

# **VERITAS Cluster Server Enterprise Agent 4.1 for DB2**

## **Installation and Configuration Guide**

**Solaris**

N15372F

March 2005

---

## **Disclaimer**

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

## **VERITAS Legal Notice**

Copyright © 1998-2005 VERITAS Software Corporation. All rights reserved. VERITAS and the VERITAS Logo are trademarks or registered trademarks of VERITAS Software Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

VERITAS Software Corporation  
350 Ellis Street  
Mountain View, CA 94043  
USA  
Phone 650-527-8000 Fax 650-527-2908  
[www.veritas.com](http://www.veritas.com)



# Contents

---

<b>Preface</b> .....	<b>vii</b>
Technical Support .....	vii
Conventions .....	viii
<b>Chapter 1. Introduction</b> .....	<b>1</b>
Version Numbers and Operating Systems .....	1
About the DB2 UDB Agent .....	1
New Features and Updates .....	2
Monitoring DB2 Instances Running in Solaris 10 Zones .....	2
Zone Agent .....	2
New Attributes .....	2
Operations of the DB2 UDB Agent .....	3
Online Operation .....	3
Offline Operation .....	3
Monitor Operation .....	4
ESE Single-Partition Instance (Earlier than Version 8) .....	4
ESE Multi-Partition Instance, SMP Hardware Configuration (Earlier than Version 8) .....	4
ESE Configurations, Version 8.0 or Later .....	5
Clean Operation .....	6
Info Entry Point Operation .....	7
Action Entry Point Operation .....	8



---

<b>Chapter 2. Preparation for Installing DB2 UDB in a VCS Cluster</b>	<b>9</b>
Prerequisites for Installing DB2 UDB, Non-MPP Versions	11
Prerequisites for Installing DB2 UDB, MPP Version	11
Creating File Systems for DB2 Instances	12
Creating the File System for the DB2 Non-MPP Instances	12
Creating the Shared Cluster File System for the DB2 MPP Instances	13
Defining DB2 User and Group Accounts	14
Creating Groups	14
Adding User Accounts	15
Installing DB2 UDB in VCS Environment	15
Setting Shared Memory Parameters	15
Installing the Binaries	15
Install the DB2 License	15
Installing the Instances	16
Setting Up the DB2 UDB Configuration	17
Checking /etc/services	17
Creating \$DB2InstHome/.rhosts	18
Modifying the \$DB2InstHome/sqlib/db2nodes.cfg File	19
Non-MPP Versions	19
MPP Versions	19
Confirming the Setup of DB2 MPP Installation	20
<b>Chapter 3. Installing the DB2 Agent</b>	<b>21</b>
Installing the DB2 UDB Agent Software	21
<b>Chapter 4. Configuring the DB2 Agent</b>	<b>23</b>
Configuring the Agent	23
Db2udb Resource Type Attributes	24
DB2 UDB Type Definition File, Db2udbTypes.cf	26
DB2 UDB Service Group for Non-MPP Configuration	27



---

DB2 UDB Service Groups for MPP Configuration .....	28
DB2 UDB Instance Configured in Solaris Zones .....	29
Zone Root on Local Disk .....	29
Configuring the DB2 UDB Agent Using Cluster Manager .....	30
Importing the Db2udbTypes.cf File .....	30
Adding Service Group for DB2 UDB MPP Using Cluster Manager .....	31
Adding Service Group for DB2 UDB Non-MPP Using Cluster Manager .....	34
Configuring the DB2 UDB Agent by Editing the main.cf File .....	37
Preparation for Editing the main.cf File .....	37
Configuring the Agent to Use the DB2 UDB MPP Configuration .....	37
Configuring the Agent to Use the DB2 UDB, Non-MPP Configurations .....	38
Verifying the Configuration .....	39
Modifying the Agent Configuration .....	41
Enabling In-Depth Monitoring of DB2 UDB Instance .....	41
Enabling In-Depth Monitoring from the Command Line .....	41
Disabling In-Depth Monitoring .....	42
<b>Chapter 5. Disabling and Removing the Agent .....</b>	<b>43</b>
Disabling the Agent .....	43
Removing the Agent .....	44
<b>Appendix A. Sample Configuration Files .....</b>	<b>45</b>
Non-MPP Configuration: DB2 UDB ESE Single-Partition Instance .....	45
Non-MPP Configuration: DB2 UDB ESE Multi-Partition Instance .....	49
MPP Configuration: DB2 UDB ESE Multi-Partition Instance .....	51
Sample DB2 Instance Running on a Solaris Zone .....	57
<b>Index .....</b>	<b>59</b>





# Preface

---

This document provides instructions on how to install and configure the VERITAS Cluster Server Agent for DB2 UDB. For information about VCS, refer to the *VERITAS Cluster Server User's Guide*.

## Technical Support

**U.S. and Canada:** Call 1-800-342-0652. You may also contact Technical Support via email at [vsupport-us@veritas.com](mailto:vsupport-us@veritas.com).

**Europe, the Middle East, or Asia:** Visit the Technical Support web site at <http://support.veritas.com> for a list of each country's contact information.

Software updates, TechNotes, product alerts, and hardware compatibility lists, are also available from <http://support.veritas.com>

To learn more about VERITAS and its products, visit <http://support.veritas.com>.



## Conventions

<b>Typeface</b>	<b>Usage</b>
<code>courier</code>	computer output, command references within text
<b><code>courier</code></b> (bold)	user input, keywords in grammar syntax
<i>italic</i>	new terms, titles, emphasis
<i>italic</i>	variables
<b>Symbol</b>	<b>Usage</b>
#	UNIX superuser prompt (for all shells)
\$	Bourne/Korn shell prompt
%	C shell prompt



# Introduction

---

1

Welcome to the VERITAS Cluster Server (VCS) agent, version 4.1, for DB2 UDB. DB2 Universal Database is a 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.

## Version Numbers and Operating Systems

The VCS enterprise agent, version 4.1 for DB2 UDB, supports DB2 Universal Database Enterprise Server Edition (ESE) versions for single and multi-partition instance. The agent supports ESE for single partition instance for versions 7.2, 8.1 and 8.2. For ESE multi-partition instance, it supports the Symmetric Multiprocessing (SMP) hardware configuration for versions 7.2, 8.1, and 8.2 and supports the Massively Parallel Processing (MPP) hardware configuration for version 8.1 and 8.2. DB2 ESE must run on Solaris 2.8 or later. The VCS version must be 4.0 and above. The memory requirements vary for different versions of DB2 being used. Check the *DB2 Universal Database Quick Beginnings Guide* for information about memory requirements.

## About the DB2 UDB Agent

The DB2 UDB enterprise agent monitors DB2 database instances while they are up and running on a given system. If the system fails, the agent detects the failure and takes the DB2 instances offline. VCS conducts failover to another system in the cluster, where the agent brings DB2 instances online.

For ESE single partition instance, the agent brings DB2 UDB database instances online, monitors database processes, and shuts down instances. For ESE multi-partition instance, it brings DB2 UDB database partitions online, monitors the database processes at the partition level, and shuts down database partitions.



## New Features and Updates

This section describes the new features and updates introduced in this version of the agent.

### Monitoring DB2 Instances Running in Solaris 10 Zones

Solaris 10 provides a means of virtualizing operating system services, allowing one or more processes to run in isolation from other activity on the system. Such a "sandbox" is called a *non-global zone*. Each non-global zone can provide a rich and customized set of services. There is also a *global zone* and processes running in this zone have the same set of privileges available on a Solaris system today.

VCS provides high availability to applications running in non-global zones by extending the failover capability to zones. VCS is installed in a global zone and all VCS agents and engine components run in the global zone. For applications running within non-global zones, agents run entry points inside the zones. If a non-global zone configured under VCS control faults, VCS fails over the entire service group containing the zone. The VCS enterprise agent for DB2 UDB is zone-aware and can monitor DB2 UDB instances running in non-global zones. Note that the MPP configuration cannot be supported in non-global zones.

### Zone Agent

The new Zone agent monitors zones, brings zones online, and takes them offline. The agent creates a user account with group administrative privileges to enable inter-zone communication if that user account doesn't exist. It renews the VxSS authentication certificate automatically before the certificate expires.

### New Attributes

Two new attributes, `ContainerType` and `ContainerName`, define the zone where the application runs.

- ◆ `ContainerType`—A static resource type attribute set to `Zone` to indicate this application is zone aware. *Do not modify this attribute.*
- ◆ `ContainerName`—A resource attribute, which you set to the name of the zone.

When you set the `ContainerName`, VCS runs the entry points for the resource in the specified zone.



## Operations of the DB2 UDB Agent

The online, offline, monitor, and clean operations performed by the DB2 agent vary depending on whether the version of DB2 UDB is ESE single-partition instance or ESE multi-partition instance and the ESE version. These operations are described in the following sections.

### Online Operation

For ESE versions earlier than version 8, the agent uses the `db2start` program to start a DB2 instance or database partition.

- ◆ The command for ESE single-partition instance is:

```
su $DB2InstOwner -c "$InstHome/sqlllib/adm/db2start"
```

- ◆ The command for ESE multi-partition instance in SMP hardware configuration is:

```
su $DB2InstOwner -c "$InstHome/sqlllib/adm/db2start $nodenum
$NodeNumber"
```

For ESE version 8, the agent uses `db2gcf` program to start either the ESE single-partition instance configuration or the ESE multi-partition instance in SMP or MPP hardware configuration. The command is:

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

### Offline Operation

For ESE versions earlier than version 8, the agent uses the `db2stop` program to stop a DB2 instance or database partition.

- ◆ The command for ESE single-partition instance is:

```
su $DB2InstOwner -c "$InstHome/sqlllib/adm/db2stop force"
```

- ◆ The command for ESE multi-partition instance in SMP hardware configuration is:

```
su $DB2InstOwner -c "$InstHome/sqlllib/adm/db2stop force $nodenum
$NodeNumber"
```

For ESE version 8, the agent uses the `db2gcf` program to stop either the ESE single-partition instance configuration or the ESE multi-partition instance in SMP or MPP hardware configuration. The command is:

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



## Monitor Operation

The commands used by the agent to monitor the DB2 instances vary depending on the DB2 version and hardware configuration.

### ESE Single-Partition Instance (Earlier than Version 8)

For versions of the ESE single-partition instance earlier than version 8, the agent executes the command:

```
db2nps 0
```

to check the processes owned by the instance owner. If the output of the command contains DB2 processes, the monitor exits with a return value of 110, indicating success. If this command doesn't exist, then the monitor tries to find the `db2sysc` process owned by the instance in the process table. If the monitor finds the process in the table, it exits with a return value of 110 and the DB2 instance continues online. Otherwise, the DB2 instance is taken offline and failed over to the next system in the service group's `SystemList` attribute (if the `RestartLimit` and `ToleranceLimit` are set to 0).

If in-depth monitoring is enabled (that is, the `InDepthMonitor` attribute is set to 1 and the `DatabaseName` attribute is not NULL), the monitor performs a query to the database indicated in the `DatabaseName` attribute. If the query succeeds, the instance continues online (monitor exit code for success is 110).

If the database query has any errors or problems, it checks the value of the `WarnOnlyIfDBQueryFailed` attribute of the `Db2udb` agent. If `WarnOnlyIfDBQueryFailed` is set to 1 (the default), the agent logs an error message containing the actual SQL error in the engine log (no more than once a day to prevent overflowing the engine log) and returns exit code 110. If `WarnOnlyIfDBQueryFailed` is set to 0, the agent logs an error message and returns exit code 100, in which case the instance restarts or fails over.

### ESE Multi-Partition Instance, SMP Hardware Configuration (Earlier than Version 8)

For versions of the ESE multi-partition instance earlier than version 8, the agent executes the command:

```
db2nps $NodeNumber
```

to check the processes owned by the instance owner for a particular database partition or node number. If the output of the command contains DB2 processes, the monitor returns an exit code of 110. Otherwise, the database partition owned by the instance is taken offline and failed over or restarted, depending on the values of the `RestartLimit` and `ToleranceLimit` attributes. If both attributes are set to 0, the instance fails over to the next system in the group's `SystemList` attribute.

If in-depth monitoring is enabled (that is, the `IndepthMonitor` attribute is set to 1 and the `DatabaseName` attribute is not NULL), the monitor performs a query to the database indicated in the `DatabaseName` attribute. The database is to be created from this partition. The database partition must be able to connect to this local database it monitors without error. The default `MonitorTimeout` value is 240 seconds. If the database connection you use is very slow, you can adjust the `MonitorTimeout` attribute to a higher value. If the query succeeds, `monitor` returns exit code 110 (`monitor` returns exit 110 for online and 100 for offline).

If the database query has any errors or problems, it checks the value of the `Db2udb` agent's `WarnOnlyIfDBQueryFailed` attribute. If `WarnOnlyIfDBQueryFailed` is set to 1 (the default), the agent logs an error message containing the actual SQL error in the engine log (no more than once a day to prevent overflowing the engine log) and returns exit code 110. If `WarnOnlyIfDBQueryFailed` is set to 0, the agent logs an error message and returns exit code 100.

## ESE Configurations, Version 8.0 or Later

For all DB2 ESE configurations, version 8.0 or later, for a single or multi-partition instance in SMP or MPP configuration, the agent executes the command:

```
db2gcf -s -i $DB2InstOwner -p $nodenum
```

to check the status of the database partition or node number. If the exit status of the `db2gcf` command is 0, the monitor returns 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`. The command `db2gcf` is only available.

If in-depth monitoring is enabled (that is, the `IndepthMonitor` attribute is set to 1 and the `DatabaseName` attribute is not NULL), the monitor performs a query to the database indicated in the `DatabaseName` attribute. The database is to be created from this partition.

If the database query has any errors or problems, it checks the value of the `Db2udb` agent's `WarnOnlyIfDBQueryFailed` attribute. If `WarnOnlyIfDBQueryFailed` is set to 1 (the default), the agent logs an error message containing the actual SQL error in the engine log (no more than once a day to prevent overflowing the engine log) and returns exit code 110. If `WarnOnlyIfDBQueryFailed` is set to 0, the agent logs an error message and returns exit code 100. The default `MonitorTimeout` value is 240 seconds. If the database connection you use is very slow, you can adjust the `MonitorTimeout` attribute to a higher value.



## Clean Operation

For ESE versions earlier than version 8, the agent uses the `db2nkill` program with instance owner's ID to kill the instance processes owned by the `$DB2InstOwner` for the instance or database partition.

- ◆ For ESE single-partition instance, the agent uses the following command to kill a DB2 instance:

```
su $DB2InstOwner -c "$InstHome/sqlllib/bin/db2nkill 0"
```

- ◆ For ESE multi-partition instance, the agent uses the following command to kill a DB2 database partition in the SMP configuration:

```
su $DB2InstOwner -c "$InstHome/sqlllib/bin/db2nkill nodenum  
$NodeNumber"
```

For ESE version 8, the agent uses the `db2gcf` program to kill a DB2 database partition with either the ESE single-partition instance configuration or the ESE multi-partition instance in SMP or MPP configuration. The command is:

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



## Info Entry Point Operation

The DB2 agent supports the `Info` entry point, which provides static and dynamic information about the database partition and its critical processes. In the example below, the entry point retrieves the database information shown by executing the following commands:

1. Make the configuration writable:

```
# haconf -makerw
```

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

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

3. Show the requested `ResourceInfo` value. In the following example output, the name value pairs for processes monitored by the agent for the DB2 resource.

```
# hares -value db2resource ResourceInfo
```

```
State Valid
```

```
Msg
```

```
PARTITION: 0
```

```
UID=db2inst1 PID=21763 PPID=21750 PNAME=db2sysc
UID=db2inst1 PID=21786 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21781 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21789 PPID=21763 PNAME=db2sysc
UID=root PID=21777 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21788 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21787 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21796 PPID=21763 PNAME=db2sysc
UID=root PID=21774 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21794 PPID=21763 PNAME=db2sysc
UID=root PID=21778 PPID=21763 PNAME=db2sysc
UID=root PID=21771 PPID=21763 PNAME=db2sysc
UID=db2inst1 PID=21795 PPID=21778 PNAME=db2sysc
```

For more information about the `info` entry point, refer to the *VERITAS Cluster Server User's Guide* and the *VERITAS Cluster Server Agent Developer's Guide*.



## Action Entry Point Operation

The DB2 agent supports the `Action` entry point, 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 the following predefined actions:

Predefined Action Token	Description
VRTS_GetInstanceName	Retrieves the DB2 instance name of the configured Db2udb resource.
VRTS_GetRunningServices	Retrieves the list of processes monitored by the agent for the Db2udb resource.

For example:

```
# hares -action db2udb0 VRTS_GetInstanceName
VCS NOTICE V-16-13323 Resource (db2udb0): action
(VRTS_GetInstanceName) completed successfully. Output is:
db2inst1

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

UID=db2inst1    PID=21763  PPID=21750  PNAME=db2sysc
UID=db2inst1    PID=21786  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21781  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21789  PPID=21763  PNAME=db2sysc
UID=root        PID=21777  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21788  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21787  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21796  PPID=21763  PNAME=db2sysc
UID=root        PID=21774  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21794  PPID=21763  PNAME=db2sysc
UID=root        PID=21778  PPID=21763  PNAME=db2sysc
UID=root        PID=21771  PPID=21763  PNAME=db2sysc
UID=db2inst1    PID=21795  PPID=21778  PNAME=db2sysc
```

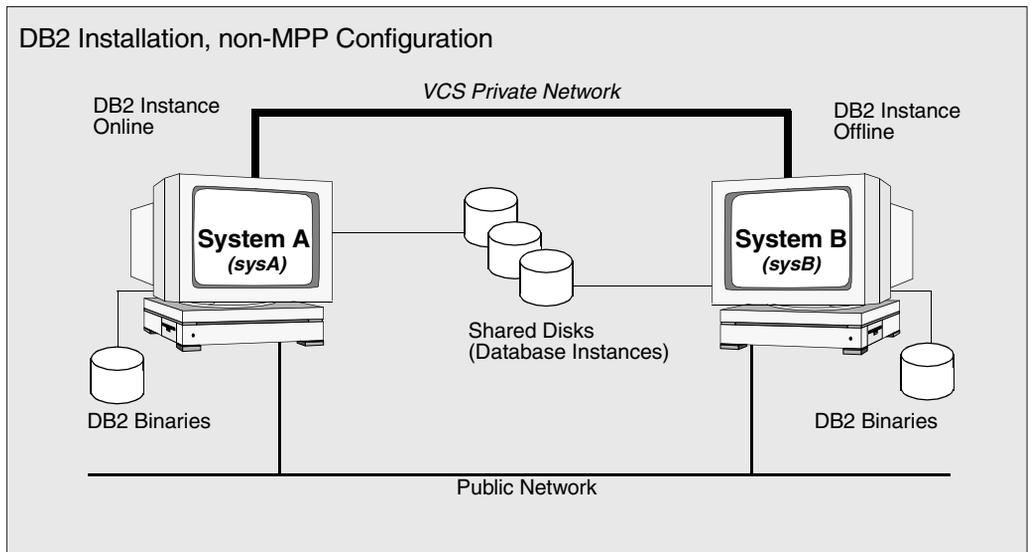
You can also add custom actions for the agent. Refer to the *VERITAS Cluster Server Agent Developer's Guide* for information on defining custom action tokens.



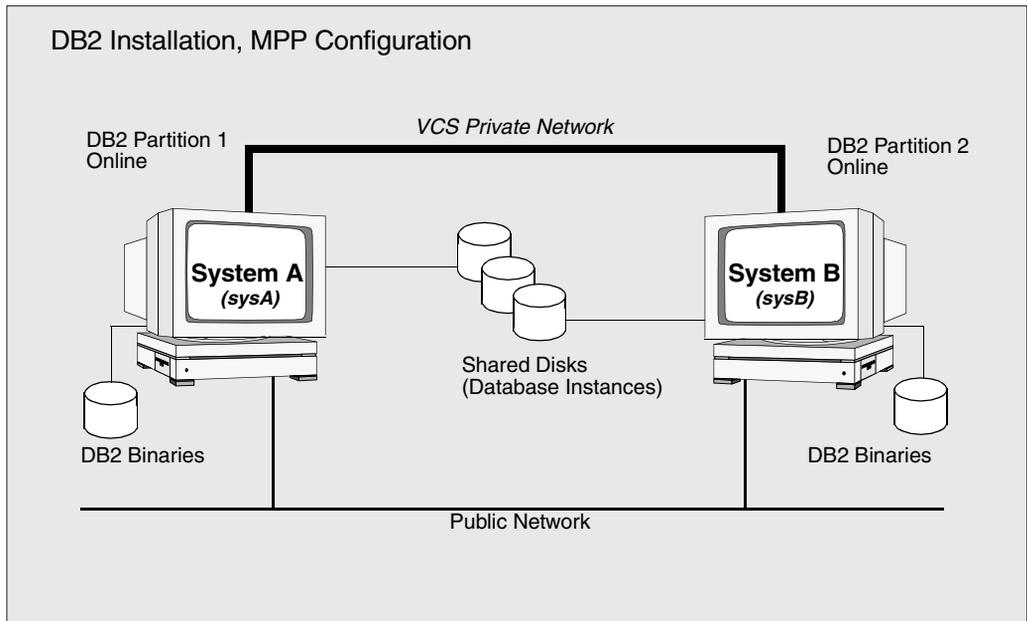
# Preparation for Installing DB2 UDB in a VCS Cluster

## 2

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, while the other system is a failover system.



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



## Prerequisites for Installing DB2 UDB, Non-MPP Versions

- ✓ Verify 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. The DB2 UDB system binaries are to be installed locally and the DB2 UDB database instances are to be installed on shared storage.
- ✓ Install and configure VCS version 4.0 or later. Refer to the *VERITAS Cluster Server Installation Guide* for instructions on running either the VERITAS installer or the `installvcs` utility.
- ✓ Before installing DB2 UDB, define DB2 UDB user and group accounts. See “[Defining DB2 User and Group Accounts](#)” on page 14. Refer also to the relevant *DB2 Universal Database Quick Beginnings* guide.

## Prerequisites for Installing DB2 UDB, MPP Version

- ✓ Verify 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. The DB2 UDB system binaries are to be installed on the local file systems on each system and the DB2 UDB database instances are to be installed on shared storage.
- ✓ The MPP configuration requires the Storage Foundation Cluster File System software. This software includes the cluster file system components required by the VCS enterprise agent 4.0 or later for DB2 UDB, and includes VERITAS Cluster Server (VCS), VERITAS Volume Manager with cluster functionality enabled (CVM), and VERITAS File System with cluster functionality enabled (CFS).  
  
Refer to the *VERITAS Storage Foundation Cluster File System Installation and Administration Guide* for detailed information on these products and instructions on running either the VERITAS installer or the `installsfdfs` utility.
- ✓ Before installing DB2 UDB, define DB2 UDB user and group accounts. See “[Defining DB2 User and Group Accounts](#)” on page 14. Refer also to the relevant *DB2 Universal Database Quick Beginnings* 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 must first create a disk group on the physically shared disk, and create a volume of sufficient size within the disk group.

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:

```
# vxvg init db2db_dg c4t0d0s2
```

Deport and import the disk group:

```
# vxvg deport db2db_dg
```

```
# vxvg 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 -o largefiles /dev/vx/rdisk/db2db_dg/db2db_vol
```

4. Create the mount point directory and mount the file system:

```
mkdir /db2_mnt/db2inst1
```

```
mount -F vxfs /dev/vx/dsk/db2db_dg/dbq2db_vol  
/db2_mnt/db2inst1
```



## Creating the Shared Cluster File System for the DB2 MPP Instances

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

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, `c5t0d0s2`. For example:

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

Deport and import the disk group:

```
# vxdg deport db2db_dg
```

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

4. Create a volume of seven GB using the `vxassist` command:

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

5. Create the file system:

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

6. Create the mount point directory and mount the file system:

```
mkdir /db2_mnt/db2inst1
```

```
mount -F vxfs -o cluster /dev/vx/dsk/db2db_dg/dbq2db_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.
- ✓ All DB2 user accounts must exist on the local systems. The use of NIS or NIS+ for users is not recommended because these services are not highly available. If their service is interrupted, VCS may not be able to work correctly.

### Creating Groups

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

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

In the following examples that show creating user accounts, the `-g` option specifies the group, `-u` specifies the user ID, `-d` the home directory, `-m` specifies that the home directory is to be created if it doesn't exist, `-s` is the user's login shell, and the final expression is the user's login.

Create the user accounts on *each* node in the cluster.

1. The first example shows creating the user, `db2inst1`, the DB2 UDB instance owner. The mount point, `/db2_mnt/db2inst1`, is used for a file system that hosts the DB2 UDB instance home directory on shared storage, accessible to each node. For example:

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

2. The next 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 VCS Environment

For installing DB2 UDB version 7.x or 8.1 ESE in a VCS environment, we recommend you follow the installation procedure documented in the relevant *DB2 Universal Database Quick Beginnings* 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

For details on setting the shared memory parameters in the `/etc/system` file on each node, refer to the relevant *DB2 Universal Database Quick Beginnings* guide.

### 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 the `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 the `db2setup` tool.

- ✓ When using `db2setup`, do not select the option to "Auto start DB2 instance at `system boot`" in the DB2 Instance Properties window (if this option exists for your DB2 version). VCS needs to bring up the resources for the DB2 instances in a specific order before bringing the instance itself online.
- ✓ The instance's home directory is a mount point on the shared storage.

## Setting Up the DB2 UDB Configuration

Use the following procedures to configure DB2 UDB ESE multi-partition instance (non-MPP) and DB2 UDB ESE multi-partition instance (MPP) in a VCS environment.

### Checking /etc/services

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

- ✓ Make sure each partition has a port number assigned. The number of ports reserved depends on the number of partitions.
- ✓ Make sure the ports are not used by any other services. 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 a DB2 UDB instance with two partitions, one port for each partition:

```
# more /etc/services
DB2_db2inst1      70000/tcp # Connection port for DB2 instance
db2inst1
DB2_db2inst1_END 70001/tcp
DB2_db2inst1a    60003/tcp # Interrupt port for DB2 instance
db2inst1
DB2_db2inst10    50002/tcp # Connection port for DB2 instance
db2inst1
DB2_db2inst10    50003/tcp # Interrupt port for DB2 instance
db2inst1
```

The file shows the ports assigned when DB2 UDB creates a new instance. The first two lines in the example show a range of ports, 70000 through 70001, for two database partitions. The next three lines assign additional ports for use by the database instance. Inspect the file and verify there are no duplicate port numbers.

The following is an example for a DB2UDB instance with four partitions, a port for each partition.

```
# more /etc/services
DB2_db2inst1      60000/tcp #
DB2_db2inst1_1    60001/tcp #
DB2_db2inst1_2    60002/tcp #
DB2_db2inst1_END  60003/tcp #
db2c_db2inst1     50000/tcp #
```

The four instances are assigned 60000 through 60003.



## 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 being prompted for 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
```

Using the command, `rsh system_name`, test that you can remotely log in with the DB2 instance (for example, `db2inst1`) account from one system in the cluster to another without being prompted for a password. Test this from each system in the cluster to all other systems.



## Modifying the \$DB2InstHome/sqlllib/db2nodes.cfg File

The \$DB2InstHome/sqlllib/db2nodes.cfg file is used by DB2 during failover from one node to another.

### Non-MPP Versions

For each DB2 UDB ESE multi-partition instance (non-MPP) database partitions, modify the file \$DB2InstHome/sqlllib/db2nodes.cfg such that you create an entry for each database partition, assigning 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 listed in the file */etc/hosts*. Make sure that the virtual IP address is up and running at this time.

### MPP Versions

For MPP versions, modify the file \$DB2InstHome/sqlllib/db2nodes.cfg with the hostname that you want each database partition to start on. The db2nodes.cfg file is automatically changed and updated by DB2 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 Installation

On the host where the shared file system is mounted, check whether you can start and stop each instance, and thus verify that DB2 is properly installed.

1. Log in as the instance owner:

```
# su - db2inst1
```

2. Attempt to start the instance:

```
$ db2start
```

The application should start on the nodes specified in the `db2nodes.cfg` file.

3. Assuming the previous command is successful, stop the instance:

```
$ db2stop
```

4. If the application does not start successfully or stop correctly on each node, check for configuration errors. Review your DB2 UDB documentation, such as the *DB2 Universal Database Quick Beginnings Guide*.

5. Create a database.

```
$ db2 create database dbname
```

6. List the database directory

```
$ db2 list database directory
```



# Installing the DB2 Agent

3

Use the following procedure to install the DB2 UDB agent software.

## Installing the DB2 UDB Agent Software

▼ To install the agent

1. Create a temporary directory for installation:

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

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

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

3. 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
```

4. Extract the compressed file from the tar file:

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

5. Install the package:

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



▼ **To install the Japanese language pack**

1. After installing the agent, insert the “Language” disc into a drive connected to the system. 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
```

## Configuring the DB2 Agent

---

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

- ◆ By using VCS Cluster Manager (the Java Console) to edit a service group template for the DB2 UDB agent. See “[Configuring the DB2 UDB Agent Using Cluster Manager](#)” on page 30.
- ◆ By using VCS commands. Refer to the *VERITAS Cluster Server User’s Guide* for information about configuring VCS from the command line.
- ◆ By editing the `main.cf` file directly, using the types configuration file and referring to the sample `main.cf` file supplied with the DB2 UDB agent (see “[Configuring the DB2 UDB Agent by Editing the main.cf File](#)” on page 37). This method requires that VCS be stopped and restarted before the configuration takes effect.

### Configuring the Agent

Configuring the DB2 UDB agent involves assigning values to the DB2 UDB resource type attributes, which are described in the following table for your review and reference. The resource type definition file, `Db2udbTypes.cf`, is also shown for reference. The sample `main.cf` configuration files are shown in [Appendix A](#).



## Db2udb Resource Type Attributes

The required and optional attributes are described below. The temporary attribute, `LastWarningDay` (not described below), is internally used by the agent to ensure that the same error messages are not repeatedly logged.

Attributes, Required	Type and Dimension	Definition
DB2InstOwner	string-scalar	User ID of Instance Owner that starts a DB2 UDB instance. Each instance requires a unique user ID.
DB2InstHome	string-scalar	Path to DB2 UDB instance home directory that contains configuration files for the DB2 instance.
Attributes, Optional	Type and Dimension	Definition
ContainerType	string-scalar	A static resource type attribute set to <code>Zone</code> to indicate this application is zone aware. <i>Do not modify this attribute.</i>
IndepthMonitor	integer-scalar	Flag indicating if in-depth monitor is enabled (1) or disabled (0). The default is 0.
DatabaseName	string-scalar	Name of the database for in-depth monitoring; required if in-depth monitor is enabled ( <code>IndepthMonitor = 1</code> ).
NodeNumber	integer-scalar	Node number or partition number of the database. Used when monitoring a specific database partition in ESE multi-partition instance environment. Default value is 0 for ESE single-partition instance and multi-partition instance configurations.
Encoding	string-scalar	Specifies operating system encoding corresponding to DB2 UDB encoding for display of DB2 UDB output. For example, if the environment variable <code>LANG</code> is set to <code>"ja,"</code> then <code>"eucJP"</code> is the Solaris value for <code>Encoding</code> . Refer to DB2 UDB and Solaris documentation for respective encoding values. The default is <code>""</code> .



<b>Attributes, Required</b>	<b>Type and Dimension</b>	<b>Definition</b>
WarnOnlyIfDBQueryFailed	boolean-scalar	Flag indicating if the DB2 resources should be faulted when the in-depth monitor fails. In-depth monitoring consists of a database query. <ul style="list-style-type: none"><li>◆ If the attribute is set to 0, the agent faults the DB2 resource if query fails.</li><li>◆ If the attribute is set to 1, the agent issues a warning message about the query failure; the resource remains online.</li></ul>
ContainerName	string-scalar	Specifies the name of the zone where the DB2 UDB instance to be monitored is running.



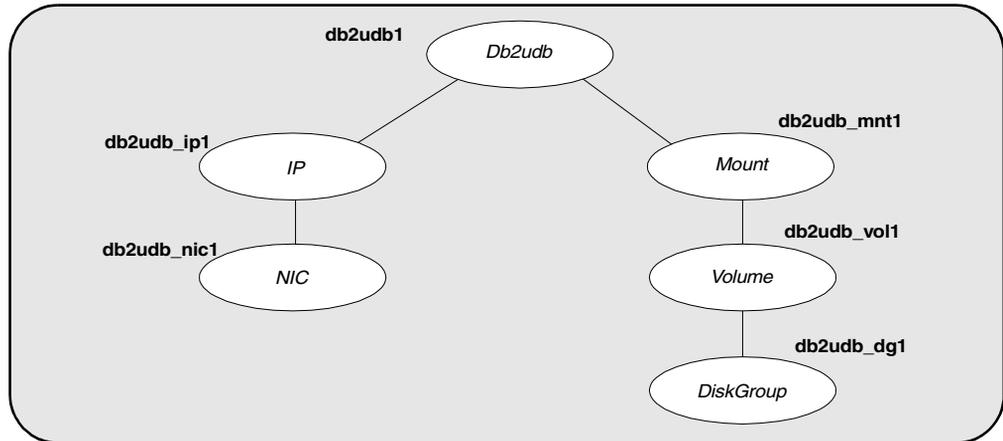
## DB2 UDB Type Definition File, Db2udbTypes.cf

```
type Db2udb (
  static str ContainerType = Zone
  static keylist SupportedActions = {VRTS_GetInstanceName,
    VRTS_GetRunningServices }
  static int CleanTimeout = 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, Encoding,
    WarnOnlyIfDBQueryFailed, LastWarningDay }
  str DB2InstOwner
  str DB2InstHome
  int IndepthMonitor
  str DatabaseName
  int NodeNumber
  str Encoding
  boolean WarnOnlyIfDBQueryFailed = 1
  temp str LastWarningDay
  str ContainerName
)
```



## DB2 UDB Service Group for Non-MPP Configuration

The figure below illustrates the dependencies among the resources configured for a non-MPP DB2 UDB instance resource group.



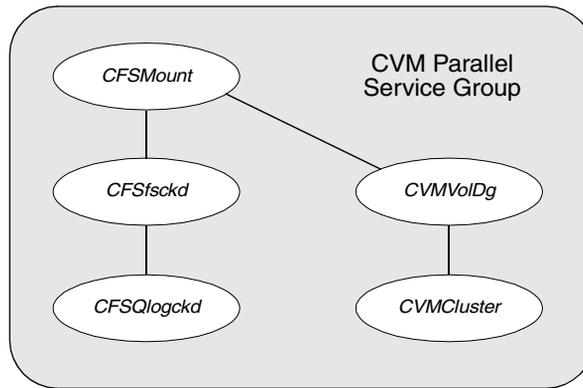
This configuration shows a service resource group for an instance of DB2 UDB. The `db2udb1` resource (the database) requires the IP resource and the Mount resource. The service group IP address for the DB2 UDB server is configured using the IP resource (`db2udb_ip1`) and the NIC resource (`db2udb_nic1`). The mount resource (`db2udb_mnt1`) requires the Volume resource (`db2udb_vol1`) which in turn requires the DiskGroup resource (`db2udb_dg1`). The DB2 UDB instance can be started after each of these resources is available.



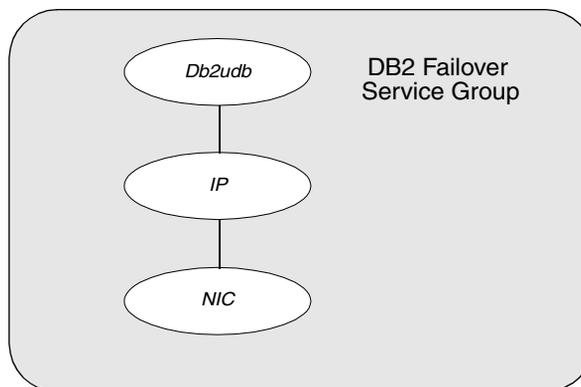
## DB2 UDB Service Groups for MPP Configuration

The DB2 UDB agent uses two service groups to support MPP configuration for DB2 version 8.1 and above.

The first service group is a parallel CVM service group. There is one CVM/Infrastructure group per cluster node. This service group has the CVM resource and the necessary resources for support of CFS. This group also contains all common components needed by DB2, such as the instance's home directory, which is shared on all the cluster nodes.



The second service group is a failover DB2 service group. This service group monitors one database partition for DB2 version 8 with MPP configuration. The failover DB2 service group depends on the parallel CVM service group with online local firm dependency.

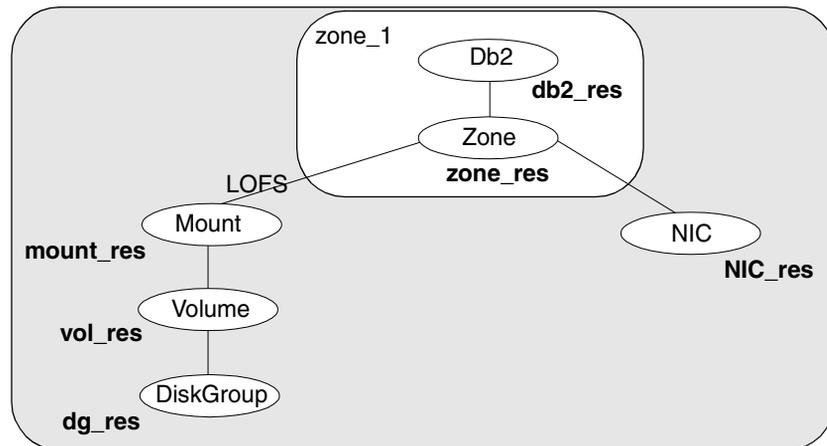


## DB2 UDB Instance Configured in Solaris Zones

The following examples describe a typical service group configured to monitor the state of a DB2 instance configured in a Solaris zone.

### Zone Root on Local Disk

If the root file system of a zone is on the local disk of each node, the file system is mounted when the system is booted. Hence, the service group does not need to have separate DiskGroup and Volume resources for the 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
  
```



## Configuring the DB2 UDB Agent Using Cluster Manager

Templates for the DB2 UDB resource groups were automatically installed when you installed the DB2 UDB enterprise agent. Templates include one for the DB2 UDB MPP configuration and one for the DB2 UDB non-MPP configuration. Using the VCS Cluster Manager (Java Console), you can use a template to configure the DB2 UDB service group, its resources, and their attributes. You can also use the Java Console to dynamically modify the attributes' values as necessary for your configuration.

Refer to the *VERITAS Cluster Server User's Guide* for information on the VCS Java Console.

### Importing the Db2udbTypes.cf File

To use the DB2 UDB templates, import the `Db2udbTypes.cf` file to the VCS engine by using Cluster Manager (Java Console):

1. On one of the systems of the cluster, start Cluster Manager.

```
# haGUI&
```

2. Log into the cluster and wait for Cluster Explorer to launch.
3. In the Cluster Explorer window, click on 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/VRTSvcs/conf/config/Db2udbTypes.cf
```

5. Click Import and wait for the file to be imported.
6. In the Cluster Explorer window, click the Save Configuration icon.

When the DB2 UDB types are imported to the VCS engine, the DB2 UDB agent can be configured.

- If you are using the DB2 UDB MPP configuration, go to [“Adding Service Group for DB2 UDB MPP Using Cluster Manager”](#) on page 31.
- If you are using the DB2 UDB non-MPP configuration, go to [“Adding Service Group for DB2 UDB Non-MPP Using Cluster Manager”](#) on page 34.



## Adding Service Group for DB2 UDB MPP Using Cluster Manager

If you have imported the `Db2udbTypes.cf` file (see [“Importing the Db2udbTypes.cf File”](#) on page 30), you can use the `Db2udb_MPP_Group` template to configure a service group.

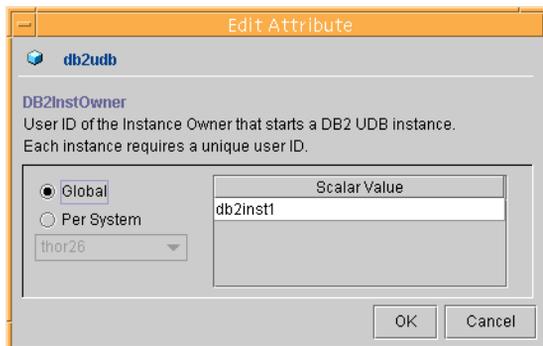
After you log into Cluster Manager, the Status tab should indicate that the CVM service group is Online on each system in the cluster. The CVM service group is automatically configured when you complete the installation of the Storage Foundation Cluster File System (SFCFS) software.

To add the service group for the DB2 UDB MPP database, do the following:

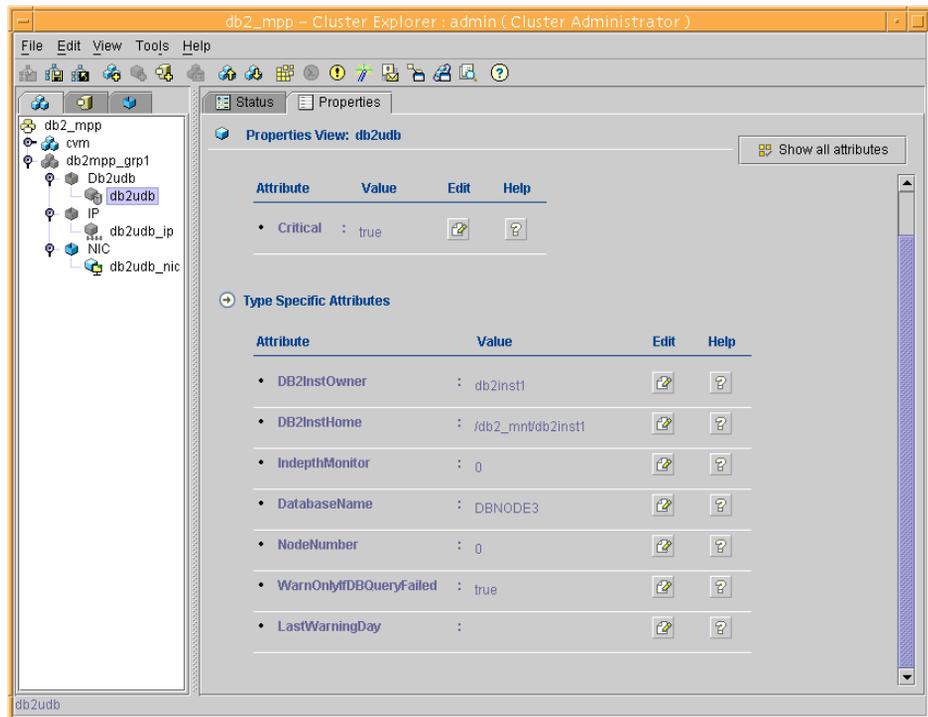
1. In the Cluster Explorer window, click the Add Service Group icon on the toolbar.
2. In the Add Service Group window, do the following:
  - a. Enter the name of the service group you want to create for the service group. For example, enter `db2mpp_grp1`. Do not press Return or Enter.
  - b. From the systems shown in the Available Systems box, double-click those you want in your configuration.
  - c. In the window showing the systems added to the configuration, click the checkbox for system on which you want to automatically start the service group.
  - d. Click the Failover radio button to specify the Service Group Type.
  - e. Click the Templates button.
  - f. In the Select Templates window, select `db2udb_mpp_grp` from the list shown 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.
  - g. Click OK on at the bottom of the Add Service Group window. The group is added. On the left pane of the Cluster Manager window, the service group is shown below the CVM service group. On the Status tab, the group is shown Offline on each system.
3. In the left pane, double-click the `db2mpp_grp1` service group. The types of resources that can be configured for the group are displayed: Db2udb, IP, and NIC.



4. Double-click the Db2udb resource type. Select the resource, db2udb, below the Db2udb type and click on the Properties tab.
5. On the Properties tab for the db2udb resource, a list of Type Specific Attributes is shown. Click on 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.



After you have assigned the attribute values, the list of Type Specific Attributes resembles the following illustration:



6. 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.
7. 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 Parent Group as db2mpp\_grp1, the Child group as CVM, the Relationship as "online local," and the Dependency Type as "firm."
8. Click OK to create the dependency link.
9. Click the icon for Save Configuration.
10. 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.



11. Click the Online Service Group icon.
12. In the window, Select the service group and the system on which you want to online. Click OK.

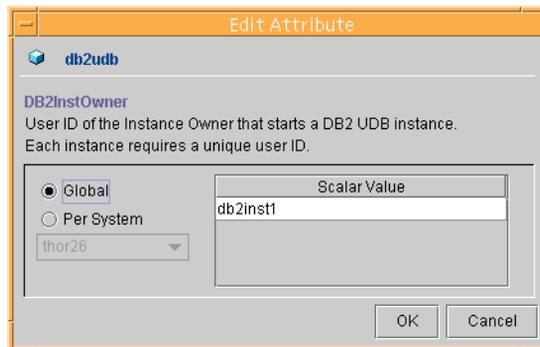
## Adding Service Group for DB2 UDB Non-MPP Using Cluster Manager

If you have imported the `Db2udbTypes.cf` file (see [“Importing the Db2udbTypes.cf File”](#) on page 30), you can use the `Db2udb_Group` template to configure a service group.

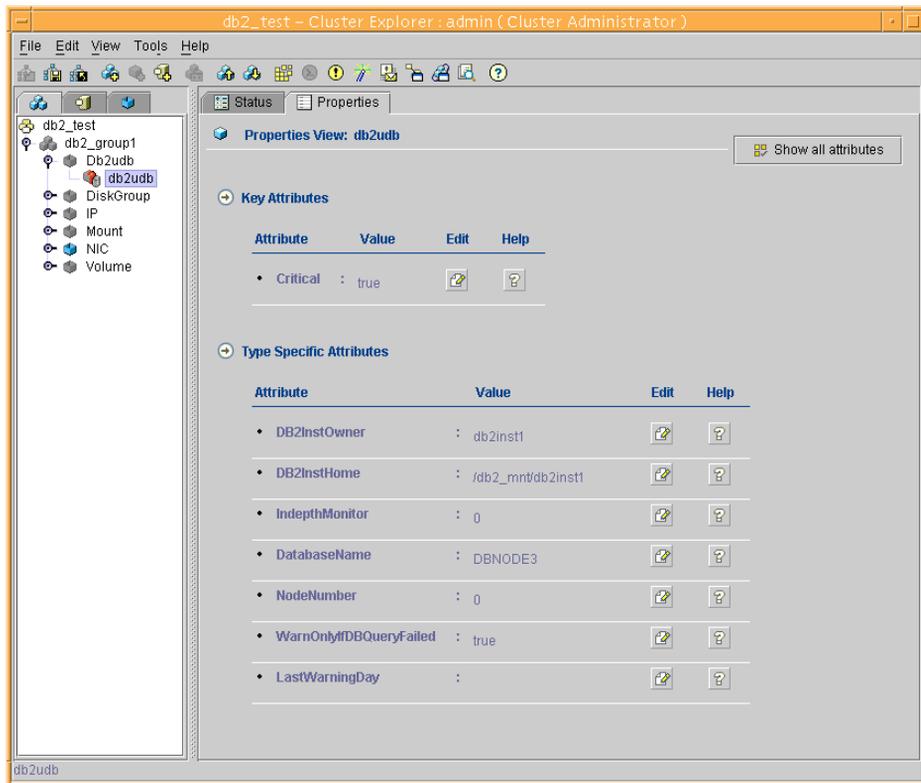
1. In the Cluster Explorer window, answer No when prompted to use the configuration wizard. Note: if you choose to use the wizard, the steps that follow are similar.
2. In the Cluster Explorer window, click the Add Service Group icon on the toolbar.
3. In the Add Service Group window, do the following:
  - a. 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.
  - b. From the systems shown in the Available Systems box, double-click those you want in your configuration.
  - c. In the window showing the systems added to the configuration, click the checkbox for a system on which you want to automatically start the service group.
  - d. Click the Failover radio button to specify the Service Group Type.
  - e. Click the Templates button.
  - f. In the Select Templates window, select `db2udb_grp` from the list shown 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.
  - g. Click OK on at the bottom of the Add Service Group window. The group is added. On the left pane of the Cluster Manager window, the service group is shown below the CVM service group. On the Status tab, the group is shown Offline on each system.



4. In the left pane, double-click the `db2_group1` service group. The types of resources that can be configured for the group are displayed: Db2udb, DiskGroup, IP, Mount, NIC, and Volume.
5. Double-click the Db2udb resource type. Select the resource, `db2udb`, below the Db2udb type and click on the Properties tab.
6. On the Properties tab for the `db2udb` resource, a list of Type Specific Attributes is shown. Click on 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`.



After you have assigned the attribute values, the list of Type Specific Attributes resembles the following illustration:



7. 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.
8. Click the icon for Save and Close Configuration.
9. 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.
10. Click the Online Service Group icon.
11. In the window, Select the service group and the system on which you want to online. Click OK.

## Configuring the DB2 UDB Agent by Editing the main.cf File

The VCS enterprise agent for DB2 UDB comes with three sample VCS configuration files installed in the `/etc/VRTSvcs/conf/sample_db2` directory. One sample is for an ESE single-partition instance configuration, another for a ESE multi-partition instance SMP configuration, and another for an ESE 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.

### Preparation for Editing 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. Ensure VCS is not running while you edit `main.cf` by using 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 DB2 agent.

### Configuring the Agent to Use the DB2 UDB MPP Configuration

Edit the `main.cf` file. Use `/etc/VRTSvcs/conf/sample_db2/main.cf.MPP` for reference. Notice that CVM service group is present in the configuration file.

1. Add an “include” line for the `Db2udbTypes.cf` file:
 

```
include "Db2udbTypes.cf"
```
2. Create service groups for the DB2 UDB resources. Refer to the sample configuration file “[MPP Configuration: DB2 UDB ESE Multi-Partition Instance](#)” on page 51. The example shows four DB2 MPP service groups and a CVM service group.



3. In the DB2 MPP service groups, include the definitions for the Db2udb, IP, and NIC resources, and assign values to the attributes for the resources to match the parameters of your configuration.

Refer to the [“Db2udb Resource Type Attributes”](#) on page 24 as well as the sample configuration files. Refer also to the *VERITAS Cluster Server Bundled Agents Reference Guide* for information about IP and NIC resources.

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. (See the *VCS User's Guide* for more information on assigning dependencies.) For example, referring to the [“MPP Configuration: DB2 UDB ESE Multi-Partition Instance”](#) on page 51, for the group db2mpp\_grp0 you would enter:

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

6. Save and close the file.

## Configuring the Agent to Use the DB2 UDB, Non-MPP Configurations

Edit the main.cf file. Use /etc/VRTSvcs/conf/sample\_db2/main.cf.EE or /etc/VRTSvcs/conf/sample\_db2/main.cf.EEE for reference:

1. Add an “include” line for the Db2udbTypes.cf file:

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

2. Create a service group for the DB2 UDB resources.

If you are using DB2 UDB ESE single-partition instance, refer to the example, [“Non-MPP Configuration: DB2 UDB ESE Single-Partition Instance”](#) on page 45 which shows two groups, “db2udb\_grp1,” and “db2udb\_grp2.”

If you are using DB2 UDB ESE multi-partition instance, refer to the example, [“Non-MPP Configuration: DB2 UDB ESE Multi-Partition Instance”](#) on page 49, which shows a group named “db2\_grp1” in which two partitions are defined.

3. Include all resources in the service groups, including the Db2udb, DiskGroup, IP, Mount, NIC, and Volume resources, and assign values to the attributes for the resources to match the parameters of your configuration.

Refer to the “[Db2udb Resource Type Attributes](#)” on page 24 as well as the sample configuration files. Refer also to the *VERITAS Cluster Server Bundled Agents Reference Guide* for information about the DiskGroup, IP, Mount, NIC, and Volume resources.

4. Assign dependencies for the newly created resources. Refer to the appropriate sample configuration file. (See the *VCS User’s Guide* for more information on assigning dependencies.) For example, referring to the “[Non-MPP Configuration: DB2 UDB ESE Single-Partition Instance](#)” on page 45, for the group db2udb\_grp1 you would 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 you would enter:

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

5. Save and close the file.

## Verifying the Configuration

After editing the main.cf file for you configuration, check the configuration using the following procedure:

1. Copy the DB2 UDB types configuration file into place:

```
# cp /etc/VRTSvcs/conf/Db2udbTypes.cf
   /etc/VRTSvcs/conf/config/Db2udbTypes.cf
```

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

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



3. Start the VCS engine on System A:

```
# hastart
```

4. Type the command `hastatus`:

```
# hastatus
```

5. When “LOCAL\_BUILD” is listed in the message column, start VCS on System B:

```
# hastart
```

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

```
# hagrps -display
```

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

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

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

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

9. Switch the DB2 UDB service group to System B:

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

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

```
# hagrps -display
```

11. 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
```



## Modifying the Agent Configuration

To dynamically reconfigure the VCS enterprise agent for DB2 UDB, use Cluster Manager 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. See the chapter on reconfiguring VCS from the command line in the *VERITAS Cluster Server User's Guide*.

## Enabling In-Depth Monitoring of DB2 UDB Instance

Shallow monitoring of a DB2 UDB instance involves either checking the “db2nps” output, which displays active processes for the instance or the database partition, or checking the exit status of the “db2gcf” command. By contrast, in-depth monitoring provides a higher level of confidence in the availability of the instance or partition and its database by making 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. It is recommended 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.

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:

```
# hagrps -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
# hagrps -unfreeze db2udb_grp1
```



## Disabling In-Depth Monitoring

You can dynamically disable in-depth monitoring as follows:

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:

```
# hagrps -freeze db2udb_grp1
```

3. Disable in-depth monitoring by assigning the `MonScript` attribute a null value. Use the `hares` command:

```
hares -modify resource IndepthMonitor 0
```

For example:

```
# hares -modify db2udb IndepthMonitor 0
# haconf -dump -makero
# hagrps -unfreeze db2udb_grp1
```



# Disabling and Removing the Agent

5

This chapter describes how to disable or remove the DB2 UDB agent.

## Disabling the Agent

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.

1. Determine if the service group is online by entering:

```
# hagrps -state service_group -sys system_name
```

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

```
# hagrps -switch service_group -to system_name
```

or:

```
# hagrps -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. See the chapter on reconfiguring VCS from the command line in the *VERITAS Cluster Server User's Guide* for more information.



## Removing the Agent

Type the following command on each system to remove the agent. Answer prompts accordingly:

```
# pkgrm VRTSvcsdb
```

# Sample Configuration Files



This chapter shows example DB2 UDB configurations in VCS configuration files.

## Non-MPP Configuration: DB2 UDB ESE Single-Partition Instance

The following configuration reflects DB2 UDB with two instances configured in a ESE single-partition instance environment.

```
include "types.cf"
include "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 db2udb_grp1 (
    SystemList = { sysA= 0, sysB = 1 }
    AutoStartList = { sysA }
)
```



```
Db2udb db2udb1 (  
    DB2InstOwner = db2inst1  
    DB2InstHome = "/db2inst1"  
    IndepthMonitor = 1  
    DatabaseName = SAMPLE  
)  
  
DiskGroup db2udb_dg1 (  
    DiskGroup = db2_dg1  
)  
  
IP db2udb_ip1 (  
    Device = hme0  
    Address = "166.98.9.163"  
)  
  
Mount db2udb_mnt1 (  
    MountPoint = "/db2inst1"  
    BlockDevice = "/dev/vx/dsk/db2_dg1/inst1_vol"  
    FSType = vxfs  
    MountOpt = rw  
    FsckOpt = "-y"  
)  
  
NIC db2udb_nic1 (  
    Device = hme0  
    NetworkType = ether  
)  
  
Volume db2udb_vol1 (  
    Volume = inst1_vol  
    DiskGroup = db2_dg1  
)  
  
db2udb1 requires db2udb_ip1  
db2udb1 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
//                         }
//                 }
//         }
//     }

group db2udb_grp2 (
    SystemList = { sysA = 0, sysB = 1 }
    AutoStartList = { sysA }
)

Db2udb db2udb2 (
    DB2InstOwner = db2inst2
    DB2InstHome = "/db2inst2"
    IndepthMonitor = 1
    DatabaseName = MYDB
)

DiskGroup db2udb_dg2 (
    DiskGroup = db2_dg2
)

IP db2udb_ip2 (
    Device = hme0
    Address = "192.2.40.21"
)
```



```

Mount db2udb_mnt2 (
    MountPoint = "/db2inst2"
    BlockDevice = "/dev/vx/dsk/db2_dg2/inst2_vol"
    FSType = vxfs
    MountOpt = rw
    FsckOpt = "-y"
)

NIC db2udb_nic2 (
    Device = hme0
    NetworkType = ether
)

Volume db2udb_vol2 (
    Volume = inst2_vol
    DiskGroup = db2_dg2
)

db2udb2 requires db2udb_ip2
db2udb2 requires db2udb_mnt2
db2udb_ip2 requires db2udb_nic2
db2udb_mnt2 requires db2udb_vol2
db2udb_vol2 requires db2udb_dg2

// resource dependency tree
//
//     group db2udb_grp2
//     {
//     Db2udb db2udb2
//     {
//     IP db2udb_ip2
//     {
//     NIC db2udb_nic2
//     }
//     Mount db2udb_mnt2
//     {
//     Volume db2udb_vol2
//     {
//     DiskGroup db2udb_dg2
//     }
//     }
//     }
//     }
//     }

```



## Non-MPP Configuration: DB2 UDB ESE Multi-Partition Instance

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

```
include "types.cf"
include "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 = hme0  
    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 = hme0  
    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 Configuration: DB2 UDB ESE Multi-Partition Instance

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 CVM service group, which manages the shared storage in the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "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 }
)

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

CFSQlogckd qlogckd (
    Critical = 0
)

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  
//     {  
//     CFSSMount 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 = hme0
    Address = "10.118.2.144"
    NetMask = "255.255.248.0"
)

NIC Db2_NIC0 (
    Device = hme0
    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 = hme0
    Address = "10.118.2.145"
    NetMask = "255.255.248.0"
)

NIC Db2_NIC1 (
    Device = hme0
    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 = hme0
    Address = "10.118.2.146"
    NetMask = "255.255.248.0"
)

NIC Db2_NIC2 (
    Device = hme0
    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 = hme0  
    Address = "10.118.2.147"  
    NetMask = "255.255.248.0"  
)  
  
NIC Db2_NIC3 (  
    Device = hme0  
    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  
//     }  
//     }  
//     }  
// }
```



## Sample DB2 Instance Running on 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 "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"
)

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
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

---

## Symbols

- .rhosts, editing for database access 18
- /etc/services, checking assigned ports 17
- /etc/system file (shared memory parameters) 15

## A

- accounts, required for DB2 UDB users 15
- Action entry point operation 8
- agent
  - configuring, overview 23
  - operations described 3
- attributes of Db2udb resource type 24

## B

- binaries, DB2 UDB, installing locally 15

## C

- clean operation 6
- Cluster Manager
  - configuring DB2 service group 30
  - using to import Db2udbTypes.cf 30
- commands
  - groupadd 14
  - hagrp 43
  - hagui 30
  - mkfs 12, 13
  - pkgrm 44
  - rsh 18
  - useradd 15
  - vxassist 12, 13
  - vxctl 13
  - vx dg 12, 13
- commands, DB2
  - create 20
  - list 20
- configuration files
  - db2nodes.cfg, host for partitions 19
  - Db2udbTypes.cf 26
  - main.cf, example 45, 49, 51

- configuring DB2 UDB agent
  - by editing configuration files 37
  - using Cluster Manager 30
- create command (DB2) 20
- CVM (infrastructure) service group 28

## D

- DatabaseName attribute 5, 24
- DB2 commands
  - create 20
  - list 20
- DB2 service group 28
- DB2 UDB EE (Enterprise Edition)
  - sample configuration 45
- DB2 UDB ESE (Enterprise Server Edition)
  - sample configuration file 49, 51
- DB2 UDB instances
  - installing on shared storage 16
  - two instances configured, example 45
- DB2 UDB type definition file 26
- db2gcf -d command 3
- db2gcf -k command (clean operation) 6
- db2gcf -u command 3
- DB2InstHome attribute 24
- DB2InstOwner attribute 24
- db2nkill program (clean operation) 6
- db2nodes.cfg, virtual host for partitions 19
- db2nps command checks processes 4
- db2setup tool, for binaries, instances 15
- db2start program 3, 20
- db2stop program 3, 20
- Db2udb resource
  - depends on IP, Mount resources 27
  - example in main.cf file 49
- Db2udbTypes.cf
  - including in main.cf file 37, 38
  - type definition file 26
- Db2udbTypes.cf file
  - importing 30



---

- dependency
  - between DB2 and CVM service groups 33
  - defining in main.cf 39
- DiskGroup resource
  - example in main.cf 50
  - required by Volume resource 27

## E

- Encoding attribute 24

## F

- file system
  - creating for MPP configuration 13
  - creating for non-MPP instances 12
- files
  - .rhosts 18
  - /etc/services 17
  - /etc/system 15
  - db2nodes.cfg 19
  - Db2udbTypes.cf 26
  - main.cf (VCS configuration) 37, 45

## G

- groupadd command 14
- groups (user), adding 14

## H

- hagrp command 43
- hagui command 30

## I

- importing Db2udbTypes.cf file 30
- in-depth monitoring
  - described 5
  - disabling 42
  - enabling 41
  - enabling with IndepthMonitor flag 24
- IndepthMonitor attribute 5, 24
- Info entry point operation 7
- installing DB2 UDB
  - confirming setup 20
  - instances on shared storage 16
  - requirements 15
  - system binaries 15
- instances
  - installing on shared disk 16
  - starting and stopping 20
- IP address, virtual 19
- IP resource
  - example in main.cf file 50
  - requires NIC resource 27

## J

- Java Console
  - using to configure service group 31, 34

## L

- list command (DB2) 20

## M

- main.cf file
  - editing 37
  - example for DB2 UDB EEE 49
  - example for DB2 UDB ESE multi-partition 51
  - examples 37
  - verifying configuration 39
- mkfs command 12, 13
- monitor operation 4
- monitoring
  - configuring in-depth monitoring 41
  - ESE multi-partition instance, MPP 5
  - ESE multi-partition instance, SMP 4
  - ESE single-partition instance 4
  - shallow and in-depth 41
- Mount resource
  - example in main.cf file 50
  - requires the Volume resource 27
- MPP (Massively Parallel Processing), support 1
- MPP multi-partition, monitoring 5
- MPP service group configuration 28

## N

- NIC resource
  - example in main.cf file 50
  - required by IP resource 27
- NodeNumber attribute 24

## O

- offline operations, commands used 3
- online operations, commands used 3

## P

- parallel mode for CVM service group 28
- partition number (NodeNumber) 24
- pkgm command 44
- ports, checking /etc/services for 17

## R

- ResourceInfo resource attribute 7
- RestartLimit group attribute 4
- rsh command 18



---

**S**

- service group
  - adding using cluster manager 31
  - adding using wizard 34
  - creating in main.cf file 37, 38
  - CVM 28
  - DB2 28
  - dependency link DB2 and CVM 33
  - MPP configuration, dependencies 28
  - non-MPP configuration, dependencies 27
- shallow monitoring, described 41
- shared memory parameters 15
- SMP (Symmetric Multiprocessing), support 1
- SMP multi-partition, monitoring 4
- starting DB2 instance 3, 20
- stopping DB2 instance 3

**T**

- template
  - for DB2 MPP service group 31
  - using for non-MPP service group 34
- testing

- remotely logging in 18
- service groups are configured 40
- ToleranceLimit group attribute 4
- type (DB2) definition file 26

**U**

- user accounts, adding 15
- useradd command 15

**V**

- VCS
  - commands for reconfiguration 41
  - example configuration files 37, 45
  - supported version 1
- virtual IP address 19
- Volume resource
  - example in mian.cf file 46
  - requires the DiskGroup resource 27
- vxassist command 12, 13
- vxctl command 13
- vxdbg command 12, 13

**W**

- WarnOnlyIfDBQueryFailed attribute 4, 5, 25

