`dist_arch`*`/vcsone/`*`vcsone_version`* Where *dist_arch* is the sol_sparc or sol_x64.

4. Enter the command to start the agent pack installer for adding resource types to the Policy Master configuration database. Use the `-addtypes` option:

   ```
   # ./installagpack -addtypes
   ```

5. When the installer prompts, enter the virtual IP address of the Policy Master.
6. Review the output as the installer verifies communication with the Policy Master system.
7. Choose whether to add the type definitions for all the agents or for specific agents. Follow the installer prompts to add the type definitions.
8. Review the output as the installer adds the agent types to the PM database configuration and copies the appropriates types.xml files to the PM system.

   You can view installation logs in the /var/VRTS/install/logs directory.

Removing the Veritas Cluster Server One agent for Data Guard on UNIX

Removing the agent package involves removing the agent files from each client system where it was installed.

You can remove the packages using the agent pack installer or the command line.

See "Removing the agent packages using the installer" on page 19.

See "Removing the agent package using CLI" on page 20.

After removing the agent packages you can remove the agent type definition from the Policy Master system.

See "Removing the agent type definition from the Policy Master system on UNIX" on page 20.

Removing the agent packages using the installer

You can remove all the agent packages or the desired agent package using the uninstallagpack program.

Note: The uninstallagpack program supports only the -responsefile and -rsh options. Symantec recommends that you do not use any of the other options from the `uninstallagpack` command help output.

To remove the agent packages from the client systems

1 Freeze the service groups that hosts the application, on the system from which you want to remove the agent package.

```
# hagrp -freeze <groupname>
```

2 Stop the agent on all client systems before you remove the agent package from the system.

```
# haagent -stop -notransition <AgentName> -sys <system_name>
```

3 Ensure that the agent operations are stopped on all the cluster systems.

```
# haagent -display <AgentName>
```

4 Access the temporary location where you downloaded the Agent Pack and navigate to the directory containing the package for the platform running in your environment:

Solaris	**cd1/solaris/*dist_arch*/vcsone/*vcsone_version***
	Where *dist_arch* is the sol_sparc or sol_x64.

5 Start the uninstallagpack program.

```
# ./uninstallagpack [-rsh]
```

6 Enter the name of the client systems on which you want to uninstall the agent pack. The names must be separated by spaces.

7 Choose whether to remove all the agent packages or a specific agent package. Follow the installer prompt to remove the agent package.

8 Review the output as the program verifies the agent pack that you installed and removes the agent packages.

You can view logs in the /var/VRTS/install/logs directory.

Removing the agent package using CLI

You can remove a desired agent package using the CLI.

Note: You must remove this agent package from each client system in the cluster.

To remove the agent for Data Guard from a client system

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

  Solaris `# pkgrm`

Removing the agent type definition from the Policy Master system on UNIX

After you remove the agent packages, you can remove the agent type definitions for agents you removed, from the Policy Master system.

To remove the agent type definition from the Policy Master system on UNIX

1 Navigate to the following directory on the client system.

```
# cd /opt/VRTS/install
```

2 Run the following command to remove the agent type definition from the Policy Master system:

```
# ./installagpack -rmtypes
```

3 When the installer prompts, enter the virtual IP address of the Policy Master.

4 Choose whether to remove the type definitions for all the agents or for specific agents. Follow the installer prompts to remove the type definitions.

You can view logs in the /var/VRTS/install/logs directory.

Chapter 3

Configuring the agent for Oracle Data Guard

This chapter includes the following topics:

Configuration concepts for the Oracle Data Guard agent

Review the resource type definition and the attribute definitions for the agents for Oracle Data Guard. The resource type for both the Oracle Data Guard agent and the Oracle Data Guard Broker agent is defined in the OraDGTypes.cf file.

Resource type and attribute definitions for the Data Guard agent

The resource type definition defines the agent in VCS One:

```
<attribute name="LinkRes" type="str" dimension="scalar">
        <insensitive>1</insensitive>
        <must_configure>1</must_configure>
</attribute>
<attribute name="Encoding" type="str" dimension="scalar">
        <insensitive>1</insensitive>
        <default><scalar>""</scalar></default>
</attribute>
<attribute name="Flashback" type="boolean"
```

```
                    dimension="scalar">
            <insensitive>1</insensitive>
            <default><scalar>0</scalar></default>
    </attribute>
    <attribute name="StandbyOnOffline" type="boolean"
                     dimension="scalar">
            <insensitive>1</insensitive>
            <default><scalar>0</scalar></default>
```

Review the description of the agent attributes. You must assign values to the required attributes.

LinkRes	Required attribute Name of the Oracle resource that manages the replicated database instance. Type-dimension: string-scalar
Encoding	Optional attribute Specifies the operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. For example, if Oracle output is in “JAPANESE_JAPAN.JA16EUC,” then “eucJP” is the Solaris value for Encoding. Refer to the Oracle and Solaris documentation for respective encoding values. Type-dimension: integer-scalar The default is ““.
Flashback	Used to enable flashback recovery when failed primary comes up. It is used by flashbackRecover action entry point. The default value is 0.
StandbyOnOffline	Whenever user want to perform migration/switchover of Oracle DataGuard based application, need to set StandbyOnOffline to 1. When StandbyOnOffline is set to 1, Oracle DataGuard Offline EP would ensure to take the database to Standby mode as part of Offline (convert the primary database to physical standby) & reset the StandbyOnOffline back to 0.

Sample configuration for the Data Guard agent

Figure 3-1 shows a sample dependency graph.

VCS service group has a resource of type Data Guard. A second service group contains all necessary resources to control the database instance. The Oracle_DG

group depends on the Oracle_DB group, which is an online local soft group dependency.

Figure 3-1 Dependency graph

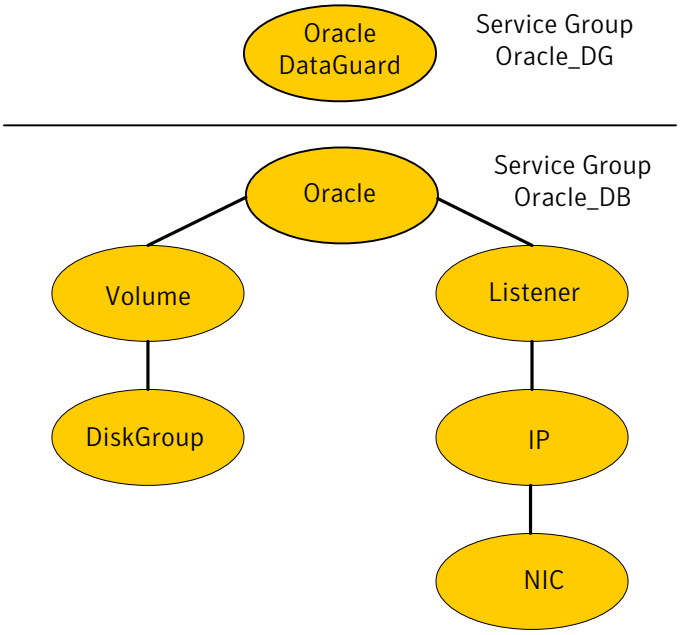

You can configure a resource of type OraDG in the main.cf file:

```
group global_db_rep (
        SystemList = { primary-sys1 = 0, primary-sys2 = 1 }
        ClusterList = { dgclus1 = 0, dgclus2 = 1 }
                )

        OraDG dg_res (
                LinkRes = ora_db_prod
                Flashback = 1
                )

        requires group oradb_prod_SG online local soft




group oradb_prod_SG (
```

```
 SystemList = { primary-sys1 = 0, primary-sys2 = 1 }
 )

        IP lsnr_ip (
         Device = eth0
         Address = "10.209.71.181"
         NetMask = "255.255.252.0"
         )

 LVMLogicalVolume ora_vol (
         LogicalVolume = OraData
         VolumeGroup = VolGroup01
         )

 LVMVolumeGroup ora_grp (
         VolumeGroup = VolGroup01
         )
Mount ora_mnt (
         MountPoint = "/u01"
         BlockDevice = "/dev/mapper/VolGroup01-OraData"
         FSType = ext3
         FsckOpt = "-y"
         )

 NIC lsnr_nic (
         Device = eth0
         )

 Netlsnr ora_db_lsnr (
         Owner = oracle
         Home = "/u01/app/oracle/product/11.2.0/db_1"
         Listener = DGUARD
         )

 Oracle ora_db_prod (
         Sid = dguard
         Owner = oracle
         Home = "/u01/app/oracle/product/11.2.0/db_1"
         StartUpOpt = CUSTOM
         )

 lsnr_ip requires lsnr_nic
 ora_db_lsnr requires lsnr_ip
```

```
ora_db_prod requires ora_mnt
ora_vol requires ora_grp
ora_mnt requires ora_vol
ora_db_prod requires ora_db_lsnr
```

Note the following variations to a standard Oracle database cluster configuration:

- The Oracle resource depends on the Listener resource. The listener process must be already active when the database instance is started because the Data Guard TCP/IP replication links use the Oracle Net Services.
- The IP and NIC resource in the database service group are optional. These resources are only necessary if a cluster on its own protects the primary database. For wide area or site failover, you can implement a transparent network client reconnect.

  To implement a transparent network client reconnect, do one of the following:

  - Use a DNS agent as part of the Data Guard service group
  - Create an alternate Oracle Net Service entries on client machines
- The Oracle resource undergoes an offline-online cycle when promoting a Data Guard standby server to become a primary database. The service group dependency must be soft.
- The name of the Oracle DataGuard resource must be the same in each global cluster configuration. Otherwise, the DemotePri action entry point that is essential for a failover will not work.

Working of Active Physical Standby feature

The Active Data Guard Option available with Oracle Database 11g Enterprise Edition enables you to open a physical standby database for read-only access for reporting, for simple or complex queries, sorting, or Web-based access while Redo Apply continues to apply changes received from the production database. All queries reading from the physical standby database execute in real time, and return current results. With Active Dataguard, you can offload any operation that requires up-to-date, read-only access to the standby database. To support active standby in Oracle Dataguard agent, we have added two action entry points, activateStandby and deactivateStandby.

activateStandby - On physical standby, it mounts the database in Read-only with Redo apply using below SQL commands:

- `ALTER DATABASE RECOVER MANAGED STANDBY DATABASE CANCEL`
- `ALTER DATABASE OPEN READ ONLY`

- `ALTER DATABASE RECOVER MANAGED STANDBY DATABASE USING CURRENT LOGFILE DISCONNECT`

deactivateStandby - It works differently for different version of Oracle. For 11gR2, only oracle resource needs to offline and then online. But, for 11gR1 it needs to send requests to primary database to send redo log to standby as connection goes down after database offline, using `ALTER SYSTEM ARCHIVE LOG CURRENT`SQL command.

Before you configure the agent for Oracle Data Guard

Before you configure the agent, review the following information:

- Review the configuration concepts, which describe the agent's type definition and attributes.
  See "Configuration concepts for the Oracle Data Guard agent" on page 21.
- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
  See "Typical Oracle Data Guard setup in a VCS cluster" on page 10.
- Verify that the clustering infrastructure is in place.
  - If you plan to configure the agent in a global cluster, make sure the global composite service group for the application is configured.
    For more information, see the *Veritas Cluster Server One User's Guide*.

Configuring the agent for Oracle Data Guard

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

- Changing the database startup profile by adding alternate log destination and creating the necessary Oracle net service entries.
- Creating a second complete database copy on the standby server.
- Adding a new service group with at least the Oracle Data Guard agent. The new service group becomes the parent of the existing Oracle database group.

See the Oracle Data Guard documentation for details on how to configure an Oracle database for Data Guard replication.

On Solaris, the Oracle Data Guard agent is zone-aware. You can configure the agent in local zone or global zone.

Chapter 4

Managing and testing clustering support for Oracle Data Guard

This chapter includes the following topics:

Failure scenarios in global clusters

Table 4-1 lists the failure scenarios in a global cluster configuration and describes the behavior of VCS One and the agent in response to the failure.

See the *Veritas Cluster Server One User's Guide* for more information on the DR configurations and the global composite service group attributes.

Table 4-1 Failure scenarios in a global cluster configuration with VCS One agent for Oracle Data Guard

Failure	Description and VCS One response
Application failure	Application cannot start successfully on any client system at the primary site. VCS One response at the secondary site: ■ Causes global composite service group at the primary site to fault and triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. Agent response: ■ Write enables the devices at the secondary site. See "Performing failback after a client system failure or an application failure" on page 32.
Client system failure	All client systems at the primary site fail. VCS One response at the secondary site: ■ Triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. Agent response: ■ Write enables the devices at the secondary site.
Site failure	All PMs, client systems, and their storage at the primary site fail. VCS One response at the secondary site: ■ Triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site. See "Performing failback after a site failure" on page 32.
Replication link failure	Replication link between the arrays at the two sites fails. VCS One response: No action.
Network failure	The network connectivity and the replication link between the sites fail. VCS One response: ■ VCS One at each site concludes that the remote cluster has faulted. ■ The global cluster failover in VCS One is manual. No action. ■ You must confirm the cause of network failure from the cluster administrator at each site and fix the issue.

Table 4-1 Failure scenarios in a global cluster configuration with VCS One agent for Oracle Data Guard *(continued)*

Failure	Description and VCS One response
Storage failure	The array at the primary site fails. VCS One response at the secondary site: ■ Causes the global composite service group at the primary site to fault and triggers a BPA event. ■ The global cluster failover in VCS One is manual. You must bring the global composite service group online at the secondary site.

Failure scenarios for Oracle Data Guard

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

All host or all application failure

If all hosts on the primary side are disabled or if the application cannot start successfully on any primary host, the service group fails over.

In global cluster environments failover requires user confirmation by default.

In global cluster environments, failover requires user confirmation by default. Multiple service groups can fail over in parallel.

Replication link failure

Data Guard detects link failures, monitors the archive logs created on the active primary. When the standby server reconnects to the primary database server, the Data Guard resynchronizes the standby database with all the archive logs. The agent resynchronizes the archive logs since the time of the link failure.

The standby database may not contain the most recent data in the following conditions:

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

However the agent is able to execute a role transition from standby to primary. The database contents at the standby site are always consistent.

After recovery of the replication link, the two replicated databases can be logically inconsistent. The database transactions can result in inconsistency in the following conditions:

- The transactions are committed on the original primary after the link failure, and
- The transactions are never replicated to the standby at the time of takeover on the original primary after the link failure

You can get both sites back into a consistent state only if Oracle flash recovery was enabled at both primary and standby database servers. Otherwise, a restart from the last consistent backup can be necessary.

Split-brain in a Data Guard environment

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

You must resynchronize the databases manually either by using flashback information or the archive logs. Similar to a replication link failure, a complete restart from a backup copy might be necessary.

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

Testing the global composite service group migration

After you configure the VCS One agent for Oracle Data Guard, verify that the global composite service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global composite service group migration in global cluster setup

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

Perform the following steps:

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

```
hacsg -switch global_csg -clus secondary_clusname
```

VCS One brings the global composite service group online at the secondary site.

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

Perform the following steps:

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

```
hacsg -switch global_csg -clus primary_clusname
```

VCS One brings the global composite service group online at the primary site.

Testing disaster recovery after site failure

Review the details on site failure and how VCS One and the agent for Oracle Data Guard behave in response to the failure.

See "Failure scenarios in global clusters" on page 27.

Depending on the DR configuration, perform one of the following procedures to test the disaster recovery in the event of site failure.

To test disaster recovery for site failure in global cluster setup

1 Halt all client systems and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

You must bring the global composite service group online at the secondary site. Run the following command:

```
hacsg -online -force global_csg
```

2 Verify that the global composite service group is online at the secondary site.

```
hacsg -state global_csg
```

Testing disaster recovery after client system failure

Review the details on client system failure and how VCS One and the agent for Oracle Data Guard behave in response to the failure.

See "Failure scenarios in global clusters" on page 27.

Depending on the DR configuration, perform one of the following procedures to test how VCS One recovers after all client systems at the primary site fail.

To test disaster recovery for client system failure in global cluster setup

1 Halt the client system at the primary site.

You must bring the global composite service group online at the secondary site. Run the following command:

```
hacsg -online -force global_csg
```

2 Verify that the global composite service group is online at the secondary site.

```
hacsg -state global_csg
```

Performing failback after a client system failure or an application failure

Review the details on client system failure and application failure and how VCS One and the agent for Oracle Data Guard behave in response to these failures.

See "Failure scenarios in global clusters" on page 27.

After the client systems at the primary site are restarted, you can perform a failback of the global composite service group to the primary site. Depending on your DR configuration, perform one of the following procedures.

To perform failback after a client system failure or an application failure in global cluster

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

```
hacsg -switch <global_csg> -clus <remote_clusname>
```

VCS One brings the global composite service group online at the primary site.

Performing failback after a site failure

See "Failure scenarios in global clusters" on page 27.

To perform failback after a site failure in global cluster

1. Take the global composite service group offline at the secondary site. At the secondary site, run the following command:

   ```
   hacsg -offline global_csg
   ```

2. Bring the global composite service group online at the primary site. Run the following command:

   ```
   hacsg -online global_csg
   ```

Index

T