

Veritas Cluster Server Agent for IBM DB2 HADR Installation and Configuration Guide

AIX, Linux, Solaris

5.0

Veritas Cluster Server Agent for IBM DB2 HADR Installation and Configuration Guide

Copyright © 2006 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

AIX is a registered trademark of IBM Corporation.
Linux is a registered trademark of Linus Torvalds.
Solaris is a trademark of Sun Microsystems, Inc.

Licensing and registration

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit <http://support.veritas.com> and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1	Introducing the IBM DB2 HADR VCS agent	
	About the IBM DB2 HADR agent	8
	Supported software and hardware	8
	Typical DB2 HADR setup in a VCS cluster	9
	DB2 HADR agent functions	10
	About the DB2 HADR agent's online function	11
	HADROper is LINK	11
	HADROper is SWITCH	12
Chapter 2	Installing and removing the agent	
	Before you install the DB2 HADR agent	14
	Installing the agent for DB2 HADR	14
	Removing the agent for DB2 HADR	15
Chapter 3	Configuring the DB2 HADR agent	
	Configuration concepts	18
	DB2 HADR resource type definition	18
	Attribute definitions for the DB2 HADR agent	19
	Required attributes	19
	Optional attributes	19
	Sample configuration	20
	Before you configure the DB2 HADR agent	22
	About cluster heartbeats	22
	About configuring system zones in replicated data clusters	22
	About preventing split-brain	23
	About DB2 HADR role transition	23
	Configuring the agent	24
	Configuring the agent in a global cluster	24
Chapter 4	Managing and testing clustering support for DB2 HADR	
	Typical test setup	28
	Testing service group migration	28
	Testing host failure	29
	Failure scenarios	29

All host or all application failure 29
Replication link failure 30

Index 31

Introducing the IBM DB2 HADR VCS agent

This chapter contains the following topics:

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

About the IBM DB2 HADR agent

The VCS enterprise agent for DB2 HADR from Symantec provides failover support and recovery in environments using 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 global clusters. DB2 UDB does not support partitioned databases (Database Partitioning Feature) and as such the DB2 HADR agent obeys the same limitation.


Supported software and hardware

The DB2 HADR agent supports VCS 5.0. The agent supports DB2 UDB version 8.2 and later. This version is the first version of DB2 UDB that has the HADR replication option. The agent also supports only a single standby database instance per configured primary database.

Typical DB2 HADR setup in a VCS cluster

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

Figure 1-1 Clustering in a DB2 environment



Clustering in a DB2 HADR environment typically consists of the following hardware infrastructure:

- The primary database instance sends redo data across a TCP/IP link to a standby database instance. A local cluster can protect the primary database to increase the availability.
- The standby database instance applies the redo information to a physical copy of the primary database.
- The primary and standby sites have to be connected through a single TCP/IP network connection. This link can be shared with VCS global cluster option heartbeat communication.
- Network heartbeats that use LLT or TCP/IP between the two data centers to determine their health.

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

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. Agent functions are also known as entry points.

online	<p>Depending on the value of the agent attribute HADROper (LINK or SWITCH), the entry point executes different sequences.</p> <p>See “About the DB2 HADR agent’s online function” on page 11.</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 is responsible to close the database and stop the instance processes.</p>
monitor	<p>Verifies that the lock file exists.</p> <p>If the lock file does not exist, the monitor entry point reports the status of the resource as offline.</p> <p>If the lock file exists, the agent report online when:</p> <ul style="list-style-type: none">■ the database is in role PRIMARY■ and uses the output from <code>db2 GET SNAPSHOT FOR ALL ON database_name</code>. <p>Monitor performs the same action if the HADROper attribute is set to LINK and the snapshot returns the database role STANDBY.</p>
open	<p>Creates the lock file: <code>\$VCS_HOME/bin/Db2HADR/.ResourceID_LCK.dont_remove</code>.</p> <p>Creates this file in the local agent directory if the database role as returned by <code>db2 get snapshot</code> is either:</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	<p>Reports the current state and role of the database using the <code>db2 GET SNAPSHOT FOR ALL ON database_name</code> command.</p>

- actions/HADStatus Extracts the current state and role of the database from a `db2 GET SNAPSHOT FOR ALL ON database_name` command.
- actions/HADRConfig Extracts the database configuration information from a `db2 GET DATABASE CONFIGURATION FOR database_name` command.

About the DB2 HADR agent's online function

The online function can equal either: **LINK** or **SWITCH**.

HADROper is LINK

The agent executes no commands if the database snapshot `db2 GET SNAPSHOT FOR ALL ON database_name` returns a PRIMARY or STANDBY database role. It executes no commands because the HADR link exists.

The agent starts the database after verifying that:

- 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 the start of a database server in primary mode without contacting the assigned standby server.

Without force, the `db2 HADR START` command returns SQL1768N if:

- 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 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: `$VCS_HOME/bin/Db2HADR/.ResourceID_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 remote cluster node `hasys -state remote_node` is FAULTED, or
- the value for the temporary resource attribute UseForce is 1

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.`

Installing and removing the agent

This chapter contains the following topics:

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

Before you install the DB2 HADR agent

Set up your cluster. For information about installing and configuring VCS, see the *Veritas Cluster Server Installation Guide*.

Set up replication and the required hardware infrastructure.

See “[Typical DB2 HADR setup in a VCS cluster](#)” on page 9

Installing the agent for 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 on AIX systems

- 1 Determine the device access name of the disc drive.

```
# cd /dev
# lsdev -C -c cdrom
```

The output resembles:

```
cd0 Available 10-60-00-4,0 16 Bit SCSI Multimedia CD-ROM Drive
```

In this example, the CD device access name is cd0.

- 2 Insert the disc into the system's drive.

- 3 Mount the disc.

```
# mkdir -p /cdrom
# mount -V cdrfs -o ro /dev/cd0 /cdrom
```

- 4 Add the filesets for the software.

```
# cd /cdrom
# installp -ac -d /cdrom/aix/replication/db2hadr_agent\
/5.0.00.0/pkgs/VRTScsdbh.rte.bff VRTScsdbh
```

To install the agent on Linux systems

- 1 Log in as superuser.
- 2 Insert the disc into the system's drive.
- 3 Mount the disc, if the disc does not automatically mount.

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

- 4 Navigate to the /mnt directory.

```
# cd /mnt/cdrom
```

- 5 Navigate to the location of the agent package.

- For all supported Red Hat distributions:

```
# cd/linux/linux/redhatlinux/replication/db2hadr_agent\
/5.0.00.0/rpms
```

- For all supported Suse distributions:

```
# cd linux/linux/suselinux/replication/db2hadr_agent\  
/5.0.00.0/rpms
```

- 6 Install the agent software:

```
# rpm -ivh agentrpm
```

The variable *agentrpm* represents the agent package in the rpms directory.

To install the agent on Solaris systems

- 1 Insert the disc into the system's drive.

```
# cd /cdrom/cdrom0
```

- 2 Navigate to the location of the agent package.

```
# cd solaris/sparc/replication/db2hadr_agent/5.0.00.0/pkgsrc/
```

- 3 Install the agent binaries.

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

Removing the agent for DB2 HADR

Before you attempt to remove the agent, make sure the application service group is not online. You must remove the agent from each node in the cluster.

To remove the agent from an AIX cluster

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

```
# installp -u VRTScsdbh
```

To remove the agent from a Linux cluster

- ◆ Type one of the following commands on each system to remove the agent

```
RHEL 4 # rpm -e VRTScsdbh-5.0-GA_RHEL4.i686.rpm
```

```
SLES 9 # rpm -e VRTScsdbh-5.0-GA_SLES9.i586.rpm
```

To remove the agent from a Solaris cluster

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

```
# pkgrm VRTScsdbh
```


Configuring the DB2 HADR agent

This chapter contains the following topics:

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

Configuration concepts

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

DB2 HADR resource type definition

The resource type defines the DB2 HADR agent.

```
type Db2HADR (  
  static keylist SupportedActions = { HADRStatus, HADRConfig }  
  static int OnlineTimeout = 600  
  static str ArgList[] = { Db2dbRes, DBName, HADROper,  
    Encoding }  
  str Db2dbRes  
  str DBName  
  str HADROper  
  str Encoding  
  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	Name of theDb2udb instance resource that manages the replicated database instance. Type-dimension: string-scalar
DBName	Name of the replicated database Type-dimension: string-scalar
HADROper	Agent operational mode [LINK SWITCH]. Type-dimension: string-scalar

Optional attributes

Set optional attributes as needed.


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

Sample configuration

The following dependency graph shows 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 Dependency tree



You can configure a resource of type Db2HADR LINK as follows 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 DB2 HADR agent

Before you configure the agent:

- Review the configuration concepts, which describe the agent's type definition and attributes.
See "[Configuration concepts](#)" on page 18.
- Verify that the agent is installed on all systems in the cluster.
- Verify the hardware setup for the agent.
See "[Typical DB2 HADR setup in a VCS cluster](#)" on page 9.
- Make sure the cluster has an effective heartbeat mechanism in place.
See "[About cluster heartbeats](#)" on page 22.
- Set up system zones in replicated data clusters.
See "[About configuring system zones in replicated data clusters](#)" on page 22
- Set up an effective heartbeat mechanism to prevent split-brain.
See "[About preventing split-brain](#)" on page 23


About cluster heartbeats

In global clusters, the VCS Heartbeat agent should send heartbeats directly across the same network that is used for the 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 failback by creating system zones. VCS attempts to fail over applications within the same system zone before failing them over across system zones. So, configure the hosts that are attached to the same database as part of the same system zone.

Figure 3-1 depicts a sample configuration where `hosta` and `hostb` are in one system zone, and `hostc` is in another system zone. Use the `SystemZones` attribute to create these zones. 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. Minimize the effects of split-brain by ensuring the cluster heartbeat links pass through similar physical infrastructure as the replication links so that if one breaks, so does the other.

About DB2 HADR role transition

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

The agent then:

- takes the resource of type `Db2HADR` with `HADROper = SWITCH` offline on the original primary and removes the lock file.

- 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.

Configuring the agent

You can adapt most DB2 database environments in VCS to 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 Db2HADR with attribute HADROper = LINK. This resource starts the HADR replication and opens the database. The Db2HADR 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 Db2HADR 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 Db2HADR with attribute HADROper = SWITCH. This group becomes parent of the existing DB2 database group.

Configuring the agent in a global cluster

If your configuration does not have the agent resource type (Db2HADR), you need to add it. You must perform [step 4](#) through [step 12](#) for each existing DB2 database service group in each cluster that uses replicated data.

To configure the agent in a global cluster

- 1 Start the Cluster Manager and log on to the cluster.
- 2 Click **Import**.
- 3 Save the configuration.
- 4 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 Db2HADR handles the database open.
- 5 Create a new resource of type Db2HADR as part of the Db2udb resource group.

- 6 Set attributes Db2dbRes to the name of the Db2udb resource, DBName to the database being replicated, and HADROper to LINK.
- 7 Make the Db2HADR resource dependent on the Db2udb resource.
- 8 Add a new group with at least one resource of type Db2HADR.
- 9 Set the attributes Db2dbRes to the name of the Db2udb resource, DBName to the database being replicated, and HADROper to SWITCH.
- 10 Create an online local hard group dependency between the new Db2HADR group and the existing DB2 database group.
- 11 Configure the Db2HADR service group using the Global Group Configuration Wizard as a global group. See the *VERITAS Cluster Server User's Guide* for more information.
- 12 Change the ClusterFailOverPolicy from the default, if necessary. Veritas recommends keeping the default, which is Manual, to minimize the chance of failing over on a split-brain.

Managing and testing clustering support for DB2 HADR

This chapter contains the following topics:

- [Typical test setup](#)
- [Testing service group migration](#)
- [Testing host failure](#)
- [Failure scenarios](#)


Typical test setup

A typical test environment includes:

- Two hosts (hosta and hostb) attached to a shared storage enclosure.
- Another host (hostc) attached to another local storage device.
- The application runs on hosta, which has access to the shared storage device.
- A replicated data cluster with two heartbeat links; a global cluster with one network heartbeat.

Figure 4-1 depicts a typical test environment.

Figure 4-1 Typical test setup



Testing service group migration

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

To perform the service group migration test

- 1 Migrate the service group to a host that is attached to the shared storage device. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 2 Click **Switch To**, and click the system that is attached to the same storage device (hostb) from the menu.

- 3 Migrate the service group to a host that attached to a different storage device. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 4 Click **Switch To**, and click the system that is attached to the another array (hostc) from the menu.
- 5 Click **Switch To**, and click the system on which the group was initially online (hosta).

Testing host failure

In this scenario, the host where the application runs is lost. Eventually all the hosts in the system zone or cluster are lost.

To perform the host failure test

- 1 Halt the or shut down host where the application runs (hosta).
- 2 Halt or shut down hostb.
In a replicated data cluster, the group fails over to hostc.
In a global cluster, a cluster down alert appears and gives you the opportunity to fail over the service group manually.
- 3 Switch the service group to its original host when VCS starts. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the service group.
- 4 Click **Switch To**, and click the system on which the service group was initially online (hosta).

Failure scenarios

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 replicated data cluster environments, the failover can be automatic, whereas in global cluster environments failover requires user confirmation by default.

Replication link failure

DB2 HADR detects link failures, monitors the archive logs created on the active primary. When the standby server reconnects to the primary database server, the DB2 HADR 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

To return both sites back into a consistent state is only possible if the former primary is completely updated from the promoted standby. Otherwise, a restart from the last consistent backup or a storage checkpoint like a Symantec Volume Manager or File System snapshot is necessary.

Index

A

- about the DB2 UDB HADR agent 8
- actions/HADRConfig 11
- actions/HADStatus 11
- agent operations 10
 - actions/HADRConfig 11
 - actions/HADStatus 11
 - clean 10
 - info 10
 - monitor 10
 - offline 10
 - online 10
 - open 10
- application failure 29
- attribute definitions 19

C

- clean 10
- cluster heartbeats 22
- configuring 17
 - sample configuration 20

D

- DB2 7
- DB2 UDB HADR Agent
 - about 8
- Db2dbRes attribute 19
- DBName attribute 19

E

- entry points 10

F

- Failure 27
- failure scenarios
 - all application failure 29
 - all host failure 29
 - replication link failure 30

H

- HADROper
 - LINK 11
 - SWITCH 12
- HADROper attribute 19
- heartbeats 22
- heartbeats, cluster 22
- host failure 29

I

- IBM 7
- IBM DB2 HADR agent
 - about 8
 - attribute definitions 19
 - operations 10
- info 10
- Installation 1
- installing 13
- installing the agent
 - AIX systems 14
 - Linux systems 14
 - Solaris systems 15

L

- LINK 11

M

- managing DB2 HADR 27
- migrating service group 28
- monitor 10

O

- offline 10
- online 10
- open 10
- operations 10
- optional attributes
 - Encoding attribute 19

R

- Removing 13
- replication link failure 30
- required attributes
 - Db2dbRes 19
 - DBName 19
 - HADROper 19
- resource type definition 18
- role transition 23

S

- sample configuration 20
- service group, migrating 28
- split-brain, handling in cluster 23
- supported
 - hardware 8
 - software 8
- SWITCH 12
- Symantec 8

T

- testing DB2 HADR 27

U

- uninstalling the agent
 - AIX systems 15
 - Linux systems 15
 - Solaris systems 15