

Cluster Server Agent for IBM DB2 HADR Installation and Configuration Guide

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.

Typical IBM DB2 HADR setup in a VCS cluster

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.

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	Executes different sequences, depending on the value of the agent attribute HADROper (LINK or SWITCH).
offline	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.
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 <code>db2 GET SNAPSHOT FOR ALL ON database_name</code> 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 <code>db2 get snapshot</code> 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	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 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 VCS agent for IBM DB2 HADR, ensure that you install and configure VCS on all nodes in the cluster.

Set up replication and the required hardware infrastructure.

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 Veritas Services and Operations Readiness Tools (SORT) 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      cdl/aix/vcs/replication/DB2_HADR_agent/  
         agent_version/pkgsg/
```

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

```
Solaris  cdl/solaris/dist_arch/vcs/replication/DB2_HADR_agent/  
         agent_version/pkgsg/
```

If you downloaded the individual agent tar file, navigate to the pkgsg directory (for AIX, HP-UX, 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.
- 3 Delete the `/etc/VRTSvcs/conf/config/Db2HADRTypes.cf` file.
- 4 Install the current version of the agent.
- 5 Copy the file `Db2HADRTypes.cf` from the `/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 the `/etc/VRTSvcs/conf/config/main.cf` configuration file.

Configure the new attributes, if applicable.

- 8 Verify the configuration by running the following command:

```
# hacf -verify config
```

- 9 Start VCS on the local node first.

```
# hastart - local
```

- 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 DB2 HADR agent from each node in the cluster.

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

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

```
Linux        # rpm -e VRTSvcsdbh
```

Solaris

```
# pkgrm VRTScsdbh
```

Note: To uninstall the agent IPS package on a Solaris 11 system:

```
# pkg uninstall 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	Name of the Db2udb 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
PeerWindowOnly	Specifies whether the agent must use the PEER WINDOW ONLY option when the agent is performing a takeover operation by force. Set the value of this attribute to 1 if you want the agent to use the PEER WINDOW ONLY option. Default: 0 Type-dimension: boolean-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 for the DB2 HADR agent

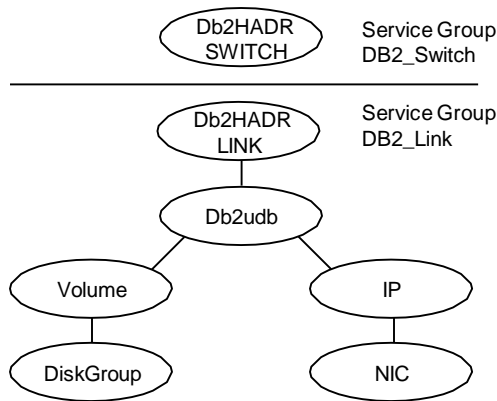
This is an example of a VCS service group that has a resource of type Db2HADR with attribute HADROper = SWITCH. The second child service group contains all

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 the cluster nodes.
- Verify the hardware setup for the agent.
- Make sure that the cluster has an effective heartbeat mechanism in place.
- Set up an effective heartbeat mechanism to prevent split-brain.

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.

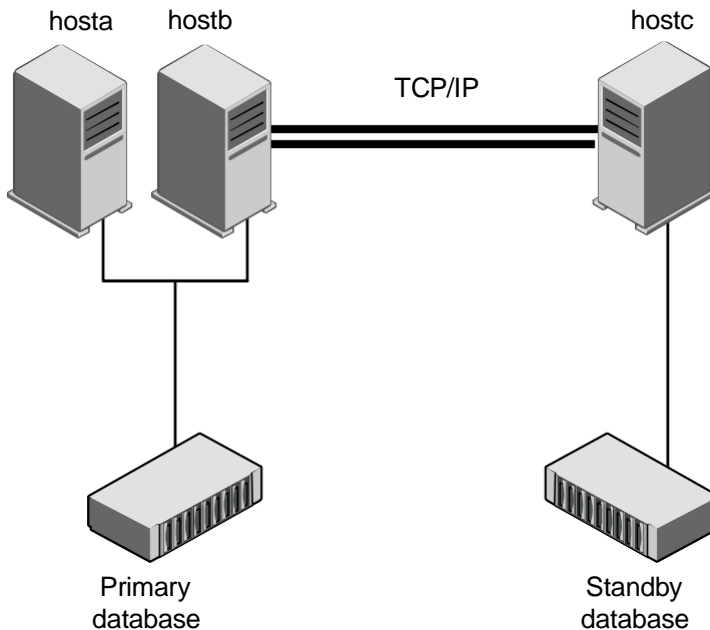
About configuring system zones in replicated data clusters

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 HADR 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 adapt most clustered applications 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 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/VRTSvcs/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 Db2HADR handles the database open.
- 6 Create a new resource of type Db2HADR 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 Db2HADR resource dependent on the Db2udb resource.
- 9 Add a new group with at least one resource of type Db2HADR.
- 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 Db2HADR group and the existing DB2 database group.
- 12 Change the ClusterFailOverPolicy from the default, if necessary. Veritas recommends keeping the Manual default, to minimize the chance of failing over on a split-brain.