

VERITAS Cluster Server Enterprise Agent 4.1.01 for IBM DB2 HADR

Installation and Configuration Guide

AIX, Linux, Solaris

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

VERITAS Legal Notice

Copyright © 2005 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS Logo, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS and the VERITAS logo, Reg. U.S. Pat. & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000 Fax 650-527-2908
www.veritas.com



Contents

Preface	v
About this book	v
How this guide is organized	v
VERITAS Cluster Server documentation	vi
Conventions	vi
Getting help	vii
Documentation feedback	vii
Chapter 1. Introduction	1
About the IBM DB2 UDB HADR agent	1
Supported software and hardware	1
Typical setup	2
Agent functions	3
online	3
offline	4
monitor	4
open	4
clean	5
info	5
actions/HADStatus	5
actions/HADRConfig	5



Chapter 2. Installing and removing the agent	7
Installing the agent	7
Removing the agent	9
Chapter 3. Configuring the agent	11
Configuration concepts	11
Resource type definition	11
Attribute definitions	12
Sample configuration	12
About cluster heartbeats	14
Individual component failure	14
All host or all application failure	15
Replication link failure	15
Split-brain	15
About DB2 HADR role transition	16
Configuring the agent	16
Configuring the agent in a global cluster	16
Chapter 4. Managing and testing clustering support	19
Testing service group migration	20
Testing host failure	21
Index	23



Preface

About this book

This document describes how to install and configure the VERITAS Cluster Server (VCS) enterprise agent for IBM DB2 HADR.

If this document is dated more than six months prior to the date you are installing the enterprise agent, contact VERITAS Technical Support to confirm you have the latest supported versions of the application and operating system.

How this guide is organized

[“Introduction”](#) on page 1

[“Installing and removing the agent”](#) on page 7

[“Configuring the agent”](#) on page 11

[“Managing and testing clustering support”](#) on page 19



VERITAS Cluster Server documentation

The following documents, along with the online help and the Release Notes, comprise the VCS documentation:

Title	File Name
<i>VERITAS Cluster Server Installation Guide</i>	vcs_install.pdf
<i>VERITAS Cluster Server User's Guide</i>	vcs_users.pdf
<i>VERITAS Cluster Server Bundled Agents Reference Guide</i>	vcs_bundled_agents.pdf
<i>VERITAS Cluster Server Agent Developer's Guide</i>	vcs_agent_dev.pdf

See the VCS Release Notes for a complete list of documents and VCS agent guides.

Conventions

The following conventions apply throughout the documentation set.

Typeface/Font	Usage
bold	names of screens, windows, tabs, dialog boxes, options, buttons
<i>italic</i>	new terms, book titles, emphasis, variables in tables or body text
Courier	computer output, command references within text
Courier (bold)	command-line user input, keywords in grammar syntax
Courier (bold, italic)	variables in a command
#	UNIX superuser prompt (all shells)



Getting help

For technical assistance, visit <http://support.veritas.com> and select phone or email support. This site also provides access to resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service. Use the Knowledge Base Search feature to access additional product information, including current and past releases of VERITAS documentation.

For license information, software updates and sales contacts, visit <https://my.veritas.com/productcenter/ContactVeritas.jsp>. For information on purchasing product documentation, visit <http://webstore.veritas.com>.

Documentation feedback

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions to clusteringdocs@veritas.com. Include the title and part number of the document (located in the lower left corner of the title page), and chapter and section titles of the text on which you are reporting. Our goal is to ensure customer satisfaction by providing effective, quality documentation. For assistance with topics other than documentation, visit <http://support.veritas.com>.





Introduction

1

The VCS enterprise agent for DB2 HADR provides failover support and recovery in environments employing HADR to replicate data between DB2 databases.

About the IBM DB2 UDB HADR agent

The VCS enterprise agent for DB2 UDB HADR monitors and manages the state of a replicated DB2 database executed on VCS nodes. The agent ensures that the system on which the DB2 HADR resource is online is replicating DB2 database transactions from a primary to a standby database server.

You can use the agent in single VCS replicated data clusters and multi-cluster environments set up using the VCS Global Cluster Option. DB2 UDB does not support partitioned databases (Database Partitioning Feature) and as such the DB2 HADR agent obeys the same limitation.

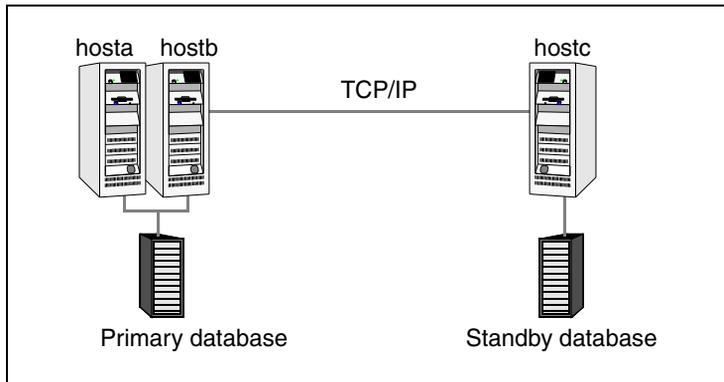
Supported software and hardware

The DB2 HADR agent supports VCS 4.0 and VCS 4.1. The agent supports DB2 UDB version 8.2 and later. This is the first version of DB2 UDB that has the HADR replication option implemented. The agent also supports only a single standby database instance per configured primary database.



Typical setup

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. The primary database can be protected by a local cluster 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, LLT or TCP/IP, between the two data centers to determine their health.

See “[About cluster heartbeats](#)” on page 14.

Agent functions

The agent performs functions called entry points.

online

Depending on the value of the agent attribute HADROper (**LINK** or **SWITCH**), the entry point executes different sequences.

HADROper = LINK

- ◆ If the database snapshot `db2 GET SNAPSHOT FOR ALL ON database_name` returns that the database role is PRIMARY or STANDBY it executes no commands, because the HADR link already exists.
- ◆ If the database configuration `db2 get db cfg for database_name` returns PRIMARY or STANDARD for HADR database role and rollforward pending is NO the agent issues the command `db2 START HADR ON DATABASE database_name AS PRIMARY`.

If the temporary attribute UseForce is 1, this db2 command string is appended with the words `BY FORCE`, and the temp attribute is reset 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.

- ◆ 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 the agent issues the command `db2 START HADR ON DATABASE database_name AS STANDBY`.
- ◆ To mark the completed online sequence, the agent creates a lock file: `$VCS_HOME/bin/Db2HADR/.ResourceID_LCK.dont_remove`.



HADROper = SWITCH

- ◆ If the database snapshot `db2 GET SNAPSHOT FOR ALL ON database_name` returns that the database role is `STANDBY` the agent assumes a failover or group switch condition.
- ◆ If the state of the remote cluster node `hasys -state remote_node` is `FAULTED` or the temporary resource attribute `UseForce` value is 1 the command `db2 TAKEOVER HADR ON DATABASE database_name BY FORCE` is executed.

In any other case, the `db2` command is issued without the `force` option, because it is assumed that the internal DB2 HADR TCP/IP link is still established to do 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`.

offline

Removes the lock file on the local node. A switch of the replication direction—promoting the standby and demoting the primary—is executed on the target node and no DB2 UDB reconfiguration happens as part of offline. In case of a complete shutdown, a standard `Db2udb` resource is responsible to close the database and stop the instance processes.

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 reports online if the database is in role `PRIMARY` using the output from `db2 GET SNAPSHOT FOR ALL ON database_name`. The same is true if the `HADROper` attribute is set to `LINK` and the database role returned by the snapshot is `STANDBY`.

open

Creates a lock file `$VCS_HOME/bin/Db2HADR/.ResourceID_LCK.dont_remove` in the local agent directory if the database role as returned by `db2 get snapshot` is either `PRIMARY` or `STANDBY` with `HADROper` attribute value set to `LINK`. This is necessary to report the correct state when VCS is started after an established HADR link.

clean

Removes the lock file for the following resource states:
OFFLINE TIMEOUT / INEFFECTIVE, ONLINE TIMEOUT, UNEXPECTED OFFLINE, or MONITOR HUNG.

info

Reports the current state and role of the database using from `db2 GET SNAPSHOT FOR ALL ON database_name`.

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.





Installing and removing the agent

2

Install the agent after setting up your cluster. For information about installing and configuring VCS, see the *VERITAS Cluster Server Installation Guide*.

Installing the agent

You must install the DB2 HADR enterprise 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 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 disc device access name is `cd0`.

2. Insert the disc into a drive connected to the system.
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/hadr_agent/
4.1.01/pkgsrc/VRTSscsdbh.rte.bff VRTSscsdbh
```



▼ To install these agent on Linux systems

1. Log in as root.
2. Insert the software disc into a drive connected to the system. If the disc does not automatically mount, type:

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

3. Navigate to the /mnt directory:

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

4. Navigate to the location of the agent binaries for the Linux distribution and architecture in your cluster:

- ◆ For all supported Red Hat distributions:

```
/linux/redhatlinux/replication/db2hadr_agent/4.1.01.0/rpms
```

- ◆ For all supported Suse distributions:

```
/linux/suselinux/replication/db2hadr_agent/4.1.01.0/rpms
```

5. Install the agent software:

```
# rpm -ivh agentrpm
```

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

▼ To install the agent on Solaris systems

1. Insert the disc into a drive connected to the system, type:

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

2. Navigate to the location of the agent binaries:

```
# cd solaris/sparc/replication/hadr_agent/4.1.01.0/pkgsrc/
```

3. Install the agent package:

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

Removing the agent

Identify your platform, and follow the instructions to remove the agent.

▼ To remove the agent from AIX systems

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

```
# installp -u VRTScsdbh
```

▼ To remove the agent from Linux systems

- ❖ Type one of the following commands, depending on the distribution, on each node. Answer the prompts accordingly.

- ◆ For all supported Red Hat distributions, type:

```
# rpm -e VRTScsdbh-4.1.01.0-GA_RHEL4.i686.rpm
```

- ◆ For all supported Suse distributions, type:

```
# rpm -e VRTScsdbh-4.1.01.0-GA_SLES9.i586.rpm
```

▼ To remove the agent from Solaris systems

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

```
# pkgrm VRTScsdbh
```





Configuration concepts

You can adapt most DB2 database environments configured 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—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.

For details on how to configure an DB2 database for HADR replication see the *IBM DB2 Universal Database Data Recovery and High Availability Guide and Reference Version 8.2*.

Resource type definition

The resource type definition defines the 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

Configure the agents behavior by changing its attributes.

Required attribute	Description
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

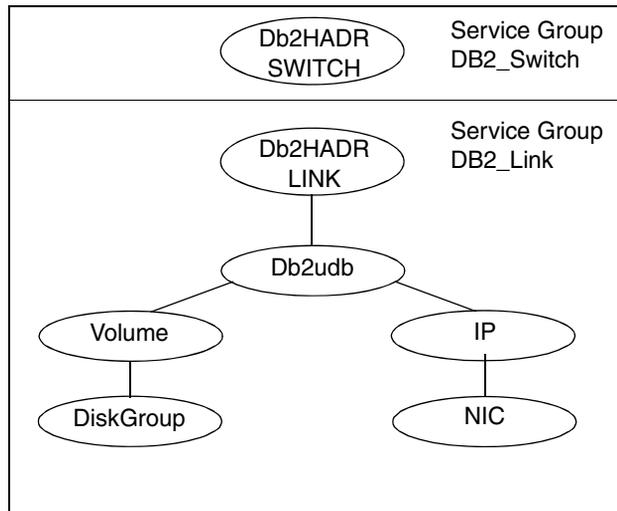
Optional attributes	Description
AgentDebug	For future use. Default: 0 Type-dimension: string-scalar
Encoding	Specifies operating system encoding that corresponds to DB2 UDB encoding for the displayed DB2 output. Default: "" Type-dimension: integer-scalar

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.



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 as follows 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 and only necessary if the primary database is protected by a local cluster on its own. 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 or creating alternate database connect entries on client machines.



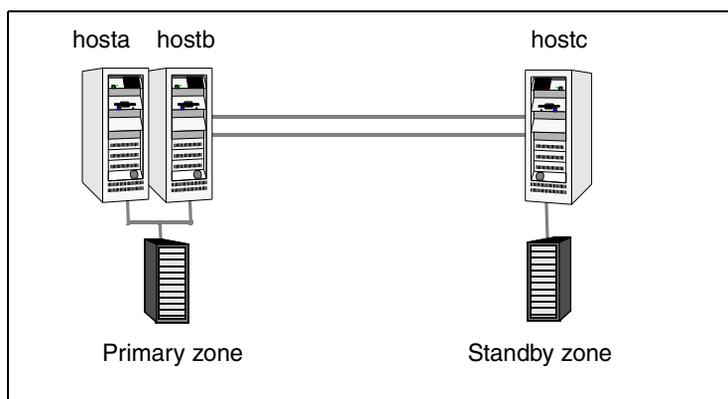
About cluster heartbeats

In global clusters, the VCS Heartbeat agent should send heartbeats directly across the same network 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.

Individual component failure

In a replicated data cluster, you can prevent unnecessary DB2 HADR failover or failback by configuring hosts attached to the same database as part of the same system zone. VCS attempts to fail over applications within the same system zone before failing them over across system zones.

In this sample, *hosta* and *hostb* are in one system zone, and *hostc* is in another system zone. Use the `SystemZones` attribute to create these zones.



Modify the `SystemZones` attribute using:

```
# hagr -modify grpname SystemZones hosta 0 hostb 0 hostc 1
```

The variable *grpname* represents the service group in the cluster.

This command creates two system zones: zone 0 with *hosta* and *hostb*, zone 1 with *hostc*.

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

Protecting the standby database server with a second node is not supported.

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 hosts, 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 since the time of the link failure.

If a failover is initiated due to a disaster at the primary site, and if a synchronization was in progress, the standby database may not contain the most recent data. Nevertheless 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. Database transactions on the original primary committed after the link failure and never replicated to the standby at the time of the takeover can result in inconsistency.

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 VERITAS Volume Manager or File System snapshot is necessary.

Split-brain

Split-brain occurs when all heartbeat links between the primary and secondary hosts are cut and each side mistakenly thinks the other side is down. To minimize the effects of split-brain, it is best if the cluster heartbeat links pass through similar physical infrastructure as the replication links so that 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) 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 HADR at either site. VCS forces you to choose which group to take offline. You must resynchronize the data manually.

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 executes the following actions. It:

- ◆ offlines resource of type Db2HADR with HADROper = SWITCH on the original primary and removes the lock file.
- ◆ onlines resource of type Db2HADR with HADROper = SWITCH on the former standby and alters the database role from standby to primary using the `db2 TAKEOVER` command.

Configuring the agent

Prior to configuring the agent:

- ✓ Verify the agent is installed on all systems in the cluster.
- ✓ Verify the hardware setup for the agent.
See “[Typical setup](#)” on page 2.
- ✓ Make sure the cluster has an effective heartbeat mechanism in place.
See “[About cluster heartbeats](#)” on page 14.
- ✓ Review the agent’s resource type definition and its attribute definitions.
- ✓ Review the configuration concepts, which present information about how VCS behaves during failover and how you can set attributes to customize VCS behavior.
See “[Individual component failure](#)” on page 14.

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 attributes IndepthMonitor and DatabaseName 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.





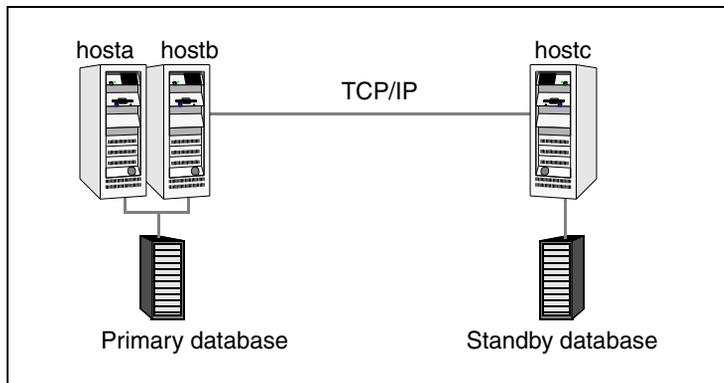
Managing and testing clustering support

4

After configuring the DB2 HADR agent in a VCS environment, you can perform some basic tests to verify the implementation.

These tests assume the following environment:

- ◆ Two hosts (hosta and hostb) attached to a share storage enclosure, and the other host (hostc) attached to another local storage device. The application is running 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. The test scenario is similar for both environments and the heartbeat links can be used for the HADR replication as well.



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 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 attached to the same storage device (hostb) from the menu.
3. Migrate the service group to a host 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 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 or shut down the 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`).





Index

A

- about the DB2 UDB HADR agent 1
- actions/HADRConfig 5
- actions/HADStatus 5
- agent operations 3
 - actions/HADRConfig 5
 - actions/HADStatus 5
 - clean 5
 - info 5
 - monitor 4
 - offline 4
 - online 3
 - open 4
- AgentDebug attribute 12
- AIX, removing the agent 9
- attribute definitions 12

C

- clean 5
- cluster heartbeats 14
- cluster, split-brain 15
- configuring 11
 - concepts 14
 - sample 12

D

- DB2 HADR service group, migrating 20
- DB2 UDB HADR Agent
 - about 1
- Db2dbRes attribute 12
- DBName attribute 12

E

- entry points 3

F

- failure
 - application 15
 - host 15
 - individual components 14
 - replication link 15

H

- HADROper
 - LINK 3
 - SWITCH 4
- HADROper attribute 12
- heartbeats 14
- heartbeats, cluster 14

I

- IBM DB2 HADR agent
 - about 1
 - attribute definitions 12
 - configuration concepts 14
 - installing, AIX 7
 - installing, Linux 8
 - operations 3
- info 5
- installing 7
 - on Solaris systems 8
- installing AIX 7
- installing Linux 8

L

- LINK 3



M

- managing DB2 HADR 19
- migrating service groups 20
- monitor 4

O

- offline 4
- online 3
- open 4
- operations 3
- optional attributes
 - AgentDebug 12
 - Encoding attribute 12

P

- pre-configuring 16

R

- removing
 - from AIX 9
 - from Linux 9
 - from Solaris 9
- required attributes
 - Db2dbRes 12
 - DBName 12
 - HADROper 12
- resource type definition 11
- role transition 16

S

- sample configuration 12
- Solaris systems
 - installing 8
- Solaris, removing the agent 9
- split-brain, handling in cluster 15
- supported
 - hardware 1
 - software 1
- SWITCH 4

T

- testing DB2 HADR 19

