

Cluster Server Agent for IBM DB2 HADR Installation and Configuration Guide

AIX, Linux, Solaris

5.0

Cluster Server Agent for IBM DB2 HADR Installation and Configuration Guide

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

Agent Version: 5.0

Document version: 5.0 Rev 0

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, the Veritas Logo, InfoScale, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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. SYMANTEC CORPORATION SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software 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, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. 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.

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 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

Technical Support	3
Chapter 1 Introducing the agent for IBM DB2 HADR	7
About the agent for IBM DB2 HADR	7
Supported software	8
Typical IBM DB2 HADR setup in a VCS cluster	8
DB2 HADR agent functions	9
About the DB2 HADR agent's online function	10
About DB2 HADR role transition	11
Chapter 2 Installing and removing the agent for IBM DB2 HADR	12
Before you install the agent for IBM DB2 HADR	12
Installing the agent for IBM DB2 HADR	12
Upgrading the agent for IBM DB2 HADR	13
Removing the agent for IBM DB2 HADR	14
Chapter 3 Configuring the agent for IBM DB2 HADR	15
Configuration concepts for the IBM DB2 HADR agent	15
Resource type definition for the DB2 HADR agent	15
Attribute definitions for the DB2 HADR agent	16
Sample configuration for the DB2 HADR agent	16
Before you configure the agent for IBM DB2 HADR	18
About cluster heartbeats	18
About configuring system zones in replicated data clusters	18
About preventing split-brain	19
Configuring the agent for IBM DB2 HADR	20
Configuring the agent manually in a global cluster	20
Index	22

Introducing the agent for IBM DB2 HADR

This chapter includes the following topics:

- [About the agent for IBM DB2 HADR](#)
- [Supported software](#)
- [Typical IBM DB2 HADR setup in a VCS cluster](#)
- [DB2 HADR agent functions](#)

About the agent for IBM DB2 HADR

The Cluster Server agent for IBM DB2 HADR provides failover support and recovery. The agent provides this support in an environment that uses DB2 HADR to replicate data between DB2 databases. It monitors and manages the state of a replicated DB2 database that runs on VCS nodes. For a system where the DB2 HADR resource is online, the agent replicates the DB2 database transactions from a primary to a standby database server.

You can use the agent in replicated data clusters and in global clusters that run VCS.

DB2 UDB does not support partitioned databases (Database Partitioning Feature) and as such the DB2 HADR agent obeys the same limitation.

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 information on the software versions that the agent for IBM DB2 HADR supports, see the Symantec Operations Readiness Tools (SORT) site:
<https://sort.symantec.com/agents>.

Typical IBM DB2 HADR setup in a VCS cluster

Figure 1-1 displays a typical cluster setup in a DB2 HADR environment.

Figure 1-1 Typical clustering setup for the agent



Clustering in a DB2 HADR 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.
See “[About cluster heartbeats](#)” on page 18.

DB2 HADR agent functions

The VCS enterprise agent for DB2 UDB HADR monitors and manages the state of a replicated DB2 database that is executed on VCS nodes. Agent functions bring resources online, take them offline, and perform different monitoring actions.

The agent functions are as follows:

online	<p>Executes different sequences, depending on the value of the agent attribute HADROper (LINK or SWITCH).</p> <p>See “About the DB2 HADR agent's online function” on page 10.</p>
offline	<p>Removes the lock file on the local node. A switch of the replication direction—promoting the standby and demoting the primary—happens on the target node. No DB2 UDB reconfiguration occurs as part of the offline function. In case of a complete shutdown, a standard Db2udb resource closes the database and stops the instance processes.</p>
monitor	<p>If HADROper is set to SWITCH:</p> <ul style="list-style-type: none">■ 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. <p>If HADROper is set to LINK, the agent reports online in the following ways:</p> <ul style="list-style-type: none">■ The database is in PRIMARY or STANDBY role■ The database uses the output from the <i>db2 GET SNAPSHOT FOR ALL ON database_name</i> command
open	<p>Creates the lock file:</p> <pre>\$VCS_HOME/bin/Db2HADR/.ResourceID_LCK.dont_remove.</pre> <p>The entry point creates this file in the local agent directory if the database role as returned by db2 get snapshot in one of the following ways:</p> <ul style="list-style-type: none">■ PRIMARY■ STANDBY with HADROper attribute value set to LINK <p>This file creation is necessary to report the correct state when VCS starts after an established HADR link.</p>
clean	<p>Removes the lock file for the following resource states: OFFLINE TIMEOUT / INEFFECTIVE, ONLINE TIMEOUT, UNEXPECTED OFFLINE, OR MONITOR HUNG.</p>

info	Reports the current state and role of the database using the <i>db2 GET SNAPSHOT FOR ALL ON database_name</i> command.
actions/HADRStatus	Extracts the current state and role of the database from a <i>db2 GET SNAPSHOT FOR ALL ON database_name</i> command.
actions/HADRConfig	Extracts the database configuration information from a <i>db2 GET DATABASE CONFIGURATION FOR database_name</i> command.
actions/PreSwitch	Reports the state of the replication link to verify if a replication-link redirection is possible without data loss.

About the DB2 HADR agent's online function

The online operation can equal either LINK or SWITCH.

HADROper is LINK

The agent does not execute commands if the database snapshot *db2 GET SNAPSHOT FOR ALL ON database_name* returns a PRIMARY or STANDBY database role. It does not execute commands because the HADR link exists.

The agent starts the database after verifying the following criteria:

- The database configuration *db2 get db cfg for database_name* returns PRIMARY or STANDARD for the HADR database role.
- Rollforward pending is NO.

The agent then issues the *db2 START HADR ON DATABASE database_name AS PRIMARY* command to start the database.

If the temporary attribute UseForce value is 1, the agent appends the words **BY FORCE** to this db2 command string. It also resets the value of the temp attribute to 0. This feature enables a database server to start in primary mode without contacting the assigned standby server.

Without force, the *db2 HADR START* command returns SQL1768N if the following criteria are met:

- The standby is not available or accepting a connection, and
- The database is not opened for access

The agent issues the command *db2 START HADR ON DATABASE database_name AS STANDBY* if the following criteria are met:

- The database configuration `db2 get db cfg for database_name` returns `STANDBY` or `STANDARD` for the HADR database role, and
- The value for rollforward pending is `DATABASE`

To mark the completed online sequence, the agent creates a lock file on the PRIMARY HOST: `$VCS_HOME/bin/Db2HADR/.ResourceID_PRI_LCK.dont_remove..`

HADROper is SWITCH

The agent assumes a failover or group switch condition if the database snapshot `db2 GET SNAPSHOT FOR ALL ON database_name` returns `STANDBY`. `STANDBY` is for the database role.

The agent executes the `db2 TAKEOVER HADR ON DATABASE database_name BY FORCE` command if the either of the following criteria are met:

- The remote cluster node `hasys -state remote_node` is `FAULTED`, or
- The value for the temporary resource attribute `UseForce` is `1`

If the `PeerWindowOnly` attribute is set to `1`, the agent appends the words `PEER WINDOW ONLY` to the `db2` command string. This feature helps ensure data consistency.

Note: The `PeerWindowOnly` attribute is only applicable to DB2 9.5 and later.

In any other case, the agent issues the `db2` command without the `FORCE` option. It issues the command because the internal DB2 HADR TCP/IP link can still perform a graceful role transition between the two sites.

To mark the completed online sequence, the agent creates a lock file:

`$VCS_HOME/bin/Db2HADR/.ResourceID_LCK.dont_remove..`

About DB2 HADR role transition

The agent supports the DB2 HADR Role Transition. In this case, you can use the `hagrp -switch` command to switch the service group where the DB2 HADR `SWITCH` resource resides.

The agent then performs the following actions:

- It takes the resource of type `Db2HADR` with `HADROper = SWITCH` offline on the original primary and removes the lock file.
- It brings the resource of type `Db2HADR` with `HADROper = SWITCH` online on the former standby. It also alters the database role from standby to primary using the `db2 TAKEOVER` command.

Installing and removing the agent for IBM DB2 HADR

This chapter includes the following topics:

- [Before you install the agent for IBM DB2 HADR](#)
- [Installing the agent for IBM DB2 HADR](#)
- [Upgrading the agent for IBM DB2 HADR](#)
- [Removing the agent for IBM DB2 HADR](#)

Before you install the agent for IBM DB2 HADR

Before you install the Cluster Server agent for IBM DB2 HADR, 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*.

See [“Typical IBM DB2 HADR setup in a VCS cluster”](#) on page 8.

Installing the agent for IBM DB2 HADR

You must install the IBM DB2 HADR agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

These instructions assume that you have already installed VCS.

To install the agent 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.

```
AIX      cd1/aix/vcs/replication/DB2_HADR_agent/  
agent_version/pkgs/
```

```
Linux    cd1/linux/generic/vcs/replication/DB2_HADR_agent/  
agent_version/rpms/
```

```
Solaris  cd1/solaris/dist_arch/vcs/replication/DB2_HADR_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.

```
AIX      # installp -ac -d VRTScsdbh.rte.bff VRTScsdbh.rte
```

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

```
Solaris  # pkgadd -d . VRTScsdbh
```

Upgrading the agent for IBM DB2 HADR

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

To upgrade the agent software

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

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

- 2 Remove the agent from the node.

See [“Removing the agent for IBM DB2 HADR”](#) on page 14.

- 3 Delete the file /etc/VRTSvcs/conf/config/Db2HADRTypes.cf.

- 4 Install the current version of the agent.

See [“Installing the agent for IBM DB2 HADR”](#) on page 12.

- 5 Copy the file Db2HADRTypes.cf from the directory /etc/VRTSvcs/conf/ to the /etc/VRTSvcs/conf/config directory.

- 6 Repeat step 2 through step 5 on each node.

- 7 From a node in the cluster, edit your configuration file /etc/VRTSvcs/conf/config/main.cf.

Configure the new attributes, if applicable.

- 8 Verify the configuration by running the following command:

```
# hacf -verify config
```

- 9 Start VCS on local node first.

- 10 Start VCS on other nodes.

Removing the agent for IBM DB2 HADR

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:

```
AIX          # installp -u VRTScsdbh.rte
```

```
Linux        # rpm -e VRTScsdbh
```

```
Solaris      # pkgrm VRTScsdbh
```

Configuring the agent for IBM DB2 HADR

This chapter includes the following topics:

- [Configuration concepts for the IBM DB2 HADR agent](#)
- [Before you configure the agent for IBM DB2 HADR](#)
- [Configuring the agent for IBM DB2 HADR](#)

Configuration concepts for the IBM DB2 HADR agent

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

Resource type definition for the DB2 HADR agent

The resource type definition defines the agent in VCS.

```
type Db2HADR (
    static keylist SupportedActions = { HADRStatus, HADRConfig,
    PreSwitch }
    static int OnlineTimeout = 600
    static str ArgList[] = { Db2dbRes, DBName, HADROper, Encoding }
    str Db2dbRes
    str DBName
    str HADROper
    str Encoding
    boolean PeerWindowOnly = 0
```

```
temp boolean UseForce = 0
)
```

Attribute definitions for the DB2 HADR agent

Review the description of the agent attributes.

Required attributes

You must assign values to required attributes.

Db2dbRes	<p>Name of theDb2udb instance resource that manages the replicated database instance.</p> <p>Type-dimension: string-scalar</p>
DBName	<p>Name of the replicated database</p> <p>Type-dimension: string-scalar</p>
HADROper	<p>Agent operational mode [LINK SWITCH].</p> <p>Type-dimension: string-scalar</p>
PeerWindowOnly	<p>Specifies whether the agent must use the PEER WINDOW ONLY option when the agent is performing a takeover operation by force.</p> <p>Set the value of this attribute to 1 if you want the agent to use the PEER WINDOW ONLY option.</p> <p>Default: 0</p> <p>Type-dimension: boolean-scalar</p>

Optional attributes

Set optional attributes as needed.

Encoding	<p>Specifies the operating system encoding that corresponds to DB2 UDB encoding for the displayed DB2 output.</p> <p>Default: ""</p> <p>Type-dimension: integer-scalar</p>
----------	--

Sample configuration for the DB2 HADR agent

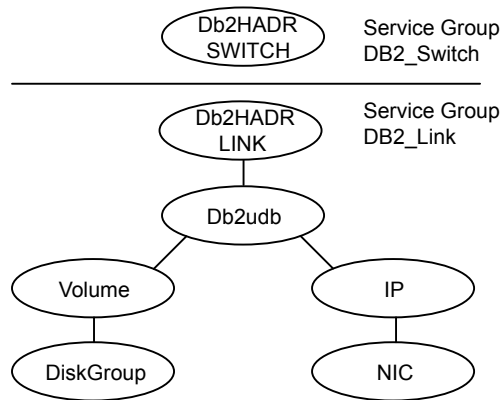
This is an example of a VCS service group that has a resource of type Db2HADR with attribute HADROper = SWITCH. The second child service group contains all

necessary resources to control the database instance and establish the HADR replication link. Here, the resource of type Db2HADR uses the attribute HADROper = LINK.

The DB2_Link group with the online local hard attribute depends on the group DB2_Switch.

Figure 3-1 shows a VCS service group that has a resource of type Db2HADR with attribute HADROper = SWITCH.

Figure 3-1 Dependency tree



You can configure a resource of type Db2HADR LINK in the `main.cf` file:

```

Db2HADR iabsolut_lnk (
    Critical = 0
    Db2dbRes = db2inst1
    DBName = iabsolut
    HADROper = LINK
)
    
```

You can configure a resource of type Db2HADR SWITCH in the `main.cf` file:

```

Db2HADR iabsolut_sw (
    Critical = 0
    Db2dbRes = db2inst1
    DBName = iabsolut
    HADROper = SWITCH
)
    
```

Note the following variations to a standard DB2 database cluster configuration:

- The Db2udb resource does not contain a reference to a specific database.
- The IP and NIC resource in the database service group are optional. These resources are only necessary if a local cluster on its own protects the primary database. For wide area or site failover, you can implement a transparent network client reconnect using a DNS agent as part of the HADR_Switch service group. You can also create alternate database connect entries on client machines.

Before you configure the agent for IBM DB2 HADR

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
See [“Typical IBM DB2 HADR setup in a VCS cluster”](#) on page 8.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 18.
- Set up an effective heartbeat mechanism to prevent split-brain.
See [“About preventing split-brain”](#) on page 19.

About cluster heartbeats

In global clusters, the VCS Heartbeat agent should send heartbeats directly across the same network that is used for database replication. With this setup, the detection of a site loss only occurs if the database replication is interrupted at the same time. No additional network addresses are necessary for the configuration.


About configuring system zones in replicated data clusters

In a replicated data cluster, you can prevent unnecessary DB2 HADR failover or fallback by creating system zones. VCS attempts to fail over applications within the same system zone before failing them over across system zones.

Configure the hosts that are attached to the same database as part of the same system zone.

[Figure 3-2](#) depicts a sample configuration where hosta and hostb are in one system zone, and hostc is in another system zone.

Figure 3-2 Example system zone configuration



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

The IBM DB2 HADR does not support protecting the standby database server with a second node.

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.

In a replicated database cluster, VCS attempts to start the application assuming a total disaster because the primary site is unreachable. Once the heartbeats are restored, VCS stops the applications on one side and restarts the VCS engine (HAD). It restarts HAD to eliminate concurrency violation of the same group being online at two places simultaneously. Administrators must resynchronize the databases manually when DB2 detects that the transaction serial numbers are

inconsistent. 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. If a fail over mistakenly occurs, the situation is similar to the replicated data cluster case. However, when the heartbeat is restored, VCS does not stop HAD at either site. VCS forces you to choose which group to take offline. You must resynchronize the data manually.

Configuring the agent for IBM DB2 HADR

You can configure clustered applications in a disaster recovery environment by:

- Changing the database configuration profile by adding the necessary HADR configuration attributes.
- Creating a second complete database copy on the standby server.
- Adding a new resource of type AgentTypes with attribute HADROper = LINK. This resource starts the HADR replication and opens the database.
The AgentTypes resource is parent of a standard Db2udb resource, whose database attribute should be blank. The resource of type Db2udb is only used to start the instance process framework, not to open the database and make it accessible. That is to say, the AgentTypes parent resource opens the database and makes it accessible depending on the database state.
- Adding a new service group with at least one resource of type AgentTypes with attribute HADROper = SWITCH. This group becomes parent of the existing DB2 database group.

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

See [“Sample configuration for the DB2 HADR agent”](#) on page 16.

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 (Db2HADR) is not added to your configuration, add it. From the Cluster Explorer **File** menu, choose **Import Types**, and select:
`/etc/VRTSvc/conf/Db2HADRTypes.cf`

- 3** Click **Import**.
- 4** Save the configuration.
- 5** Update the Db2udb resource so that the IndepthMonitor and DatabaseName attributes are empty (""). This resource only handles the DB2 database manager start and stop. A resource of type AgentTypes handles the database open.
- 6** Create a new resource of type AgentTypes as part of the Db2udb resource group.
- 7** Set attributes Db2dbRes to the name of the Db2udb resource, DBName to the database being replicated, and HADROper to LINK.
- 8** Make the AgentTypes resource dependent on the Db2udb resource.
- 9** Add a new group with at least one resource of type AgentTypes.
- 10** Set the attributes Db2dbRes to the name of the Db2udb resource, DBName to the database being replicated, and HADROper to SWITCH.
- 11** Create an online local hard group dependency between the new AgentTypes group and the existing DB2 database group.
- 12** Change the ClusterFailOverPolicy from the default, if necessary. Symantec recommends keeping the Manual default, to minimize the chance of failing over on a split-brain.
- 13** Perform step [5](#) through step [12](#) for each existing DB2 database service group in each cluster that uses replicated data.

Index

A

- actions/HADRConfig 9
- actions/HADStatus 9
- agent functions
 - actions/HADRConfig 9
 - actions/HADStatus 9
 - clean 9
 - info 9
 - monitor 9
 - offline 9
 - online 9
 - open 9
- attribute definitions 16

C

- clean 9
- cluster
 - heartbeats 18
- configuring
 - sample configuration 16

D

- Db2dbRes attribute 16
- DBName attribute 16

H

- HADROper
 - LINK 10
 - SWITCH 11
- HADROper attribute 16

I

- IBM DB2 HADR agent
 - attribute definitions 16
- info 9
- installing the agent
 - AIX systems 12
 - Linux systems 12
 - Solaris systems 12

L

- LINK 10

M

- monitor 9

O

- offline 9
- online 9
- open 9
- optional attributes
 - Encoding attribute 16

R

- required attributes
 - Db2dbRes 16
 - DBName 16
 - HADROper 16
- resource type definition 15
- role transition 11

S

- sample configuration 16
- split-brain
 - handling in cluster 19
- SWITCH 11

U

- uninstalling the agent
 - AIX systems 14
 - Linux systems 14
 - Solaris systems 14