# Veritas™ Cluster Server Agent for DB2 Installation and Configuration Guide

AIX, Linux, Solaris

5.0

12067555

symantec.

# Veritas Cluster Server Agent for DB2 Installation and Configuration Guide



Veritas Cluster Server 5.0



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.

## Technical support

For technical assistance, visit http://www.symantec.com/enterprise/support/assistance_care.jsp 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 3 Installing, upgrading, and removing the agent for DB2 UDB



## Chapter 4 Configuring VCS service groups for DB2 UDB



## Chapter 5 Administering VCS service groups for DB2 UDB

Chapter 1

# Introducing the Veritas High Availability Agent for DB2

This chapter contains the following topics:



## About the Veritas High Availability Agent for DB2

The Veritas High Availability agent, version 5.0, for DB2 UDB. DB2 Universal Database is a high availability solution for the relational database management system.

This guide describes the agent for DB2 UDB, its modes of operation, and its attributes. It describes how to install and configure the agent. It also describes how to administer service groups where the resource runs.

# How the agent makes DB2 highly available

The agent for DB2 monitors DB2 database processes at the partition level. If the system fails, the agent detects the failure and takes the DB2 instances offline. VCS fails over the node to another system in the cluster, where the agent brings the DB2 UDB database partition or partitions online.

The agent can perform different levels of monitoring for DB2, which you can configure. You can also configure many of the different actions it performs. For more information on configuring the different functions that the agent has:

See "DB2 agent functions" on page 11.

# Supported software

The Veritas High Availability agent, version 5.0 for DB2 UDB, supports the DB2 Universal Database Enterprise Server Edition (ESE).

For the ESE multi-partition instance, it supports both of the following:

- Symmetric Multiprocessing (SMP) hardware configuration
- Massively Parallel Processing (MPP) hardware configuration.

The agent for DB2 ESE must supports the following platforms or single and multi-partition instances:

**Table 1-1** Supported operating systems and versions

| Operating system | DB2 UDB version |
|---|---|
| AIX 5.2 or later | 8.1, 8.2, and 9.1 |
| RHEL 4.0 | 8.1, 8.2, and 9.1 |
| SLES 9.0 | 8.1, 8.2, and 9.1 |
| Solaris SPARC 2.8 or later | 8.1, 8.2, and 9.1 |
| Solaris x64 2.10 or later | 8.2 |

The Veritas Cluster Server (VCS) version must be 5.0 or later. The memory requirements vary for different versions of DB2 being used. Check the relevant IBM DB2 guide for information about memory requirements.

# DB2 agent functions

The agent can perform different operations or functions on the database. These functions are online, offline, monitor, clean info, and action. With the action agent function, you can add custom actions for the agent. For information on defining custom actions, refer to the *Veritas Cluster Server Agent Developer's Guide.*

Online The agent uses db2gcf program to start a DB2 instance or database partition. The command is:

```
su $DB2InstOwner -c "$InstHome/sqllib/bin/db2gcf -u -i
    $DB2InstOwner -p $nodenum
```

Offline The agent uses the db2gcf program to stop a DB2 database partition. For a database with only one partition, it stops the instance. The command is:

```
su $DB2InstOwner -c "$InstHome/sqllib/bin/db2gcf -d -i
    $DB2InstOwner -p $nodenum
```

Monitor

The agent executes the `db2gcf -s -i $DB2InstOwner -p $nodenum` command to check the status of the database partition or node number. If the exit status of the `db2gcf` command is 0, the monitor returns the exit code 110. Otherwise, the monitor returns an exit code of 100 and the resource is taken offline. The agent then restarts or fails over the resource, depending on other type-independent attributes, such as RestartLimit or ToleranceLimit.

Set the IndepthMonitor attribute to 1 for in-depth monitoring. The agent looks for the custom_monitor_$db2instance_$nodenum file in the /opt/VRTSagents/ha/bin/Db2udb directory. It executes this user-defined customized in-depth monitor file if the file exists and is executable. You can find samples of custom monitor scripts in the /etc/VRTSAgents/ha/conf/sample_db2udb directory.

**Note:** When you run the DB2 instance inside of a Solaris 10 non-global zone, create the custom_monitor_$db2instance_$nodenum file inside the non-global zone. Create the file after doing a "`zlogin zonename`" to login to the zone. You then change to the /etc/VRTSagents/ha/bin/Db2udb directory, which is relative to the local zone.

If the custom monitor has any errors or problems, it checks the value of the WarnOnlyIfDBQueryFailed attribute of the Db2udb agent. If you have a db2error.dat file in the /opt/VRTSagents/ha/bin/Db2udb directory, the agent checks this file, and handles the error according to the error configuration. For error handling information:

See "Handling DB2 error codes during in-depth monitoring" on page 51.

If you set the WarnOnlyIfDBQueryFailed attribute to 1 (its default), and you have configured the Notifier resource, the agent:

- Sends a notification
- Returns the exit code 110

If you set the WarnOnlyIfDBQueryFailed attribute to 0, it performs error handling in the db2error.dat file. Note that the file needs to exist to perform error handling. If it does not exist, it returns the exit code 100, which is the default.

Clean

The agent uses the `db2gcf` program to kill a DB2 database partition. The command is:

```
su $DB2InstOwner -c "$InstHome/sqllib/bin/db2gcf -k -i
    $DB2InstOwner -p $nodenum"
```

Info

The agent for DB2 supports the Info operation, which provides static and dynamic information about the database partition and its critical processes. For more information about the info agent function, see the:

- *Veritas Cluster Server User's Guide*
- *Veritas Cluster Server Agent Developer's Guide*

For an example of the info agent function retrieving database information:

See "Running the Info agent function" on page 13.

Action The agent for DB2 supports the Action operation, which enables you to perform predefined actions or custom actions on a resource. To perform an action on a resource, type the following command:

```
# hares -action res token [-actionargs arg1 ...] \
[-sys system]   [-rclus cluster]
```

The agent supports these predefined actions:

- The VRTS_GetInstanceName token retrieves the DB2 instance name of the configured Db2udb resource.
- The VRTS_GetRunningServices token retrieves the list of processes that the agent monitors for the Db2udb resource.

For example:

```
# hares -action db2udb1 VRTS_GetInstanceName \
-sys systemName
    VCS NOTICE V-16-13323 Resource (db2udb0): action
    (VRTS_GetInstanceName) completed successfully. Output
    is:

    db2inst1
# hares -action db2udb0 VRTS_GetRunningServices \
-sys systemName

    VCS NOTICE V-16-13323 Resource (db2udb0): action
    (VRTS_GetRunningServices)
    completed successfully. Output is:
    PARTITION: 0

        PID   TTY   TIME  CMD
        9800  ?     0:06  db2sysc
```

## Running the Info agent function

You can run the Info agent function to return database information. In this example, the Info agent function retrieves the database information.

### To retrieve database information

1 Make the configuration writable:

```
# haconf -makerw
```

2 Specify the periodic interval in seconds that the info agent function is invoked. The default value of 0 means info agent function is not invoked.

```
# hatype -modify Db2udb InfoInterval 300
```

In the above command, Db2udb is the name of the DB2 resource type. InfoInterval 300 is the duration (in seconds) after which the Info agent function executes the Info script. The script gets the processes' information that the agent monitors.

3 Show the requested ResourceInfo value. The following example output shows the processes that the agent monitors for the DB2 resource.
Note that ResourceInfo refreshes every 300 seconds (five minutes), since you set the InfoInterval to 300 in the previous step.

```
# hares -value db2udb1 ResourceInfo
State Valid
Msg
PARTITION: 0
    PID    TTY  TIME CMD
  413924    -   0:00 db2sysc

TS Fri Jan 14 18:11:52 2005
```

# Typical DB2 configuration in a VCS cluster

In the following examples, VCS is configured on a two-system cluster. DB2 UDB system binaries are installed identically on local file systems on System A and System B. The instance home directory, instance binaries, and the database reside on shared storage, available to either node. In the case of the non-MPP configuration, an instance is online on only one system at a time. The other system is the failover system.

**Figure 1-1** DB2 installation with a non-MPP configuration

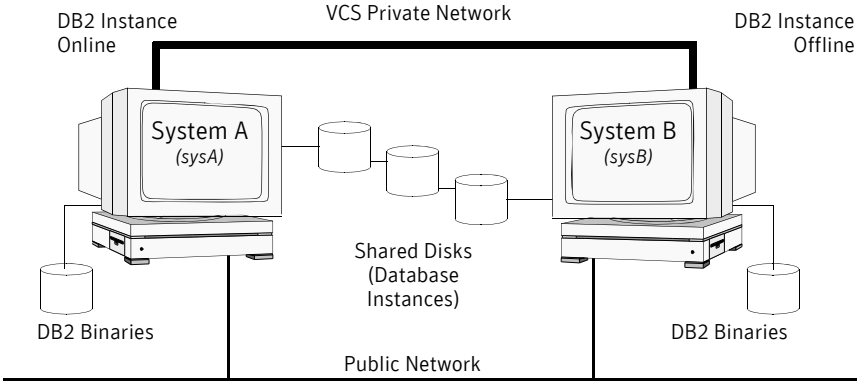


In the case of the MPP configuration, a database partition can run on each system and each system can become a failover system.

**Figure 1-2** DB2 installation with an MPP configuration

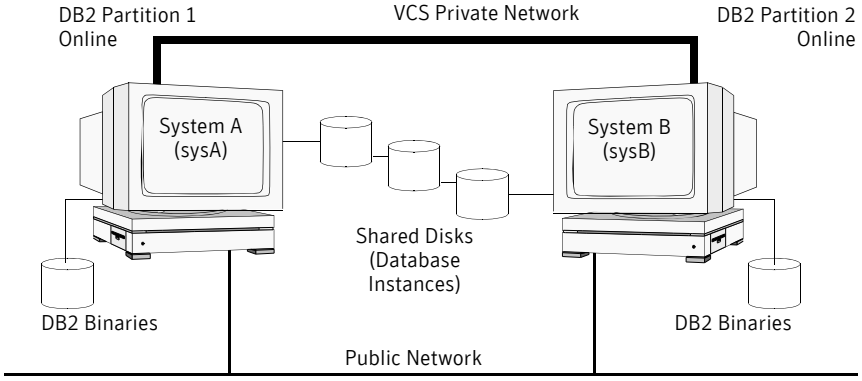


# Road map for setting up a DB2 UDB cluster

Use the following road map when you set up the cluster.

**Table 1-2** Setting up a DB2 UDB cluster

| Task | Reference |
|---|---|
| Review the following:<br>■ The supported software<br>■ The agent functions<br>■ The agent requirements | ■ See "Supported software" on page 10.<br>■ See "DB2 agent functions" on page 11.<br>■ See "VCS requirements for installing DB2 UDB" on page 17. |
| Install and set up DB2. | See "Installing DB2 UDB in a VCS environment" on page 25.<br>See "Setting up the DB2 UDB configuration" on page 25. |
| Install the VCS agent for DB2. | See "Installing the agent for DB2 UDB" on page 32. |
| Configure the service groups for the VCS agent for DB2.<br>Optionally, configure in-depth monitoring and any automated actions. | See "Configuring the DB2 UDB agent using Cluster Manager (Java Console)" on page 43.<br>See "Setting up in-depth monitoring of DB2 UDB instance" on page 50. |
| Bring the service group online. | See "Bringing the service group online" on page 53. |

Chapter 2

# Installing and configuring DB2 UDB

This chapter contains the following topics:



## VCS requirements for installing DB2 UDB

Review the following prerequisites and requirements before you install DB2.

### Prerequisites for installing DB2 UDB, non-MPP versions

Perform the following prerequisites before installing the non-MPP versions of DB2.

- Verify that all systems have enough resources, such as shared memory, to run DB2 UDB. Check the DB2 memory requirements, which vary depending on the version and hardware configuration of DB2.
- Install the DB2 UDB system binaries locally.
- Install the DB2 UDB database instances on shared storage.
  - Install and configure VCS version 5.0 on all nodes in the cluster. For installation instructions, see to the *Veritas Cluster Server Installation Guide*.

- Before installing DB2 UDB, define DB2 UDB user and group accounts, see:
  - "Defining DB2 user and group accounts" on page 22
  - For your particular version of DB2, refer to the appropriate DB2 UDB guide.

## Prerequisites for installing DB2 UDB, MPP version

Perform the following prerequisites before installing the MPP versions of DB2.

- Verify that all systems have enough resources, such as shared memory, to run DB2 UDB. Check the DB2 memory requirements, which vary depending on the version and hardware configuration of DB2.
- Install the DB2 UDB system binaries on the local file systems on each system.
- Install the DB2 UDB database instances on shared storage.
- The MPP configuration requires the Storage Foundation Cluster File System software. This software includes the:
  - Veritas Cluster Server (VCS)
  - Veritas Volume Manager with cluster functionality enabled (CVM)
  - Veritas File System with cluster functionality enabled (CFS)

  For information and installation instructions for these products, see the:
  - *Veritas Storage Foundation Cluster File System Installation Guide*
  - *Veritas Storage Foundation Cluster File System Administration Guide*
- Before installing DB2 UDB, define DB2 UDB user and group accounts. See the:
  - "Defining DB2 user and group accounts" on page 22
  - The relevant DB2 UDB guide

## Creating file systems for DB2 instances

The following sections describe examples of creating disk groups for the DB2 database instances.

### Creating the file system for the DB2 non-MPP instances

To create a file system, you first create a disk group on the physically shared disk. You then create a volume of sufficient size within the disk group.

### To create a file system on AIX systems for non-MPP instances

1. Create a disk group on the shared disk. List the disks using the `lsdev -Cc disk` command. In this case the group consists of one disk, hdisk5. For example:

   ```
   # vxdg init db2db_dg hdisk5
   ```

   Deport and import the disk group:

   ```
   # vxdg deport db2db_dg
   # vxdg import db2db_dg
   ```

2. Create a volume of three GB using the vxassist command:

   ```
   # vxassist -g db2db_dg make db2db_vol 3g
   ```

3. Create the file system:

   ```
   # mkfs -V vxfs -o largefiles /dev/vx/dsk/db2db_dg/db2db_vol
   ```

4. Create the mount point directory and mount the file system. Make sure that the mount point exists on all nodes in the cluster on the local file system—not on shared storage.

   ```
   # mkdir /db2_mnt/db2inst1
   # mount -V vxfs /dev/vx/dsk/db2db_dg/db2db_vol \
     /db2_mnt/db2inst1
   ```

### To create a file system on Linux systems for non-MPP instances

1. Create a disk group on the shared disk. List the disks using the `vxdisk list` command. In this case the group consists of one disk, sdc. For example:

   ```
   # vxdg init db2db_dg /dev/sdc
   ```

   Deport and import the disk group:

   ```
   # vxdg deport db2db_dg
   # vxdg import db2db_dg
   ```

2. Create a volume of three GB using the vxassist command:

   ```
   # vxassist -g db2db_dg make db2db_vol 3g
   ```

3. Create the file system:

   ```
   # mkfs -t vxfs /dev/vx/dsk/db2db_dg/db2db_vol
   ```

4. Create the mount point directory and mount the file system. Make sure that the mount point exists on all nodes in the cluster on the local file system—not on shared storage.

   ```
   # mkdir /db2_mnt/db2inst1
   # mount -t vxfs /dev/vx/dsk/db2db_dg/db2db_vol \
     /db2_mnt/db2inst1
   ```

#### To create a file system on Solaris systems for non-MPP instances

1. Create a disk group on the shared disk. List the disks using the `vxdisk list` command. In this case the group consists of one disk, c4t0d0s2. For example:

   ```
   # vxdg init db2db_dg c4t0d0s2
   ```

   Deport and import the disk group:

   ```
   # vxdg deport db2db_dg
   # vxdg import db2db_dg
   ```

2. Create a volume of three GB using the vxassist command:

   ```
   # vxassist -g db2db_dg make db2db_vol 3g
   ```

3. Create the file system:

   ```
   # mkfs -F vxfs /dev/vx/rdsk/db2db_dg/db2db_vol
   ```

4. Create the mount point directory and mount the file system. Make sure that the mount point exists on all nodes in the cluster on the local file system—not on shared storage.

   ```
   # mkdir /db2_mnt/db2inst1
   # mount -F vxfs /dev/vx/dsk/db2db_dg/db2db_vol \
     /db2_mnt/db2inst1
   ```

## Creating the shared cluster file system for the DB2 MPP instances

To create a shared file system, you first create a shared disk group on the physically shared disk. You then create a volume of sufficient size within the disk group. You must have installed the Storage Foundation Cluster File System software.

#### To create a shared file system on AIX for MPP instances

1. You must issue the commands to create a shared disk group from the CVM master node. To determine whether a node is the master or the slave, enter the command:

   ```
   # vxdctl -c mode
   ```

   In the output, look for:

   ```
   cluster active - MASTER
   ```

   Or

   ```
   cluster active - SLAVE
   ```

2. From the master node, create the disk group. List the disks using the `vxdisk list` command.

3. Create a shared disk group. In this case, the group consists of one disk. In this example the disk is hdisk5:

   ```
   # vxdg -s init db2db_dg hdisk5
   ```

4 Deport and import the disk group:

```
# vxdg deport db2db_dg
# vxdg -s import db2db_dg
```

5 Use the `vxassist` command to create a 7-GB volume:

```
# vxassist -g db2db_dg make db2db_vol 7g
```

6 Create the file system:

```
# mkfs -V vxfs -o largefiles /dev/vx/rdsk/db2db_dg/db2db_vol
```

7 Create the mount point directory and mount the file system.

```
# mkdir /db2_mnt/db2inst1
# mount -V vxfs -o cluster /dev/vx/dsk/db2db_dg/db2db_vol \
  /db2_mnt/db2inst1
```

### To create a shared file system on Linux for MPP instances

1 You must issue the commands to create a shared disk group from the CVM master node. To determine whether a node is the master or the slave, enter the command:

```
# vxdctl -c mode
```

In the output, look for:

```
cluster active - MASTER
```

Or

```
cluster active - SLAVE
```

2 From the master node, create the disk group. List the disks using the `vxdisk list` command.

3 Create a shared disk group. In this case, the group consists of one disk. In this example the disk is sdc:

```
# vxdg -s init db2db_dg hdisk5
```

4 Deport and import the disk group:

```
# vxdg deport db2db_dg
# vxdg -s import db2db_dg
```

5 Use the `vxassist` command to create a 7-GB volume:

```
# vxassist -g db2db_dg make db2db_vol 7g
```

6 Create the file system:

```
# mkfs -t vxfs -o largefiles /dev/vx/rdsk/db2db_dg/db2db_vol
```

7 Create the mount point directory and mount the file system.

```
# mkdir /db2_mnt/db2inst1
# mount -t vxfs -o cluster /dev/vx/dsk/db2db_dg/dbq2db_vol \
  /db2_mnt/db2inst1
```

**To create a shared file system on Solaris for MPP instances**

1. You must issue the commands to create a shared disk group from the CVM master node. To determine whether a node is the master or the slave, enter the command:

   ```
   # vxdctl -c mode
   ```

   In the output, look for:

   ```
   cluster active - MASTER
   ```

   Or

   ```
   cluster active - SLAVE
   ```

2. From the master node, create the disk group. List the disks using the `vxdisk list` command.

3. Create a shared disk group. In this case, the group consists of one disk. In this example the disk is c5t0d0s2:

   ```
   # vxdg -s init db2db_dg c5t0d0s2
   ```

4. Deport and import the disk group:

   ```
   # vxdg deport db2db_dg
   # vxdg -s import db2db_dg
   ```

5. Use the `vxassist` command to create a 7-GB volume:

   ```
   # vxassist -g db2db_dg make db2db_vol 7g
   ```

6. Create the file system:

   ```
   # mkfs -F vxfs -o largefiles /dev/vx/rdsk/db2db_dg/db2db_vol
   ```

7. Create the mount point directory and mount the file system.

   ```
   mkdir /db2_mnt/db2inst1
   mount -F vxfs -o cluster /dev/vx/dsk/db2db_dg/db2db_vol \
    /db2_mnt/db2inst1
   ```

## Defining DB2 user and group accounts

Before installing DB2 UDB binaries and creating instances, you must define DB2 UDB user and group accounts for each instance on each system. Note the following requirements:

- The IDs for DB2 users and groups must be exactly the same across all cluster systems.
- The DB2 instance owner's home directory must exist locally on each node. This directory is the mount point that the DB2 instance uses. The database that you want to mount must be on shared storage. Create the mount point directory locally on each node if it does not already exist.

- All DB2 user accounts must exist on the local systems. Symantec does not recommend the use of NIS or NIS+ for users because these services are not highly available. If their service is interrupted, VCS may not be able to work correctly.

## Creating user group accounts

Three user group accounts are required on each node in the cluster.

### To create the group accounts on *each* node in the cluster for AIX systems

1. Create a group for the DB2 UDB instance owner. For example, enter:
   ```
   # mkgroup id=999 db2iadm1
   ```
2. Create a group for the user to execute fenced user-defined functions (UDFs) or store procedures. For example, enter:
   ```
   # mkgroup id=998 db2fadm1
   ```
3. Create a group for the database administration server. For example, enter:
   ```
   # mkgroup id=997 db2asgrp
   ```

### To create the group accounts on *each* node in the cluster for Linux and Solaris systems

1. Create a group for the DB2 UDB instance owner. For example, enter:
   ```
   # groupadd -g 999 db2iadm1
   ```
2. Create a group for the user to execute fenced user-defined functions (UDFs) or store procedures. For example, enter:
   ```
   # groupadd -g 998 db2fadm1
   ```
3. Create a group for the database administration server. For example, enter:
   ```
   # groupadd -g 997 db2asgrp
   ```

## Adding user accounts for AIX systems

Create the user accounts on each node in the cluster.

### To create the user accounts on each node in the cluster for AIX systems

- This example shows creating the user, db2inst1, who is the DB2 UDB instance owner. The instance's home directory is also the mount point, /db2_mnt/db2inst1. The file system that hosts the DB2 UDB instance home directory on shared storage uses this mount point. The DB2 UDB instance home directory must exist on every node. For example:
  ```
  # mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/ \
  db2_mnt/db2inst1 db2inst1
  ```

- These examples show creating user accounts for db2fenc1 and db2as. These users' home directories are under /home in the local file system on each node.

```
# mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/home/ \
db2fenc1 db2fenc1
# mkuser id=1002 pgrp=db2asgrp groups=db2asgrp home=/home/ \
db2as db2as
```

## Adding user accounts for Linux and Solaris systems

In the following examples that show creating user accounts, you can use the following options:

- The -g option specifies the group
- The -u option specifies the user ID
- The -d option specifies the home directory
- The -m option creates the home directory if it doesn't exist
- The -s option is the user's login shell
- The final expression is the user's login.

Create the user accounts on each node in the cluster.

### To create the user accounts on each node in the cluster for Linux and Solaris systems

- This example shows creating the user, db2inst1, who is the DB2 UDB instance owner. The instance's home directory is also the mount point: /db2_mnt/db2inst1. The file system that hosts the DB2 UDB instance home directory on shared storage uses this mount point. The DB2 UDB instance home directory must exist on every node. For example:

```
# useradd -g db2iadm1 -u 1004 -d /db2_mnt/db2inst1 -m -s
/bin/ksh/ db2inst1
```

- These examples show creating user accounts for db2fenc1 and db2as. These users' home directories are under /home in the local file system on each node.

```
# useradd -g db2fadm1 -u 1003 -d /home/db2fenc1 -m -s /bin/
ksh db2fenc1
# useradd -g db2asgrp -u 1002 -d /home/db2as -m -s /bin/ksh
db2as
```

# Installing DB2 UDB in a VCS environment

For installing DB2 UDB in a VCS environment, Symantec recommends that you follow the installation procedure in the relevant IBM DB2 UDB guide.

Install binaries on local disks of each node, and the database instances on shared storage, accessible by each cluster node.

## Setting shared memory parameters

Refer to the relevant IBM DB2 UDB guide to make sure that memory requirements are met. On Solaris systems, set the memory parameters in the /etc/system file.

## Installing the binaries

Install the DB2 UDB system binaries on local disks on each node (mirrored disks are recommended) not on shared storage. You can use IBM's db2setup tool.

## Install the DB2 license

Install the DB2 license on each node. For example, enter:

```
# /opt/IBM/db2/V8.1/adm/db2licm -a db2ese.lic
```

## Installing the instances

Install the database instances on the shared storage only on the one node where the instance's home directory is currently mounted. You can choose to install single-partition instance or multi-partition instance. You can use IBM's db2setup tool.

- When you use the db2setup, do not select the option to Auto start DB2 instance at system boot in the DB2 Instance Properties window. Note that this option does not exist on all DB2 versions. VCS needs to bring up the resources for the DB2 instances in a specific order before it brings the instance online.
- The instance's home directory is a mount point on the local system.

# Setting up the DB2 UDB configuration

Use the following procedures to configure DB2 UDB in a VCS environment.

## Checking /etc/services

On each system in the cluster, use the `more` command to check the file /etc/services.

- Make sure each partition has a port number assigned. The number of reserved ports depends on the number of partitions.
- Make sure that no other services use the ports. Manually assign new numbers if necessary.
- Make sure all systems in the cluster have the same entries in the /etc/services file.

The following is an example for two DB2 UDB instances: db2inst1 and db2inst2. Both instances have two partitions each. Each instance requires two ports plus one port per partition, hence four lines per instance.

```
# more /etc/services
DB2_db2inst1      60000/tcp
DB2_db2inst1_1    60001/tcp
DB2_db2inst1_2    60002/tcp
DB2_db2inst1_END  60003/tcp
DB2_db2inst2      60004/tcp
DB2_db2inst2_1    60005/tcp
DB2_db2inst2_2    60006/tcp
DB2_db2inst2_END  60007/tcp
```

Inspect the file and verify that no duplicate port numbers exist.

## Creating $DB2InstHome/.rhosts

On each system, create a file named $DB2InstHome/.rhosts, and place a "+" character within it. This file permits a system to access the database without the use of a password.

If security is a concern, put the hostname and userid inside the .rhosts file, as shown in the following examples:

```
dbmach01   db2inst1
dbmach02   db2inst1
dbmach03   db2inst1
dbmach04   db2inst1
```

Or

```
+   db2inst1
```

With the `rsh` *`system_name`* command, you can test passwordless remote login. From one system in the cluster to another, the command tests that you can remotely log in with the DB2 instance (for example, db2inst1) account. You should not be prompted for a password. Test this command from each system in the cluster to all other systems.

## Configuring ssh on Suse

Perform the following procedure to configure ssh on Suse.

### To configure ssh on Suse

1. Log on to the system from which you want to install VCS.
2. Generate a DSA key pair on this system by running the following command:
   ```
   # ssh-keygen -t dsa
   ```
3. Accept the default location: ~/.ssh/id_dsa
4. At the prompt, enter a passphrase and confirm it.
5. Change the permissions of the .ssh directory, type:
   ```
   # chmod 755 ~/.ssh
   ```
6. The file ~/.ssh/id_dsa.pub contains a line beginning with ssh_dss and ending with the name of the system on which it was created. Copy this line to the /root/.ssh/authorized_keys2 file on all systems where VCS is to be installed.

   **Note:** If the local system is part of the cluster, make sure to edit the authorized_keys2 file on that system.

7. Run the following commands on the system from which the installation takes place:
   ```
   # exec /usr/bin/ssh-agent $SHELL
   # ssh-add
   ```

   **Note:** This step is shell-specific and is valid for the duration the shell is alive.

8. When the installer prompts you, enter your DSA passphrase.
   You are ready to install VCS. You can install it on several systems by running the installvcs script on any one of them or on an independent machine outside the cluster.
   To avoid running the ssh-agent on each shell, run the X-Window system and configure it so that it does not prompt you for the passphrase. Refer to the Red Hat documentation for more information.

# Modifying the $DB2InstHome/sqllib/db2nodes.cfg file

DB2 uses the $DB2InstHome/sqllib/db2nodes.cfg file during failover from one node to another.

## Non-MPP versions

For each DB2 UDB instance (non-MPP) database partition, modify the file $DB2InstHome/sqllib/db2nodes.cfg. You need to create an entry for each database partition, and to assign the virtual IP address as the hostname. For example:

```
0 virtualhostname 0
1 virtualhostname 1
```

Note that the *virtualhostname* corresponds to the virtual IP address in the /etc/hosts file. Make sure that the virtual IP address is up and running at this time.

## MPP versions

For MPP versions, modify the file $DB2InstHome/sqllib/db2nodes.cfg with the hostname that you want each database partition to start on. DB2 automatically changes and updates the db2nodes.cfg file to enable the database partitions to fail over from one node to another. DB2 adds a fourth column for the “netname,” which is, by default, the hostname. The virtual IP is not used in the db2nodes.cfg file for MPP configurations.

For example:

```
0 sysA 0
1 sysB 0
2 sysC 0
3 sysD 0
```

Make sure that the relative port number in the third column is unique for each partition on a host. For example:

```
0 sysA 0
1 sysA 1
2 sysB 0
3 sysC 0
4 sysD 0
```

## Confirming the setup of DB2 MPP and non-MPP installations

On the host where the shared file system is mounted, check whether you can start and stop each instance. Do this procedure to verify the DB2 installation.

### To check if a DB2 instance can start and stop

1. Log in as the instance owner:
   ```
   # su - db2inst1
   ```
2. Attempt to start the instance:
   ```
   $ db2start
   ```
   DB2 should start on the partitions in the db2nodes.cfg file. If DB2 does not start, check the error codes.
3. Assuming that the previous command is successful, stop the instance:
   ```
   $ db2stop
   ```
4. If the application does not stop correctly on each node, check for configuration errors. Review the DB2 UDB documentation for error codes.
5. Create a database.
   ```
   $ db2 create database dbname
   ```
6. List the database directory
   ```
   $ db2 list database directory
   ```

### To check the rest of the DB2 configuration in the cluster

1. For each node in the VCS cluster, import the disk group and start all the volumes in the disk group.
2. Mount the file system for the volume containing the DB2 instance and database.
3. Unmount and deport the disk group.
4. Repeat this procedure for each node in the cluster.

Chapter 3

# Installing, upgrading, and removing the agent for DB2 UDB

This chapter contains the following topics:



## Before installing or upgrading the agent

Meet the prerequisites to install or upgrade the Veritas High Availability Agent for DB2.

- See "VCS requirements for installing DB2 UDB" on page 17.
- See "Installing DB2 UDB in a VCS environment" on page 25.
- Make sure that VCS is installed in the cluster.
- Verify that DB2 is installed and configured.
- Symantec recommends installing the VCS graphical user interface, the Cluster Manager (Java Console). For more information, see to the *Veritas Cluster Server Installation Guide*.

# Installing the agent for DB2 UDB

You now need to mount the disc and install the agent.

### To install the agent on AIX systems

1. Log in as superuser.
2. Determine the device access name of the disc drive. For example, enter:

   ```
   # 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, cd0 is the drive's device access name.
3. Insert the software disc containing the DB2 UDB agent software into the system's disc drive.
4. Mount the software disc using the device access name that you found in step 2:

   ```
   # mkdir -p /cdrom
   # mount -V cdrfs -o ro /dev/cd0 /cdrom
   # cd /cdrom
   ```
5. Install the DB2 UDB agent software by entering:

   ```
   # installp -ac -d /cdrom/VRTSvcsdb.rte.bff VRTSvcsdb.rte
   ```
6. Repeat this procedure on each system that you want to be a part of the DB2 service group.

### To install the agent on Linux systems

1. Log in as superuser.
2. Insert the software disc the system's drive. The disc automatically mounts. If the disc does not automatically mount, enter:

   ```
   # mount -o ro /dev/cdrom /mnt/cdrom
   ```
3. Go to the directory where rpm is present:
   - On RHEL 4:

     ```
     # cd /mnt/cdrom/rhel4_ppc64/cluster_server_agents/\
     db2_agent/rpms
     ```
   - On SLES 9:

     ```
     # cd /mnt/cdrom/sles9_ppc64/cluster_server_agents/\
     db2_agent/rpms
     ```
4. Install the DB2 UDB agent software:

   ```
   # rpm -i VRTSvcsdb-5.0.12.00-MP1_GENERIC.noarch.rpm
   ```

5 Verify that the package is installed, enter:

```
# rpm -q VRTSvcsdb
```

Where the output resembles:

```
VRTSvcsdb-5.0.00.12-MP1_GENERIC
```

### To install the agent on Solaris systems

1 Log in as superuser.

2 Create a temporary directory for installation:

```
# mkdir /tmp/install
```

3 Insert the disc into a system drive.

- If you are running Solaris volume-management software, the software automatically mounts the disc as /cdrom/cdrom0. Type the following command to go to the location: # **cd /cdrom/cdrom0**
- If you are not running Solaris volume-management software, you must mount the disc manually. For example:

  ```
  # mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
  ```

  Where, in this example, /dev/dsk/c0t6d0s2 is the default for the CD drive.

  Once the disc is mounted, type the following commands to go to the location: # **cd /cdrom**

4 Copy the compressed package files from the software disc to the temporary directory:

```
# cp -r db2_agent/pkgs/* /tmp/install
```

5 Go to the temporary directory and unzip the compressed package file:

---

**Note:** If your system does not have the gunzip utility, copy it from the disc:
**# cp /cdrom_path/gnu/gunzip /tmp/install**

---

```
# cd /tmp/install
# gunzip VRTS*.gz
```

6 Extract the compressed file from the tar file:

```
# tar -xvf VRTSvcsdb.tar
```

7 Install the package:

```
# pkgadd -d . VRTSvcsdb
```

### To install the Japanese language pack on Solaris systems

1 After you have installed the agent, insert the language disc into the system's drive. Type the command:

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

2 Copy the compressed package files from the software disc to the temporary directory:

```
# cp -r ja/db2_agent/pkgs/* /tmp/install
```

3 Go to the temporary directory and unzip the compressed package file:

```
# cd /tmp/install
# gunzip VRTSjacsb.tar.gz
```

4 Extract the compressed file from the tar file:

```
# tar -xvf VRTSjacsb.tar
```

5 Install the Japanese package:

```
# pkgadd -d . VRTSjacsb
```

# Upgrading the agent for DB2 UDB

You can only upgrade the HA agent for DB2 manually. The installvcs program does not automatically upgrade the VRTSvcsdb package.

To upgrade without saving previous configuration information is to disable the agent, remove it, and re-install it.

The steps to upgrade and re-use previous configuration information follow.

### To upgrade from DB2 agent 4.0 or 4.1

1 Disable and remove the agent.

- See “Disabling the agent for DB2 UDB” on page 35.
- See “Removing the agent for DB2 UDB” on page 35.

2 From the disc that has the HA agent for DB2, add the new package.
See “Installing the agent for DB2 UDB” on page 32.

3 Copy the new "Db2udbTypes.cf" from the /etc/VRTSagents/ha/conf directory to the /etc/VRTSagents/ha/conf/config directory.

4 Update the location of the Db2udbTypes.cf file in your main.cf file’s include statement. For example, change this statement:

```
include "Db2udbTypes.cf"
```

To read:

```
include "/etc/VRTSagents/ha/conf/config/DB2udbTypes.cf"
```

5 To continue to use in-depth monitoring, use the custom monitoring sample script or any user-defined scripts.
See “Handling DB2 error codes during in-depth monitoring” on page 51.

# Disabling the agent for DB2 UDB

To disable the agent on a system, you must first change the DB2 UDB service group to an OFFLINE state on the system. You can stop the application completely, or switch the service group to another system.

### To disable the agent

1 Determine if the service group is online by entering:

```
# hagrp -state service_group -sys system_name
```

2 If the service group is online, take it offline by entering:

```
# hagrp -switch service_group -to system_name
```

Or

```
# hagrp -offline service_group -sys system_name
```

3 Stop the agent on the system by entering:

```
# haagent -stop service_group -sys system_name
```

When you get the message "`Please look for messages in the log file`," check the file /var/VRTSvcs/log/engine_A.log for a message confirming the agent has stopped.

You can also use the `ps` command to confirm the agent is stopped.

When the agent is stopped, you can remove the system, the service group, or the resource type from the VCS configuration. For more information, see the chapter on reconfiguring VCS from the command line in the *Veritas Cluster Server User's Guide*.

# Removing the agent for DB2 UDB

Before you remove the agent, you must disable it.

See "Disabling the agent for DB2 UDB" on page 35.

Perform the following instructions to remove the agent.

### To remove the agent from AIX systems

- On each system that has the agent, type:

```
# installp -u VRTSvcsdb.rte
```

### To remove the agent from Linux systems

- On each system that has the agent, type:

```
# rpm -e VRTSvcsdb
```

### To remove the agent from Solaris systems

- On each system that has the agent, type:

```
# pkgrm VRTSvcsdb
```

Chapter 4

# Configuring VCS service groups for DB2 UDB

This chapter contains the following topics:



## About configuring service groups for DB2 UDB

This chapter describes how you can configure the DB2 UDB agent. You can configure the agent using the following three methods:

- Use the VCS Cluster Manager (the Java Console)
  You can use the Cluster Manager (Java Console) to edit a service group template for the DB2 UDB agent.
  See "Configuring the DB2 UDB agent using Cluster Manager (Java Console)" on page 43.
- Use the VCS Cluster Management Console
  You can use the Cluster Management Console to configure service groups. For more information on the Cluster Management Console, see the *Veritas Cluster Server User's Guide*.

- Use VCS commands
  From the command line, you can configure the agent. For information about configuring service groups from the command line, see the *Veritas Cluster Server User's Guide*.
- Edit the main.cf file
  You can edit the main.cf file. Refer to the sample main.cf file supplied with the DB2 UDB agent for help. This method requires that you stop and restart VCS before the new configuration takes effect.
  You can use the haconf -verify command to ensure that the modified main.cf file works. For more information on this command, see the *Veritas Cluster Server User's Guide*.
  See "Configuring the DB2 UDB agent by editing the main.cf file" on page 46.

# About DB2 UDB configurations in VCS

You can configure DB2 UDG service groups in MPP and non-MPP configurations. Solaris Zones are also available for use with service groups.

## DB2 UDB service group for non-MPP configuration

Figure 4-1 illustrates the dependencies among the resources that are configured for a non-MPP DB2 UDB instance resource group.

**Figure 4-1** Dependency tree for a Db2udb resource

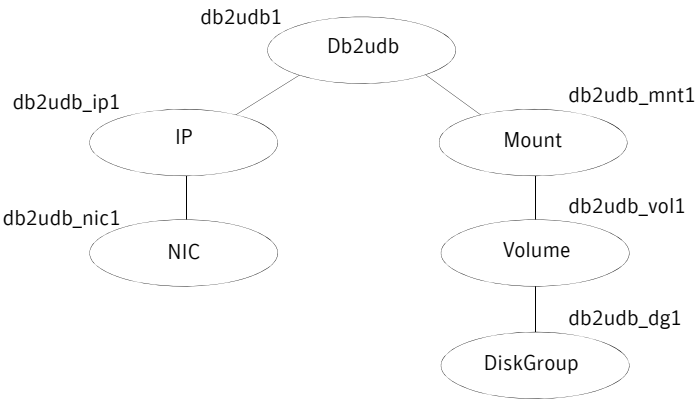


This configuration shows a service group for a Db2udb resource. The db2udb1 resource (the database) requires the IP resource and the Mount resource. The service group IP address for the DB2 UDB server is configured using the

IP resource (db2udb_ip1) and the NIC resource (db2udb_nic1). The mount resource (db2udb_mnt1) requires the Volume resource (db2udb_vol1) which in turn requires the DiskGroup resource (db2udb_dg1). You can start the service group after each of these resources is available.

# DB2 UDB service groups for MPP configuration

The DB2 UDB agent uses two service groups to support MPP configuration. These service groups are as follows:

- Parallel CVM
- DB2 failover

## Parallel CVM service group

Figure 4-2 illustrates the parallel CVM service group, which is the first of two service groups that support the MPP configuration. One CVM/Infrastructure group per cluster node exists. This service group has the CVM resource and the necessary resources for support of CFS. This group also contains all common components that DB2 needs, such as the instance's home directory, which is shared on all the cluster nodes.

**Figure 4-2** Parallel CVM service group

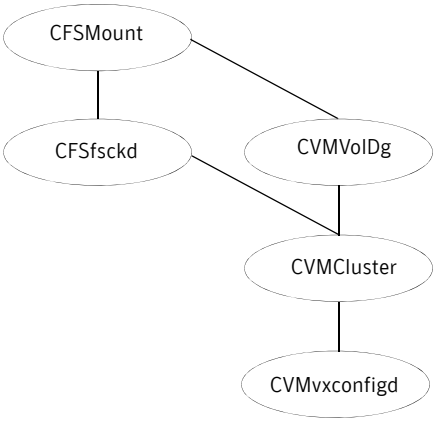

### DB2 failover service group

Figure 4-3 illustrates the DB2 failover service group, which is the second of two service groups that support the MPP configuration. This service group monitors one database partition with MPP configuration. The failover DB2 service group depends on the parallel CVM service group with online local firm dependency.

**Figure 4-3** DB2 failover service group

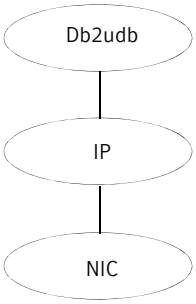


## DB2 UDB instances that are configured in Solaris zones

The following examples describe a service group that monitors the state of a DB2 instance in a Solaris zone.

Figure 4-4 illustrates the dependencies among the resources that are configured for a Db2udb resource that is configured in a Solaris zone.

**Figure 4-4** Db2udb resource that is configured in a Solaris zone

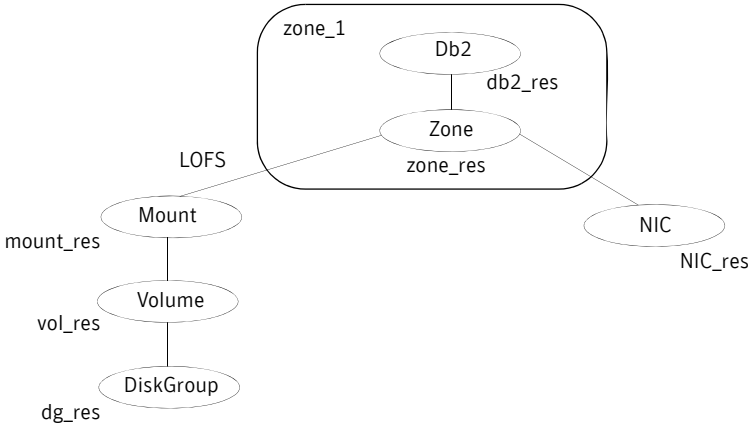


The shared disk groups and volumes in the cluster are configured as resources of type DiskGroup and Volume respectively. The volumes are mounted using the Mount agent. The Solaris Zone is monitored through a Zone resource, which is dependent on the Mount and NIC resources. The DB2 server can be started after each of these resources is brought online.

The DB2 instance's home directory is mounted in the global zone. In order to make this file system available to the non-global zone, you must execute the following command on the global zone.

The lines in the following example specify that you mount /zones/db2data in the global zone as /db2inst1 in the non-global zone named zone1. The file system type to use is LOFS. The /db2inst1 directory in this example is the home directory for the DB2 instance.

```
# zonecfg -z zone1
zonecfg:zone1> add fs
zonecfg:zone1:fs> set dir=/db2inst1
zonecfg:zone1:fs> set type=lofs
zonecfg:zone1:fs> set special=/zones/db2data
zonecfg:zone1:fs> end
```

For more information on setting up VCS to work with Solaris 10 local zones, see the appendix in the *Veritas Cluster Server User's Guide*.

# Before configuring the service group using Cluster Manager (Java Console)

Before you configure the DB2 service group, you must:

- Verify that VCS is installed and configured on all nodes in the cluster where you plan to configure the service group. For more information on installing VCS, see the *Veritas Cluster Server Installation Guide*.
- Verify that DB2 is installed and configured identically on all nodes in the cluster.
  See "Installing DB2 UDB in a VCS environment" on page 25.
- Verify that the Veritas High Availability Agent for DB2 UDB is installed on all nodes in the cluster.
  See "Installing the agent for DB2 UDB" on page 32.

## Importing the Db2udbTypes.cf file

Before you use the DB2 UDB templates, use the Cluster Manager (Java Console) to import the Db2udbTypes.cf file to the VCS engine.

### To import the Db2udbTypes.cf file

1. On one of the systems of the cluster, start the Cluster Manager (Java Console).
   ```
   # hagui&
   ```
2. Log into the cluster and wait for Cluster Explorer to launch.
3. In the Cluster Explorer window, click File and select Import Types from the drop down menu. Switch to the read/write mode if prompted.
4. In the Import Type dialog box, enter the pathname for the `Db2udbTypes.cf` file in the File Name box:
   **`/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf`**
5. Click **Import** and wait for the file to import.
6. In the Cluster Explorer window, click the Save Configuration icon.
   When the Db2udb types are imported to the VCS engine, the Db2udb agent can be configured.
   - If you use the DB2 UDB MPP configuration:
     See "Adding service group for DB2 UDB MPP using the Cluster Manager (Java Console)" on page 43.

- If you use the DB2 UDB non-MPP configuration:
  See "Adding a service group for DB2 UDB non-MPP with the Cluster Manager (Java Console)" on page 45.

# Configuring the DB2 UDB agent using Cluster Manager (Java Console)

Templates for the DB2 UDB resource groups were automatically installed when you installed the agent for DB2. Using the Cluster Manager (Java Console), you can use the template (/etc/VRTSagents/ha/Templates/Db2udbGroup.tf) to configure the DB2 UDB service group, its resources, and their attributes. You can also use the Cluster Manager (Java Console) to dynamically modify the attributes' values as necessary for your configuration.

For information on the Cluster Manager (Java Console), see the *Veritas Cluster Server User's Guide*.

## Adding service group for DB2 UDB MPP using the Cluster Manager (Java Console)

If you have imported the Db2udbTypes.cf file, you can use the template (/etc/VRTSagents/ha/Templates/Db2udbGroup.tf) to configure a service group.

See "Importing the Db2udbTypes.cf file" on page 42.

After you log into the Cluster Manager, the Status tab indicates that the CVM service group is online on each system in the cluster. The CVM service group is automatically configured when you complete the installation of the Storage Foundation Cluster File System (SFCFS) software.

**To add the service group for the DB2 UDB MPP database**

1. In the Cluster Explorer window, click the **Add Service Group** icon on the toolbar.
2. In the Add Service Group window, enter the name of the service group you want to create for the service group. For example, enter `db2mpp_grp1`. Do not press **Return** or **Enter**.
3. In the Available Systems box, double-click the systems that you want in your configuration.
4. Select the system where you want to automatically start the service group. In the window that shows the systems in the configuration, select the check box next to the system.
5. Click the **Failover** radio button to specify the Service Group Type.

6 Click the Templates button.

7 In the Select Templates window, select db2udb_mpp_grp from the list in the Templates box. The Dependency graph information and the Types information should change to reflect the template choice. Click **OK**. In the Add Service Group window, the name of the template is now shown as selected.

8 Click **OK** on at the bottom of the Add Service Group window. The group is added. On the left pane of the Cluster Manager (Java Console) window, the service group is under the CVM service group. On the Status tab, the group is shown Offline on each system.

9 In the left pane, double-click the db2mpp_grp1 service group. The types of resources that you can configure for the group appear: Db2udb, IP, and NIC.

10 Double-click the Db2udb resource type. Select the resource, db2udb, which is under the Db2udb type. Click the **Properties** tab.

11 On the Properties tab for the db2udb resource, a list of Type Specific Attributes is shown. click the **Edit** icon for each attribute you want to configure. In the Edit Attribute window, enter the necessary attribute value information. For example, enter the db2inst1 as the value for DB2InstOwner.

12 Assign values for the IP and NIC resources in the same manner as you assigned values to the db2udb resource: double-click the type to display the resource and select the resource. With the Properties tab visible, you can edit the Type Specific Attributes for each resource.

13 Right-click the db2mpp_grp1 service group in the left pane. Click Link in the drop-down menu. The Link Service Groups window shows the following:

- The Parent Group as db2mpp_grp1
- The Child group as CVM
- The Relationship as “online local”
- The Dependency Type as “firm”

14 Click **OK** to create the dependency link.

15 Click the **Save Configuration** icon.

16 Enable the db2udb and IP resources. Right click a resource and select **Enabled** in the drop-down menu. If necessary, make the configuration read/write.

17 Click the **Online Service Group** icon.

18 In the window, select the service group and the system that you want to bring online. Click **OK**.

# Adding a service group for DB2 UDB non-MPP with the Cluster Manager (Java Console)

If you have imported the Db2udbTypes.cf file, you can use the Db2udb_Group template to configure a service group.

See "Importing the Db2udbTypes.cf file" on page 42.

**To configure a service group if you have imported the Db2udbTypes.cf file**

1. In the Cluster Explorer window, answer **No** when prompted to use the configuration wizard. Note that if you choose to use the wizard, the steps that follow are similar.
2. In the Cluster Explorer window, click the **Add Service Group** icon on the toolbar.
3. In the Add Service Group window, enter the name of the service group you want to create for the service group. For example, enter db2_group1. Do not press **Return** or **Enter**.
4. From the systems in the Available Systems box, double-click those that you want in your configuration.
5. Select the system where you want to automatically start the service group. In the window that shows the systems in the configuration, select the check box next to the system.
6. Click the **Failover** radio button to specify the Service Group Type.
7. Click the **Templates** button.
8. In the Select Templates window, select db2udb_grp from the list in the Templates box. The Dependency graph information and the Types information should change to reflect the template choice. Click **OK**. In the Add Service Group window, the name of the template is now shown as selected.
9. Click **OK** on at the bottom of the Add Service Group window. The group is added. On the left pane of the Cluster Manager (Java Console) window, the service group is shown under the cluster name. On the Status tab, the group is shown Offline on each system.
10. In the left pane, double-click the **db2_group1** service group. The types of resources that you can configure for the group appear: Db2udb, DiskGroup, IP, Mount, NIC, and Volume.
11. Double-click the **Db2udb** resource type. Select the resource, **db2udb**, which is under the Db2udb type. Click the **Properties** tab.

12 On the Properties tab for the db2udb resource, a list of Type Specific Attributes is shown. Click the **Edit** icon for each attribute you want to configure. In the Edit Attribute window, enter the necessary attribute value information. For example, enter the db2inst1 as the value for DB2InstOwner.
DB2InstOwner and DB2InstHome are required attributes. You must edit these attributes.

13 Assign values for the DiskGroup, IP, Mount, NIC, and Volume resources in the same manner as you assigned values to the db2udb resource: double-click the type to display the resource and select the resource. With the Properties tab visible, you can edit the Type Specific Attributes.
For a list of the required attributes, and their descriptions, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

14 Enable the resources in db2_group1. Right click each resource and select **Enabled** in the drop-down menu. If necessary, make the configuration read/write.

15 Click the **Save and Close Configuration** icon.

16 Click the **Online Service Group** icon.

17 In the window, select the service group and the system that you want to bring online. Click the system where you want to bring it online. Click **OK**. Click **Yes** at the confirmation question.

# Configuring the DB2 UDB agent by editing the main.cf file

The VCS agent for DB2 comes with three sample VCS configuration files, which are in the /etc/VRTSagents/ha/conf/Db2udb/sample_db2udb directory. The samples are as follows:

- A single-partition instance configuration
- A multi-partition instance SMP configuration
- A multi-partition instance MPP configuration.

The appropriate file can be used as reference to directly modify your present main.cf configuration file. When you use this method, you must stop and restart VCS to implement the configuration.

### To prepare to edit the main.cf file

1. Log in to System A as root.
2. Save your existing configuration to prevent any changes while you modify the main.cf file:

   ```
   # haconf -dump -makero
   ```

3. Make sure to stop VCS while you edit main.cf. Use the `hastop` command to stop the VCS engine on all systems and leave the resources available:

   ```
   # hastop -all -force
   ```

4. Make a backup copy of the main.cf file:

   ```
   # cd /etc/VRTSvcs/conf/config
   # cp main.cf main.cf.orig
   ```

Depending on your configuration, go to one of the following sections that describe configuring the agent for DB2.

## Configuring the agent to use the DB2 UDB MPP configuration

Edit the main.cf file. Use /etc/VRTSagents/ha/conf/Db2udb/sample_db2udb/main.cf.MPP for reference. Notice that CVM service group is present in the configuration file.

### To configure the agent to use the DB2 UDB MPP configuration

1. Add the fully qualified path to the Db2udbTypes.cf file.

   ```
   include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"
   ```

2. Create service groups for the Db2udb resources. Refer to the sample configuration file. The example shows four DB2 MPP service groups and a CVM service group. See the following examples:
   - See "MPP main.cf configuration for AIX" on page 63.
   - See "MPP main.cf configuration for Linux" on page 70.
   - See "MPP main.cf configuration for Solaris" on page 77.
3. In the DB2 MPP service groups, include the definitions for the Db2udb, IP, and NIC resources. Assign values to the attributes for the resources to match the parameters of your configuration.
   For more information, see:
   - "Db2udb resource type attributes" on page 56
   - Sample configuration files.
   - For information about IP and NIC resources, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

4 Assign the online local firm service group dependency of the db2udb service group for the cvm service group. For example:

```
requires group cvm online local firm
```

5 Immediately following the service group dependency, assign dependencies for the newly created resources. Refer to the appropriate sample configuration file. For the group `db2mpp_grp0`, enter:

```
db2udb0 requires Db2_IP0
Db2_IP0 requires Db2_NIC0
```

For more information on assigning dependencies, see the *Veritas Cluster Server User's Guide*.

6 Save and close the file.

## Configuring the agent to use the DB2 UDB, non-MPP configurations

Edit the main.cf file. Use /etc/VRTSagents/ha/conf/Db2udb/sample_db2udb/main.cf.EE or /etc/VRTSagents/ha/conf/Db2udb/sample_db2/main.cf.EEE for reference.

**To configure the agent to use the DB2 UDB, non-MPP configurations**

1 Add the fully qualified path to the Db2udbTypes.cf file.

```
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"
```

2 Create a service group for the DB2 UDB resources.
This example that shows a group named "db2_grp1" in which two partitions are defined.
See "Non-MPP main.cf configuration for AIX" on page 61.

3 Include all resources in the service groups, including the Db2udb, DiskGroup, IP, Mount, NIC, and Volume resources. Assign values to the attributes for the resources to match the parameters of your configuration.
For more information, see:

- "Db2udb resource type attributes" on page 56
- The sample configuration files

For information about the DiskGroup, IP, Mount, NIC, and Volume resources, see the *Veritas Cluster Server Bundled Agents Reference Guide*.

4 Assign dependencies for the newly created resources. For the group db2udb_grp1, enter:

```
db2udb1 requires db2udb_ip1
db2udb1 requires db2udb_mnt1
db2udb_ip1 requires db2udb_nic1
db2udb_mnt1 requires db2udb_vol1
db2udb_vol1 requires db2udb_dg1
```

And for group db2udb_grp3, enter:

```
db2udb3 requires db2udb_ip3
db2udb3 requires db2udb_mnt3
db2udb_ip3 requires db2udb_nic3
db2udb_mnt3 requires db2udb_vol3
db2udb_vol3 requires db2udb_dg3
```

For more information, refer to the appropriate sample configuration file. For more information on assigning dependencies. see the *Veritas Cluster Server User's Guide.*

5 Save and close the file.

## Verifying the configuration

After editing the main.cf file for you configuration, check the configuration.

### To check the configuration

1 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# cd /etc/VRTSvcs/conf/config
# hacf -verify .
```

2 Start the VCS engine on System A:

```
# hastart
```

3 Type the hastatus command:

```
# hastatus
```

4 When "LOCAL_BUILD" is listed in the message column, start VCS on System B:

```
# hastart
```

5 Verify that all DB2 UDB service group resources are brought online on System A:

```
# hagrp -display
```

6 Take the service groups offline on System A and verify that all resources are stopped:

```
# hagrp -offline db2udb_grp1 -sys sysa
# hagrp -offline db2udb_grp3 -sys sysa
# hagrp -display
```

7 Bring the service groups online again on System A and verify that all resources are available:

```
# hagrp -online db2udb_grp1 -sys sysa
# hagrp -online db2udb_grp3 -sys sysa
# hagrp -display
```

8 Switch the DB2 UDB service group to System B:

```
# hagrp -switch db2udb_grp1 -to sysb
# hagrp -switch db2udb_grp3 -to sysb
```

9 Verify that all DB2 UDB service group resources are brought online on System B:

```
# hagrp -display
```

10 On all the systems, look at the following log files for any errors or status:

```
/var/VRTSvcs/log/engine_A.log
/var/VRTSvcs/log/Db2udb_A.log
```

# Setting up in-depth monitoring of DB2 UDB instance

To dynamically reconfigure the Veritas agent for DB2, use Cluster Manager (Java Console) or the VCS command line. The following description of changing the configuration to include in-depth monitoring shows the use of VCS commands from the command line. For information on reconfiguring VCS from the command line, see the *Veritas Cluster Server User's Guide*.

## Enabling in-depth monitoring of DB2 UDB instance

Shallow monitoring of a DB2 UDB instance involves checking the exit status of the `db2gcf` command.

In contrast, in-depth monitoring provides a higher level of confidence in the availability of the instance or partition and its database. It makes additional queries to the database to verify whether the database is available.

## Enabling in-depth monitoring from the command line

You can dynamically configure in-depth monitoring. Symantec recommends that you successfully run DB2 UDB with the agent's default (shallow) monitoring before you start the in-depth monitoring. In the MPP configuration, make sure the database can be accessible locally by the database partition.

For locales other than English, you need to add the following lines to the $INSTHOME/sqllib/*userprofile* file.

The following example adds Japanese language support for Solaris:

```
LANG=ja
export LANG
```

The following example adds Japanese language support for AIX:

```
export LANG=Ja_JP
```

### To start the in-depth monitor for a given instance

1 Make the VCS configuration writable: **# haconf -makerw**

2 Freeze the service group so VCS does not perform actions automatically based on an incomplete reconfiguration:

```
# hagrp -freeze db2udb_grp1
```

3 Enable in-depth monitoring using the command:

```
hares -modify resource DatabaseName name
hares -modify resource IndepthMonitor 1
```

For example:

```
# hares -modify db2udb DatabaseName SAMPLE
# hares -modify db2udb IndepthMonitor 1
# haconf -dump -makero
# hagrp -unfreeze db2udb_grp1
```

**Note:** You need to have custom monitoring scripts.
See "IndepthMonitor" on page 58.

## Handling DB2 error codes during in-depth monitoring

The agent for DB2 comes with enhanced handling of DB2 errors during in-depth monitoring. The agent classifies DB2 errors according to their severity and associates predefined actions with each error code.

You can create a custom error handling file, db2error.dat. The file lists the DB2 errors and the associated actions that you want the agent to take when it encounters an error.

The file stores information in the following format:

```
SQL_error_string:action_to_be_taken
```

For example:

```
SQL1034N: IGNORE
SQL1039N: WARN
SQL1234N: FAILOVER
```

**Table 4-1** Available actions for in-depth monitoring

| Action | Description |
|---|---|
| IGNORE | Ignores the error. |
| UNKNOWN | Marks the resource state as UNKNOWN and sends a notification if the Notifier resource is configured. For more information about VCS notification, see the *Veritas Cluster Server User's Guide*.<br>This action is typically associated with configuration errors. |
| WARN | Marks the resource state as ONLINE and sends a notification if the Notifier resource is configured.<br>This action is typically associated with low-severity errors. |

**Table 4-1** Available actions for in-depth monitoring

| Action | Description |
|---|---|
| FAILOVER<br>(Default) | Marks the resource state as OFFLINE. This faults the service group, which fails over to the next available system.<br>This action is the agent's default behavior. If the DB2 error code that the agent encounters does not exist in the db2error.dat file, then the agent assumes this default behavior. |
| NOFAILOVER | Freezes the service group temporarily and marks the resource state as OFFLINE. The agent also sends a notification if the Notifier resource is configured.<br>This action is typically associated with the errors that are not system-specific. For example, if a database was corrupted, failing it over to another node does not help. |

## Disabling in-depth monitoring

You can dynamically disable in-depth monitoring.

### To dynamically disable in-depth monitoring

1. Make the VCS configuration writable:
   ```
   # haconf -makerw
   ```
2. Freeze the service group so VCS does not perform actions automatically based on an incomplete reconfiguration:
   ```
   # hagrp -freeze db2udb_grp1
   ```
3. Assign the InDepthMonitor attribute a null value to disable in-depth monitoring. Use the command:
   ```
   hares -modify resource IndepthMonitor 0
   ```
   For example:
   ```
   # hares -modify db2udb IndepthMonitor 0
   # haconf -dump -makero
   # hagrp -unfreeze db2udb_grp1
   ```

Chapter 5

# Administering VCS service groups for DB2 UDB

This chapter contains the following topics:



## About administering service groups for DB2 UDB

You can administer service groups with the Cluster Manager (Java Console), the Cluster Management Console, or from the command line.

## Bringing the service group online

Perform the following procedure from Cluster Manager (Java Console) to bring the service group online.

**To bring a service group online**

1. From Cluster Explorer, click the **Service Groups** tab in the configuration tree.
2. Right-click the service group and click **Enable Resources**.
3. Right-click the service group, pause over **Enable**, and select either the system or all the systems where you want to enable the service group.

4 Save and close the configuration. Click **File** > **Save Configuration**, then **Close Configuration**.

5 Right-click the service group, pause over Online and select the system where you want to bring the service group online.

# Taking the service group offline

Perform the following procedure from Cluster Manager (Java Console) to take the service group offline.

### To take a service group offline

1 In the Cluster Explorer configuration tree with the Service Groups tab selected, right-click the service group that you want to take offline.

2 Choose **Offline,** and select the appropriate system from the pop-up menu.

# Switching the service group

The process of switching a service group involves taking it offline on its current system and bringing it online on another system. Perform the following procedure from Cluster Manager (Java Console) to switch the service group.

### To switch a service group

1 In the Cluster Explorer configuration tree with the Service Groups tab selected, right-click the service group.

2 Choose **Switch To**, and select the appropriate system from the pop-up menu.

# Disabling the agent

For instructions on disabling the agent, see:

See "Disabling the agent for DB2 UDB" on page 35.

# Modifying the agent configuration

You can dynamically configure the Veritas High Availability Agent for DB2 UDB in several ways, including the configuration wizard, the command-line interface, Cluster Manager Java Console, and the Cluster Management Console. For more information, see the *Veritas Cluster Server User's Guide*.

Appendix A

# Resource type definition for DB2 UDB

This appendix contains the following topics:



## About the resource type and attribute definitions for DB2 UDB

Configuring the DB2 UDB agent involves assigning values to the DB2 UDB resource type attributes. The attributes are described in the following table. The resource type definition file, Db2udbTypes.cf, is also shown for reference. The sample main.cf configuration files are shown in Appendix B, "Sample configuration files".

### DB2 UDB type definition file: Db2udbTypes.cf for Solaris

The following type definition is for Solaris.

```
type Db2udb (
    static str ContainerType = Zone
    static str AgentDirectory = "/opt/VRTSagents/ha/bin/Db2udb"
    static str AgentFile = "/opt/VRTSagents/ha/bin/Db2udb/
    Db2udbAgent"
    static keylist SupportedActions = { VRTS_GetInstanceName,
    VRTS_GetRunningServices }
    static int CleanTimeout = 240
    static int MonitorTimeout = 240
    static int OfflineTimeout = 240
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 180
```

```
    static int OnlineWaitLimit = 1
    static int RestartLimit = 3
    static int ToleranceLimit = 1
    static str ArgList[] = { DB2InstOwner, DB2InstHome,
    IndepthMonitor, DatabaseName, NodeNumber, StartUpOpt,
    ShutDownOpt, AgentDebug, Encoding, WarnOnlyIfDBQueryFailed,
    LastWarningDay, ContainerName }
    str DB2InstOwner
    str DB2InstHome
    int IndepthMonitor
    str DatabaseName
    int NodeNumber
    str StartUpOpt = START
    str ShutDownOpt = STOP
    boolean AgentDebug = 0
    str Encoding
    boolean WarnOnlyIfDBQueryFailed = 1
    temp str LastWarningDay
    str ContainerName
)
```

# Db2udb resource type attributes

The DB2 resource has several required and optional attributes.

**Table A-1** Required attributes for the agent for DB2

| Required attributes | Description |
|---|---|
| DB2InstHome | Path to DB2 UDB instance home directory that contains critical data and configuration files for the DB2 instance.<br>Type and dimension: string-scalar |
| DB2InstOwner | User ID of Instance Owner that starts a DB2 UDB instance. Each instance requires a unique user ID.<br>Type and dimension: string-scalar<br>**Caution:** Incorrect changes to this attribute can result in DB2 entering an inconsistent state. |

**Table A-2** Optional attributes for the agent for DB2

| Optional attributes | Description |
|---|---|
| DatabaseName | Name of the database for in-depth monitoring; required if in-depth monitor is enabled (IndepthMonitor = 1).<br>**Note:** Be careful when you change the DataBase name attribute as you can fault all the partitions in the database. Do not change the DataBaseName attribute to an invalid or incorrect value.<br>Type and dimension: string-scalar |
| NodeNumber | Node number or partition number of the database. Used when monitoring a specific database partition.<br>Default: 0<br>Type and dimension: integer-scalar |
| StartUpOpt | Provides start up options. The allowed values are: START, ACTIVATEDB, or CUSTOM.<br>■ START (default)<br>Starts the DB2 instance or partition.<br>■ ACTIVATEDB<br>Performs activate database command after db2 processes start.<br>■ CUSTOM<br>The agent leaves all the online operation completely to the user when the StartUpOpt attribute is set to CUSTOM. It looks for a file named start_custom_$db2instance_$nodenum in the /opt/VRTSagents/ha/bin/Db2udb directory. If this file exists and is executable, it executes this customized online file instead.<br>Example:<br>To customize the online operation for partition/nodenum 1 for the db2 instance named db2inst1, the agent for DB2 runs this customized file start_custom_db2inst1_1 under the /opt/VRTSagents/ha/bin/Db2udb directory.<br>Type and dimension: string-scalar |

**Table A-2** Optional attributes for the agent for DB2

| Optional attributes | Description |
|---|---|
| ShutDownOpt | The allowed values for this attribute are STOP and CUSTOM.<br>■ STOP<br>Shuts the Db2 instance or partition down in the usual way.<br>■ CUSTOM<br>Leaves all the offline operation completely to the user when the ShutDownOpt is set to CUSTOM. It looks for a file named stop_custom_$db2instance_$nodenum in the /opt/VRTSagents/ha/bin/Db2udb directory.<br>If this file exists and is executable, it executes this customized offline file instead.<br>Example:<br>You want to customize the offline operation for partition/ nodenum 0 for the db2 instance named db2inst1. You have the agent for DB2 run this customized file: stop_custom_db2inst1_0. The file is in the /opt/VRTSagents/ ha/bin/Db2udb directory.<br>Type and dimension: string-scalar |
| IndepthMonitor | Set the value of the IndepthMonitor attribute to 1 to enable in-depth monitoring. Before this release, IndepthMonitor performed a default SQL query to the database. In 5.0, this default query no longer exists. The agent now looks for the custom_monitor_$db2instance_$nodenum file in the /opt/VRTSagents/ha/bin/Db2udb directory.<br>It executes this customized indepth monitor file if the file exists and is executable. You can find samples of custom monitor scripts in the sample_db2udb directory.<br>Type and dimension: string-integer |
| Encoding | Specifies the operating system encoding corresponding to DB2 UDB encoding for display of DB2 UDB output. For example, if the environment variable LANG is set to “ja,” then “eucJP” is the Solaris value for Encoding. Refer to DB2 UDB and Solaris documentation for respective encoding values. The default is “”.<br>Type and dimension: string-scalar |
| AgentDebug | When the value of this attribute is 1, it causes the agent to log additional debug messages.<br>Type and dimension: boolean-scalar |

**Table A-2** Optional attributes for the agent for DB2

| Optional attributes | Description |
|---|---|
| WarnOnlyIfDBQueryFailed | This attribute either logs SQL errors, or checks the errors for special handling.<br>Set the value of the WarnOnlyIfDBQueryFailed attribute to 1 to enable it. When this attribute is enabled, it ignores all SQL errors and logs a warning message in the agent log once a day.<br>Set the value of the WarnOnlyIfDBQueryFailed attribute to 0 to disable it. When disabled, it checks if an error code has special handling in the db2error.dat file. If the error code does not exist in the db2error.dat file, then it returns OFFLINE for monitor. Otherwise, it follows the action of that particular error code in the db2error.dat file.<br>Type and dimension: boolean-scalar |
| ContainerName | Name of the Solaris zone (Solaris 10 only)<br>Type and dimension: string-scalar |

**Table A-3** Internal attributes for the agent for DB2

| Required attributes | Description |
|---|---|
| AgentDirectory | Specifies the location of other files and scripts that are related to the agent.<br>Do not use. For internal use only. |

Appendix B

# Sample configuration files

This appendix contains the following topics:



## AIX sample configuration files

This section shows example DB2 UDB configurations for AIX.

### Non-MPP main.cf configuration for AIX

The following main.cf configuration file reflects DB2 UDB in an ESE multi-partition instance SMP environment. Two database partitions are shown.

```
include "types.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2_clus (
    UserNames = { admin = "cDRpdxPmHpzS." }
    Administrators = { admin }
    CounterInterval = 5
    )

system sysA (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
    ActionTimeLimit = 0, Action = NONE,
    NotifyThreshold = 0, NotifyTimeLimit = 0 }
    )

system sysB (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
    ActionTimeLimit = 0, Action = NONE,
    NotifyThreshold = 0, NotifyTimeLimit = 0 }
```

```
    )
group db2_grp1 (
    SystemList = { sysA = 0, sysB = 1 }
    AutoStartList = { sysA }
    )

    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = DWCNTRL
        NodeNumber = 0
        )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 0
        NodeNumber = 1
        )

    DiskGroup db2dg1 (
        DiskGroup = db2dg1
        )

    IP db2ip1 (
        Device = en0
        Address = "166.98.9.188"
        NetMask = "255.255.220.0"
        )

    Mount db2mnt1 (
        MountPoint = "/db2_mnt/db2inst1"
        BlockDevice = "/dev/vx/dsk/db2dg1/db2dg1home"
        FSType = vxfs
        MountOpt = rw
        )

    NIC db2nic1 (
        Device = en0
        NetworkHosts = { "166.98.128.180" }
        )

    Volume db2vol1 (
        Volume = db2dg1home
        DiskGroup = db2dg1
        )

    db2ip1 requires db2nic1
    db2mnt1 requires db2vol1
    db2udb1 requires db2ip1
    db2udb1 requires db2mnt1
```

```
    db2vol1 requires db2dg1
    db2udb2 requires db2ip1
    db2udb2 requires db2mnt1
```

## MPP main.cf configuration for AIX

The following configuration file reflects DB2 UDB in an ESE multi-partition instance MPP environment. Four database partitions are shown. One partition is configured on each cluster node. Each database service group depends on the same CVM service group, which manages the shared storage in the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2_aix_mpp (
    UserNames = { admin = gpqIpkPmqLqqOyqKpn }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
    CounterInterval = 5
    )

system sysA (
)

system sysB (
)

group cvm (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sysA, sysB, sysC, sysD }
    )

    CFSMount db2cfsmnt (
        MountPoint = "/db2_mnt/db2inst1"
        BlockDevice = "/dev/vx/dsk/cdb2dg1/cdb2dg1home"
        MountOpt = "cluster"
        NodeList = { sysA, sysB, sysC, sysD }
        )

    CFSfsckd vxfsckd (
    )
```

```
    CVMCluster cvm_clus (
        CVMClustName = db2_aix_mpp
        CVMNodeId = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
        CVMTransport = gab
        CVMTimeout = 200
        )

    CVMVolDg db2dg (
        CVMDiskGroup = cdb2dg1
        CVMActivation = sw
        )

    CVMVxconfigd cvm_vxconfigd (
        Critical = 0
        CVMVxconfigdArgs = { syslog }
        )

    cvm_clus requires cvm_vxconfigd
    db2cfsmnt requires db2dg
    db2cfsmnt requires vxfsckd
    db2dg requires cvm_clus
    vxfsckd requires cvm_clus

    // resource dependency tree
    //
    //          group cvm
    //          {
    //          CFSMount db2cfsmnt
    //              {
    //              CVMVolDg db2dg
    //               {
    //               CVMCluster cvm_clus
    //                  {
    //                   CVMVxconfigd cvm_vxconfigd
    //                   }
    //              }
    //               CFSfsckd vxfsckd
    //                   {
    //                   CVMCluster cvm_clus
    //                       {
    //                       CVMVxconfigd cvm_vxconfigd
    //                       }
    //                   }
    //               }
    //          }

group db2mpp_grp0 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysA }
    )
```

```
    Db2udb db2udb0 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = sample
        )

    IP Db2_IP0 (
        Device = en0
        Address = "11.192.10.32"
        NetMask = "255.255.244.0"
        )

    NIC mynic0 (
        Device = en0
        NetworkHosts = { "11.192.11.90" }
        )

    requires group cvm online local firm
    Db2_IP0 requires mynic0
    db2udb0 requires Db2_IP0

    // resource dependency tree
    //
    //          group db2mpp_grp0
    //          {
    //          Db2udb db2udb0
    //               {
    //               IP Db2_IP0
    //                   {
    //                   NIC mynic0
    //                   }
    //               }
    //          }

group db2mpp_grp1 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysB }
    )

    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = test1
        NodeNumber = 1
        )
```

```
    IP Db2_IP1 (
        Device = en0
        Address = "11.192.10.33"
        NetMask = "255.255.244.0"
        )

    NIC mynic1 (
        Device = en0
        NetworkHosts = { "11.192.11.90" }
        )

    requires group cvm online local firm
    Db2_IP1 requires mynic1
    db2udb1 requires Db2_IP1

    // resource dependency tree
    //
    //          group db2mpp_grp1
    //          {
    //          Db2udb db2udb1
    //               {
    //               IP Db2_IP1
    //                   {
    //                   NIC mynic1
    //                   }
    //               }
    //          }

group db2mpp_grp2 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysC }
    )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = test2
        NodeNumber = 2
        )

    IP Db2_IP2 (
        Device = en0
        Address = "11.192.10.33"
        NetMask = "255.255.244.0"
        )

    NIC mynic2 (
        Device = en0
        NetworkHosts = { "11.192.11.90" }
        )
```

```
    requires group cvm online local firm
    Db2_IP2 requires mynic2
    db2udb2 requires Db2_IP2


    // resource dependency tree
    //
    //          group db2mpp_grp2
    //          {
    //          Db2udb db2udb2
    //               {
    //               IP Db2_IP2
    //                   {
    //                   NIC mynic2
    //                   }
    //               }
    //          }

group db2mpp_grp3 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysD }
    )

    Db2udb db2udb3 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        NodeNumber = 3
        )

    IP Db2_IP3 (
        Device = en0
        Address = "11.192.10.33"
        NetMask = "255.255.244.0"
        )

    NIC mynic3 (
        Device = en0
        NetworkHosts = { "11.192.11.90" }
        )

    requires group cvm online local firm
    Db2_IP3 requires mynic3
    db2udb3 requires Db2_IP3


    // resource dependency tree
    //
    //          group db2mpp_grp3
    //          {
    //          Db2udb db2udb3
    //               {
    //               IP Db2_IP3
```

```
    //                      {
    //                      NIC mynic3
    //                      }
    //                  }
    //          }
```

# Linux sample configuration files

This section shows example DB2 UDB configurations for Linux.

## Non-MPP main.cf configuration for Linux

The following main.cf configuration file reflects DB2 UDB in an ESE multi-partition instance SMP environment. Two database partitions are shown.

```
include "types.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster vcs (
    CounterInterval = 5
    )

system vcstc1 (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
    ActionTimeLimit = 0, Action = NONE, NotifyThreshold = 0,
    NotifyTimeLimit = 0 }
    )

system vcstc2 (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
    ActionTimeLimit = 0, Action = NONE, NotifyThreshold = 0,
    NotifyTimeLimit = 0 }
    )

group db2udb_grp1 (
    SystemList = { vcstc1= 0, vcstc2 = 1 }
    AutoStartList = { vcstc1 }
    )

    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2inst1"
        IndepthMonitor = 1
        DatabaseName = SAMPLE
        NodeNumber = 0
        )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2inst1"
```

```
    IndepthMonitor = 0
    NodeNumber = 1
    )

DiskGroup db2udb_dg1 (
    DiskGroup = db2_dg1
    )

IP db2udb_ip1 (
    Device = eth0
    Address = "166.98.9.163"
    NetMask = "255.255.252.0"
    )

Mount db2udb_mnt1 (
    MountPoint = "/db2inst1"
    BlockDevice = "/dev/vx/dsk/db2_dg1/inst1_vol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-n"
    )

NIC db2udb_nic1 (
    Device = eth0
    )

Volume db2udb_vol1 (
    Volume = inst1_vol
    DiskGroup = db2_dg1
    )

db2udb1 requires db2udb_ip1
db2udb1 requires db2udb_mnt1
db2udb2 requires db2udb_ip1
db2udb2 requires db2udb_mnt1
db2udb_ip1 requires db2udb_nic1
db2udb_mnt1 requires db2udb_vol1
db2udb_vol1 requires db2udb_dg1
```

```
// resource dependency tree
//
//          group db2udb_grp1
//          {
//          Db2udb db2udb1
//                {
//                IP db2udb_ip1
//                      {
//                      NIC db2udb_nic1
//                      }
//                Mount db2udb_mnt1
//                      {
//                      Volume db2udb_vol1
//                            {
//                            DiskGroup db2udb_dg1
//                            }
//                      }
//                }
//          Db2udb db2udb2
//                {
//                IP db2udb_ip1
//                      {
//                      NIC db2udb_nic1
//                      }
//                Mount db2udb_mnt1
//                      {
//                      Volume db2udb_vol1
//                            {
//                            DiskGroup db2udb_dg1
//                            }
//                      }
//                }
//          }
```

## MPP main.cf configuration for Linux

The following configuration file reflects DB2 UDB in an ESE multi-partition instance MPP environment. Four database partitions are shown. One partition is configured on each cluster node. Each database service group depends on the same CVM service group, which manages the shared storage in the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2_mpp (
    CounterInterval = 5
    )
```

```
system vcstc1 (
    )

system vcstc2 (
    )

system vcstc3 (
    )

system vcstc4 (
    )

group cvm (
    SystemList = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { vcstc1, vcstc2, vcstc3, vcstc4 }
    )

    CFSMount db2cfsmnt (
        MountPoint = "/db2_mnt/db2inst1"
        BlockDevice = "/dev/vx/dsk/cdb2dg1/cdb2dg1home"
        )

    CFSfsckd vxfsckd (
        )

    CVMCluster cvm_clus (
        CVMClustName = db2_mpp
        CVMNodeId = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3
        }
        CVMTransport = gab
        CVMTimeout = 200
        )

    CVMVolDg db2dg (
        CVMDiskGroup = cdb2dg1
        CVMVolume = { cdb2dg1home }
        CVMActivation = sw
        )

    CVMVxconfigd cvm_vxconfigd (
        Critical = 0
        CVMVxconfigdArgs = { syslog }
        )

    cvm_clus requires cvm_vxconfigd
    db2cfsmnt requires db2dg
    db2cfsmnt requires vxfsckd
    db2dg requires cvm_clus
    vxfsckd requires cvm_clus
```

```
    // resource dependency tree
    //
    //          group cvm
    //          {
    //          CFSMount db2cfsmnt
    //                {
    //                CVMVolDg db2dg
    //                      {
    //                      CVMCluster cvm_clus
    //                            {
    //                            CVMVxconfigd cvm_vxconfigd
    //                            }
    //                      }
    //                CFSfsckd vxfsckd
    //                      {
    //                      CVMCluster cvm_clus
    //                            {
    //                            CVMVxconfigd cvm_vxconfigd
    //                            }
    //                      }
    //                }
    //          }

group db2mpp_grp0 (
    SystemList = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3 }
    AutoStartList = { vcstc1 }
    AutoStart = 1
    )

    Db2udb db2udb0 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = SAMPLE
        NodeNumber = 0
        )

    IP Db2_IP0 (
        Device = eth0
        Address = "10.118.2.144"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC0 (
        Device = eth0
        )

    requires group cvm online local firm
    Db2_IP0 requires Db2_NIC0
    db2udb0 requires Db2_IP0
```

```
    // resource dependency tree
    //
    //          group db2mpp_grp0
    //          {
    //          Db2udb db2udb0
    //                {
    //                IP Db2_IP0
    //                      {
    //                      NIC Db2_NIC0
    //                      }
    //                }
    //          }

group db2mpp_grp1 (
    SystemList = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3 }
    AutoStartList = { vcstc2 }
    AutoStart = 1
    )
    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = TEST1
        NodeNumber = 1
        )

    IP Db2_IP1 (
        Device = eth0
        Address = "10.118.2.145"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC1 (
        Device = eth0
        )

    requires group cvm online local firm
    Db2_IP1 requires Db2_NIC1
    db2udb1 requires Db2_IP1

    // resource dependency tree
    //
    //          group db2mpp_grp1
    //          {
    //          Db2udb db2udb1
    //                {
    //                IP Db2_IP1
    //                      {
    //                      NIC Db2_NIC1
    //                      }
    //                }
    //          }
```

```
group db2mpp_grp2 (
    SystemList = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3 }
    AutoStartList = { vcstc3 }
    AutoStart = 1
    )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = TEST2
        NodeNumber = 2
        )

    IP Db2_IP2 (
        Device = eth0
        Address = "10.118.2.146"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC2 (
        Device = eth0
        )

    requires group cvm online local firm
    Db2_IP2 requires Db2_NIC2
    db2udb2 requires Db2_IP2


    // resource dependency tree
    //
    //          group db2mpp_grp2
    //          {
    //          Db2udb db2udb2
    //                {
    //                IP Db2_IP2
    //                      {
    //                      NIC Db2_NIC2
    //                      }
    //                }
    //          }

group db2mpp_grp3 (
    SystemList = { vcstc1 = 0, vcstc2 = 1, vcstc3 = 2, vcstc4 = 3 }
    AutoStartList = { vcstc4 }
    AutoStart = 1
    )
```

```
Db2udb db2udb3 (
    DB2InstOwner = db2inst1
    DB2InstHome = "/db2_mnt/db2inst1"
    NodeNumber = 3
    )
IP Db2_IP3 (
    Device = eth0
    Address = "10.118.2.147"
    NetMask = "255.255.248.0"
    )

NIC Db2_NIC3 (
    Device = eth0
    )

requires group cvm online local firm
Db2_IP3 requires Db2_NIC3
db2udb3 requires Db2_IP3


// resource dependency tree
//
//          group db2mpp_grp3
//          {
//          Db2udb db2udb3
//                {
//                IP Db2_IP3
//                      {
//                      NIC Db2_NIC3
//                      }
//                }
//          }
```

# Solaris sample configuration files

This section shows example DB2 UDB configurations for Solaris.

## Non-MPP main.cf configuration for Solaris

The following main.cf configuration file reflects DB2 UDB in an ESE multi-partition instance SMP environment. Two database partitions are shown.

```
include "types.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2_clus (
    UserNames = { admin = "cDRpdxPmHpzS." }
    Administrators = { admin }
    CounterInterval = 5
    )
```

```
system sysA (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
     ActionTimeLimit = 0, Action = NONE,
     NotifyThreshold = 0, NotifyTimeLimit = 0 }
    )

system sysB (
    CPUUsageMonitoring = { Enabled = 0, ActionThreshold = 0,
     ActionTimeLimit = 0, Action = NONE,
     NotifyThreshold = 0, NotifyTimeLimit = 0 }
    )

group db2_grp1 (
    SystemList = { sysA = 0, sysB = 1 }
    AutoStartList = { sysA }
    )

    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = DWCNTRL
        NodeNumber = 0
        )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 0
        NodeNumber = 1
        )

    DiskGroup db2dg1 (
        DiskGroup = db2dg1
        )

    IP db2ip1 (
        Device = bge0
        Address = "192.2.40.21"
        )

    Mount db2mnt1 (
        MountPoint = "/db2_mnt/db2inst1"
        BlockDevice = "/dev/vx/dsk/db2dg1/db2dg1home"
        FSType = vxfs
        MountOpt = rw
        FsckOpt = "-y"
        )
```

```
    NIC db2nic1 (
        Device = bge0
        NetworkType = ether
        )

    Volume db2vol1 (
        Volume = db2dg1home
        DiskGroup = db2dg1
        )

    db2ip1 requires db2nic1
    db2mnt1 requires db2vol1
    db2udb1 requires db2ip1
    db2udb1 requires db2mnt1
    db2vol1 requires db2dg1
    db2udb2 requires db2ip1
    db2udb2 requires db2mnt1
```

## MPP main.cf configuration for Solaris

The following configuration file reflects DB2 UDB in an ESE multi-partition instance MPP environment. Four database partitions are shown. One partition is configured on each cluster node. Each database service group depends on the same CVM service group, which manages the shared storage in the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2_mpp (
    CounterInterval = 5
    )

system sysA (
    )

system sysB (
    )

system sysC (
    )

system sysD (
    )

group cvm (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { sysA, sysB, sysC, sysD }
    )
```

```
    CFSMount db2cfsmnt (
        MountPoint = "/db2_mnt/db2inst1"
        BlockDevice = "/dev/vx/dsk/cdb2dg1/cdb2dg1home"
        Primary = sysD
        )

    CFSfsckd vxfsckd (
        )

    CVMCluster cvm_clus (
        Critical = 0
        CVMClustName = db2_mpp
        CVMNodeId = { sysA = 0, sysB = 1, sysC = 2,
           sysD = 3 }
        CVMTransport = gab
        CVMTimeout = 200
        )

    CVMVolDg db2dg (
        CVMDiskGroup = cdb2dg1
        CVMVolume = { cdb2dg1home }
        CVMActivation = sw
        )

    db2cfsmnt requires db2dg
    db2cfsmnt requires vxfsckd
    db2dg requires cvm_clus
    vxfsckd requires qlogckd

    // resource dependency tree
    //
    //          group cvm
    //          {
    //          CFSMount db2cfsmnt
    //              {
    //              CVMVolDg db2dg
    //                  {
    //                  CVMCluster cvm_clus
    //                  }
    //              CFSfsckd vxfsckd
    //                  {
    //                  CFSQlogckd qlogckd
    //                  }
    //              }
    //          }

group db2mpp_grp0 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysA }
    )
```

```
    Db2udb db2udb0 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = SAMPLE
        )

    IP Db2_IP0 (
        Device = bge0
        Address = "10.118.2.144"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC0 (
        Device = bge0
        NetworkHosts = { "10.118.11.90" }
        )

    requires group cvm online local firm
    Db2_IP0 requires Db2_NIC0
    db2udb0 requires Db2_IP0


    // resource dependency tree
    //
    //          group db2mpp_grp0
    //          {
    //          Db2udb db2udb0
    //              {
    //              IP Db2_IP0
    //                  {
    //                  NIC Db2_NIC0
    //                  }
    //              }
    //          }


group db2mpp_grp1 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysB }
    )


    Db2udb db2udb1 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = TEST1
        NodeNumber = 1
        )

    IP Db2_IP1 (
```

```
        Device = bge0
        Address = "10.118.2.145"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC1 (
        Device = bge0
        NetworkHosts = { "10.118.11.90" }
        )

    requires group cvm online local firm
    Db2_IP1 requires Db2_NIC1
    db2udb1 requires Db2_IP1


    // resource dependency tree
    //
    //          group db2mpp_grp1
    //          {
    //          Db2udb db2udb1
    //              {
    //              IP Db2_IP1
    //                  {
    //                  NIC Db2_NIC1
    //                  }
    //              }
    //          }


group db2mpp_grp2 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysC }
    )

    Db2udb db2udb2 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        IndepthMonitor = 1
        DatabaseName = TEST2
        NodeNumber = 2
        )

    IP Db2_IP2 (
        Device = bge0
        Address = "10.118.2.146"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC2 (
        Device = bge0
        NetworkHosts = { "10.118.11.90" }
        )
```

```
    requires group cvm online local firm
    Db2_IP2 requires Db2_NIC2
    db2udb2 requires Db2_IP2


    // resource dependency tree
    //
    //          group db2mpp_grp2
    //          {
    //          Db2udb db2udb2
    //              {
    //              IP Db2_IP2
    //                  {
    //                  NIC Db2_NIC2
    //                  }
    //              }
    //          }

group db2mpp_grp3 (
    SystemList = { sysA = 0, sysB = 1, sysC = 2, sysD = 3 }
    AutoStartList = { sysD }
    )

    Db2udb db2udb3 (
        DB2InstOwner = db2inst1
        DB2InstHome = "/db2_mnt/db2inst1"
        NodeNumber = 3
        )

    IP Db2_IP3 (
        Device = bge0
        Address = "10.118.2.147"
        NetMask = "255.255.248.0"
        )

    NIC Db2_NIC3 (
        Device = bge0
        NetworkHosts = { "10.118.11.90" }
        )

    requires group cvm online local firm
    Db2_IP3 requires Db2_NIC3
    db2udb3 requires Db2_IP3

    // resource dependency tree
    //
    //          group db2mpp_grp3
    //          {
    //          Db2udb db2udb3
    //              {
    //              IP Db2_IP3
```

```
    //                  {
    //                  NIC Db2_NIC3
    //                  }
    //              }
    //          }
```

## DB2 instance running in a Solaris zone

This sample DB2 instance running on Solaris gives the following configuration, which reflects a DB2 UDB instance running in a Solaris 10 zone environment.

```
include "types.cf"
include "/etc/VRTSagents/ha/conf/Db2udb/Db2udbTypes.cf"

cluster db2zone (
    UserNames = { "z_zoneres@vcs_lzs@sysA.engba.veritas.com" = Gn,
        "z_z1@vcs_lzs@sysA.engba.veritas.com" = aH }
    ClusterAddress = "10.178.6.32"
    SecureClus = 1
    CredRenewFrequency = 0
    CounterInterval = 5
    )

system sysA (
    )

system sysB (
    )

group Db2grp (
    SystemList = { sysA = 0, sysB = 1 }
    AutoStartList = { sysA, sysB }
    Administrators = { "z_zoneres@vcs_lzs@sysA.engba.veritas.com" }
    )

    DiskGroup z-dg (
        DiskGroup = db2dg1
        )

    IP ipres (
        Device = bge0
        Address = "10.178.6.28"
        ContainerName = "zone1"
        )

    Mount z-mnt (
        MountPoint = "/zones/db2data"
        BlockDevice = "/dev/vx/dsk/db2dg1/db2dg1data"
        FSType = vxfs
        FsckOpt = "-y"
        )
```

```
NIC z-nic (
    Device = bge0
    NetworkType = ether
    NetworkHosts = { "10.178.2.4" }
    )

Volume z-vol (
    Volume = db2dg1data
    DiskGroup = db2dg1
    )

Zone zoneres (
    ZoneName = zone1
    )

Db2udb db2udb1 (
    ContainerName = "zone1"
    DB2InstOwner = "db2inst1"
    DB2InstHome = "/db2inst1"
    )

ipres requires z-nic
ipres requires zoneres
z-mnt requires z-vol
z-vol requires z-dg
zoneres requires z-nic
zoneres requires z-mnt
db2udb1 requires zoneres

// resource dependency tree
//
//          group Db2grp
//          {
//          IP ipres
//              {
//              NIC z-nic
//              }
//          Mount z-mnt
//              {
//              Volume z-vol
//                  {
//                  DiskGroup z-dg
//                  }
//              }
//          Zone zoneres
//              {
//              NIC z-nic
//              }
//          }
```

# Index

## T



## U



## V



## W