DB2 Universal Database Version 8.2 for Linux HOWTO

Ian Hakes

Rory Hinton

Dan Scott

Revision History

Revision 1.1.7	2006−01−13	Revised by: dbs
Add Fedora Core 4, OpenSuse 10.0, and Ubuntu 5.10.		
Revision 1.1.6	2005−11−09	Revised by: dbs
Minor updates for Ubuntu installation steps.		
Revision 1.1.5	2005−10−31	Revised by: idmh
Updated user setup instructions for non−DB2 created users.		
Revision 1.1.4	2005−07−13	Revised by: dbs
Added Ubuntu instructions.		
Revision 1.1.3	2005−04−20	Revised by: idmh
Added Red Hat Enterprise 4 instructions. Corrected minor cross−reference link issues.		
Revision 1.1.2	2005−02−07	Revised by: idmh
Updated non−RPM install instructions. Minor updates to clarify audience and add some late breaking information.		
Revision 1.1.1	2005−01−25	Revised by: idmh
Added Gentoo install instructions (provided by Dan Scott)		
Revision 1.0.1	2005−01−06	Revised by: idmh
Fixed typos in Red Hat 2.6 Enhancements section		
Revision 1.0.0	2004−12−09	Revised by: idmh
Initial release of DB2 V8.2 HOWTO.		

This HOWTO is intended for anyone interested in installing and using 32−bit IBM DB2® Universal Database on Linux®. If you're interested in 64−bit or other architectures, we refer you to the official documentation sources (see *Resources* for a link).

Table of Contents

Table of Contents

Table of Contents

1. Introduction

1.1. Who should read this HOWTO

This HOWTO is intended for anyone interested in installing and using 32−bit IBM DB2® Universal Database on Linux®. If you're interested in 64−bit or other architectures, we refer you to the official documentation sources (see *Resources* for a link).

The information found in this document is based on our experiences installing and configuring DB2 UDB Version 8.2.

1.2. Disclaimer

No liability for the contents of this document can be accepted. Use the concepts, examples and information at your own risk. There may be errors and inaccuracies, that could be damaging to your system. Proceed with caution, and although it is highly unlikely that accidents will happen because of following advice or procedures described in this document, the authors do not take any responsibility for any damage claimed to be caused by doing so.

All copyrights are held by their by their respective owners, unless specifically noted otherwise. Use of a term in this document should not be regarded as affecting the validity of any trademark or service mark. Naming of particular products or brands should not be seen as endorsements.

This HOWTO is not official IBM documentation, and cannot be construed as such. IBM accepts no responsibility, liability or coverage for the material or instructions contained in this HOWTO. The views expressed in this document are those of the authors and not IBM Corporation.

1.3. Copyright and license

1.4. Credits

Ian Hakes is the current owner and maintainer of this HOWTO. Please send all suggestions for improvement, criticisms or questions to db2howto@REMgmail.com (remove the REM).

Big thanks to Dan Scott, who kindly handed this project over with his blessing, and even found time to handle the DocBook conversion for this version. Dan is responsible for getting both of us interested in DB2 UDB on Linux throughout the past four years. His work on the original DB2 UDB Version 7.1 for Linux HOWTO, published by the LDP, stands as a testament to his hard work and dedication and we hope this document can

live up to those high standards.

We'd also to thank: Rav Ahuja, Grant Hutchison, Ryan Chase, and Darin McBride for their input and patience in answering our endless stream of questions.

Most of all, we must thank our wives and children for allowing us to spend weeknights and weekends clattering away on our laptops and servers to get this document together. We couldn't have done it without their love and support.

1.5. About the authors

Ian Hakes, R.H.C.E., has worked for IBM in DB2 UDB Information Development since 1999. Currently, he is the team lead for the DB2 UDB Information Development Infrastructure team, developing solutions to improve writers lives as they put together the official documentation. Before that, he was the writer responsible for the DB2 UDB SQL Reference manual. His free time is spent mainly on computers, thinking about how much nicer it would be to be outside sailing or hiking.

Rory A.A. Hinton, Ph.D., is presently employed at ATI Technologies. He works in Application Engineering, writing BIOS development and register programming reference guides. In his spare time, he cuts his lawn, drinks Ricard with ice and water, and hosts a weekly talk radio show on CIUT 89.5 FM in Toronto. He tries to live by the motto: "I want to work, provide for my family, and die quietly without pain."

2. Prerequisites

Prerequisites are the elements required on your system before you can install and use DB2 UDB Version 8.2. These elements include hardware, software, and communication pieces required for DB2 UDB Version 8.2 to run successfully. This section covers:

- Hardware requirements:
  - Supported processors
  - Disk space requirements
  - Memory requirements
- Software requirements:
  - Linux distribution
  - Kernel levels and libraries requirements
  - Additional software requirements
  - The IBM Developer Kit (SDK) for Java"

2.1. Hardware requirements

2.1.1. Supported processors

- x86 (Intel® Pentium®, Intel Xeon, and AMD Athlon)
- x86−64 (Intel EM64T and AMD64)
- IA64 (Intel Itanium 2 or later)
- POWER" (any iSeries" or pSeries® that supports Linux)
- zSeries

2.1.2. Disk space requirements

Take into account the disk space requirements before you install and configure DB2 UDB on your distribution:

Table 1. Disk space requirements

Install type	Description	Required disk space
Typical	DB2 UDB is installed with most features and functionality, including graphical tools such as the Control Center and the Development Center.	450 to 500 MB
Compact	DB2 UDB is installed with basic features and functionality, without the graphical tools.	350 to 400 MB
Custom	This option allows you to select the features and functionality that you want to install.	350 to 700 MB

It's important to allocate enough disk space for the required software prerequisites, along with the product documentation (if desired) and the databases you create. For example, in Version 8.2, DB2 UDB product documentation is accessed through the DB2 UDB Information Center. This is a separately installable product that requires at least 100 MB when installing the core English language files. Additional space is required for additional languages. However, if disk space is an issue, you can configure DB2 UDB to access the Information Center from the Web. More detail about this great feature is provided in *Installing the DB2 UDB Information Center*.

2.1.3. Memory requirements

You should allocate a minimum of 256 MB of RAM for a product like the DB2 UDB Enterprise Server Edition, but additional memory should be allocated for other software and communication processes. When determining memory requirements, consider the following points:

- Your SWAP space should be at least twice as much as your RAM.
- Additional memory should be allocated for any non−DB2 UDB software that may be running on your system.
- Additional memory is required to support database clients and database activity.
- These memory requirements will be affected by the size and complexity of your database system, as well as specific performance requirements.

2.2. Software requirements

2.2.1. Distribution levels that are supported by DB2 UDB Version 8.2

In DB2 UDB Version 6, only 32−bit Intel architectures were supported. With the release of DB2 UDB Version 8.2, IBM has demonstrated its support for the Linux platform by validating on a much broader range of architectures. For the sake of convenience, this table lists the recommended and validated distributions that DB2 UDB officially supports in Version 8.2.

Table 2. Recommended and validated distributions

Distributions	Platforms							
	x86	x86_64		IA64	POWER		zSeries®	
		AMD64	EM64T					
	32−bit	64−bit	64−bit	64−bit	32−bit	64−bit	31−bit	64−bit
Red Hat Enterprise Linux (RHEL) 3	x	x	x	x	x	x	x	x
SuSE Linux Enterprise Server (SLES) 8	x	x		x	x		x	x
SuSE Linux Enterprise Server (SLES) 9	x	x	x	x	x	x	x	x

For a list of all of the distributions that are officially supported, as well as the latest information about kernels, libraries, integrated cluster environments, and papers, visit the DB2 UDB for Linux validation Web site listed in *Resources*.

2.2.2. Required kernel levels and libraries

2.2.2.1. Kernel parameter values

To successfully run DB2 UDB on your Linux distribution, you may be required to update some of the default kernel parameter settings. For example, the 2.4.x series kernel message queue parameter msgmni has a default value that allows only a limited number of simultaneous connections to DB2 UDB.

These are the recommended values for DB2 UDB to run optimally:

- kernel.shmmax=268435456 for 32−bit
- kernel.shmmax=1073741824 for 64−bit
- kernel.msgmni=1024
- fs.file−max=8192
- kernel.sem="250 32000 32 1024"

Fortunately, DB2 UDB Version 8.2 has a new feature that checks the values of the kernel.semmni, kernel.msgmni, and kernel.shmmax parameters when you enter the db2start command, and changes them for you if the current values are not optimal. This new feature makes these changes:

- The semmni kernel parameter is changed to 1024
- The msgmni kernel parameter is changed to 1024
- The shmmax kernel parameter is changed to 268435456 (32−bit) or 1073741824 (64−bit)

For example, after you issue the db2start command for the first time, you should receive output like the following messages in your db2diag.log file:

```
ADM0506I DB2 has automatically updated the "semmni" kernel
parameter from "128" to the recommended value "1024".

2004-07-31-16.38.59.074791 Instance:db2inst1
Node:000
PID:15996(db2sysc) TID:8192 Appid:none
base sys utilities sqlesysc_main Probe:9

ADM0506I DB2 has automatically updated the "msgmni" kernel
parameter from "16" to the recommended value "1024".

2004-07-31-16.38.59.076916 Instance:db2inst1
Node:000
PID:15996(db2sysc) TID:8192 Appid:none
base sys utilities sqlesysc_main Probe:9

ADM0506I DB2 has automatically updated the "shmmax" kernel
parameter from "33554432" to the recommended value "268435456".

2004-07-31-16.39.01.262594 Instance:db2inst1
Node:000
PID:15994(db2star2) TID:8192 Appid:none
base sys utilities startdbm Probe:911

ADM7513W Database manager has started.
```

Notice that db2start did not update the fs.file−max parameter. Any kernel at 2.4.18 and later automatically updates this parameter. Because of this new feature, it is no longer necessary for you to manually update the kernel.shmmax, kernel.msgmni, and kernel.sem parameters prior to installing DB2 UDB.

2.2.2.2. Manually updating kernel parameters

If, for some reason, you need to change these default settings, you can update them manually. To check your current shared memory segment, semaphore array, and message queue limits, enter the ipcs –l command. Your output should look something like this:

```
        ------ Shared Memory Limits --------
        max number of segments = 4096
        max seg size (kbytes) = 32768
        max total shared memory (kbytes) = 8388608
        min seg size (bytes) = 1
        ------ Semaphore Limits --------
        max number of arrays = 1024
        max semaphores per array = 250
        max semaphores system wide = 32000
        max ops per semop call = 32
        semaphore max value = 32767
        ------ Messages: Limits --------
        max queues system wide = 1024
        max size of message (bytes) = 8192
        default max size of queue (bytes) = 16384
```

For example, here are instructions on how to modify the kernel.shmmax, kernel.sem, and the kernel.msgmni parameters on Red Hat Linux 8, and how to keep them set after a reboot.

1. Log on as a user with root authority
2. Open up `/etc/sysctl.conf` in a text editor and add entries:
   ```
   kernel.shmmax=268435456
   kernel.msgmni=1024
   kernel.sem="250 32000 32 1024"
   ```
3. Enter the sysctl –p command to load in sysctl settings from `/etc/sysctl.conf`.
4. Enter the ipcs –l to view the updated kernel parameters in sysctl.

To view all sysctl settings, enter: **sysctl –a**. You can also temporarily update the kernel parameters (during run time) by using the **sysctl –w** command. For example, to change the kernel.msgmni parameter to 1024, enter the command:

```
sysctl -w kernel.msgmni=1024
```

These settings will not remain after the next reboot unless they are saved in the `/etc/sysctl.conf` file.

Again, consider the same procedure on SuSE Linux. Modifying kernel parameters on SuSE Linux is a little different from the method on Red Hat Linux. These instructions will explain how to update the kernel.shmmax, kernel.sem, and the kernel.msgmni parameters, and how to set them for reboot.

1. Log in as a user with root authority.
2. Some SuSE Linux distributions do not have a `/etc/sysctl.conf` file. If so, you need to create one manually using a text editor.
3. In the `/etc/sysctl.conf` file, add entries:
   ```
           kernel.shmmax=268435456
           kernel.msgmni=1024
           fs.file-max=8129
           kernel.sem="250 32000 32 1024"
   ```
4. Enter the **sysctl –p** command to load in sysctl settings from `/etc/sysctl.conf`.

5. Add sysctl −p to a system initialization file to set kernel parameters after each reboot. To do this, write a script and configure it to run automatically at runlevel 5. Specifically, you need to create an executable file in `/etc/init.d/`, and then add pointers to this script in `/etc/init.d/rc5.d`. For example, in `/etc/init.d` create an executable file named **kerneldb2** that contains the script:

```
#! /bin/sh
#
#
# /etc/init.d/kerneldb2
#
### END INIT INFO
touch /var/lock/subsys/kerneldb2
/sbin/sysctl −p >> /var/lock/subsys/kerneldb2
```

   Then in `/etc/init.d/rc5.d`, we added pointers to the kerneldb2 script by entering the commands:

```
cd /etc/init.d/rc5.d
ln −s ../kerneldb2 S99kerneldb2
ln −s ../kerneldb2 K99kerneldb2
```

   Like Red Hat Linux, you can temporarily update the kernel parameters (during run time) by using the **sysctl −w** command. For example, to change the kernel.sem parameter, enter the command:

```
sysctl −w kernel.sem="250 32000 32 1024"
```

   Again, these settings will not remain after the next reboot unless they are saved in the `/etc/sysctl.conf` file.

2.2.3. The IBM Developer Kit for Java (SDK), Version 1.3.1 and 1.4.1

If the IBM SDK for Java is required by a component being installed, and it's not already installed on your computer, the SDK for Java will be installed if you use either the DB2 UDB Setup wizard or a response file to install the product. If you use the db2_install utility to install DB2 UDB, the SDK won't get installed. In that case, you must install the rpm file manually.

You require the appropriate level of IBM Software Development Kit for Java (SDK) to use Java−based tools like the Control Center and the Development Center, and to create and run Java applications, including stored procedures and user−defined functions.

Table 3. SDK levels by operating system

Operating system	SDK for Java level
Linux x86	SDK 1.4.1 Service Release 2
Linux IA64	SDK 1.4.1 Service Release 2
Linux x86−64 (hybrid)	32−bit SDK 1.4.1 Service Release 2
Linux 390	SDK 1.4.1 Service Release 2
	SDK 1.4.1 Service Release 2
Linux PPC (hybrid)	SDK 1.4.1 Service Release 2

To manually install the SDK rpm file:

1. Run one of these rpm commands appropriate for your Linux operating system from the `/cdrom/db2/linux/Java-1.4` directory:

Table 4. RPM commands

Operating system	Command	Destination Directory
Linux IA32	rpm −ivh IBMJava2−SDK−1.4.1−2.0.i386.rpm	/opt/IBMJava2−141
Linux IA64	rpm −ivh IBMJava2−SDK−1.4.1−2.0.ia64.rpm	/opt/IBMJava2−141
Linux 390	rpm −ivh IBMJava2−SDK−1.4.1−2.0.s390.rpm	/opt/IBMJava2−s390−141
Linux 390 64−bit	rpm −ivh IBMJava2−SDK−1.4.1−2.0.s390x.rpm	/opt/IBMJava2−s390−141
Linux PPC 32−bit	rpm −ivh IBMJava2−SDK−1.4.1−2.0.ppc.rpm	/opt/IBMJava2−ppc−141
Linux PPC 64−bit	rpm −ivh IBMJava2−SDK−1.4.1−2.0.ppc64.rpm	/opt/IBMJava2−ppc64−141

2. Set up the Java environment for all users by adding these lines to your `/etc/profile` file:
   ```
   export PATH=$PATH:/opt/dest-dir/jre/bin
   ```
   where `dest-dir` corresponds to the destination directory listed in the table above.

When you install DB2 UDB on Linux, the RPM−based installation attempts to install IBM's version of Java (IBM SDK 1.4.1 SR2). If a later installed level of the RPM, such as IBM SDK 1.5.1 SR1, is detected, the back−level RPM is not installed.

However, in this case, the installation leaves the JDK_PATH database configuration parameter pointing to the back−level path. As a result, none of the Java−dependent functionality, including the installation of the DB2 UDB Tools Catalog, will work.

To solve this problem, run the following commands as the instance owner (for example, db2inst1):

```
db2 UPDATE DBM CFG USING JDK_PATH existing SDK path
db2 UPDATE ADMIN CFG USING JDK_PATH existing SDK path
```

where `existing SDK path` points to the previous Java path. These changes will point DB2 UDB to the correct IBM Developer Kit for the Java−based tools.

2.2.4. Additional software requirements

Depending on your DB2 UDB requirements, you may be required to install additional software packages for DB2 UDB to function properly. Make sure that these software packages are installed prior to using DB2 UDB.

- X Window System software, capable of rendering a graphical user interface. You need this if you want to use the DB2 UDB Setup wizard graphical installer, or any of the DB2 UDB graphical tools.
- A Web browser, to view topics in the DB2 UDB Information Center.
- If you are using NIS or NIS+ for security authentication on your system, you must create DB2 UDB user accounts before installing DB2 UDB.
- Additional packages, as outlined in the following table:

  Table 5. Package requirements

Package name	RPM name	Description
pdksh	pdksh (version 5.2 or later)	Korn Shell. Required for multiple−partition database environments.
rsh−server	rsh (version 0.17 or later)	Contains a set of programs that allow users to run commands on a remote computer. Required for multiple−partition environments.
nfs−utils	nfs−utils (version 1.0 or later)	Network File System support package. It allows access for local files to remote machines. Required for systems utilizing NFS mounted drives.

To check whether you have these packages installed, use the **rpm −q** commands:

```
rpm −qa | grep pdksh
rpm −qa | grep rsh
rpm −qa | grep nfs
```

For each command, if you have these packages installed, your system should return the following output:

```
pdksh−5.2.14−19
rsh−0.17−10
rsh−server−0.17−10
redhat−config−nfs−1.0.1−3
nfs−utils−1.0.1−2
```

To install these packages on Red Hat Linux, use the Red Hat Linux installation CDs, or the **up2date** utility, and the **rpm** command. For example, to install pdksh on from the Red Hat Linux CD, mount the Red Hat Linux CD #3, and enter this command as root:

```
rpm −ivh /mnt/cdrom/RedHat/RPMS/pdksh−5.2.14−19.i386.rpm
```

To install using the up2date package, run this command as root:

```
up2date pdksh
```

To install these packages on SuSE Linux, you can also use **rpm** or use the handy YaST (Yet another Setup Tool) utility. For example, to install pdksh on SuSE Linux Pro 8 using YaST:

1. Log on as root
2. From the YaST Control Center, select Software −> Install or Remove Software, then choose the appropriate packages.

3. Installation considerations and planning

After verifying that your system meets the basic requirements, you need to plan your installation by considering installation methods, storage planning, and the creation of users and groups:

- Installation choices:
    - db2setup installation
    - db2_install installation
    - response file installation
    - RPM command installation
- Storage planning:
    - Raw devices versus file systems
        - Raw device configuration
        - File system configuration
    - Log storage
    - Temporary table space storage
- User and group setup:
    - User and group requirements
    - Creating users − single−partition
    - Creating users − multiple−partition
        - DAS user considerations for a multiple−partition environment

3.1. Installation methods

You can use four methods to install DB2 UDB:

- **db2setup** utility
- **db2_install** utility
- A DB2 response file
- Linux **rpm** command

Each method has its own advantages and disadvantages. The preferred method often depends on your level of expertise and type of environment, but in general, if a graphical terminal is available, using db2setup is recommended.

3.1.1. db2setup installation

The **db2setup** command starts the DB2 UDB Setup wizard, a Java−based graphical tool that installs DB2 UDB. It lays down the DB2 UDB filesets, the IBM SDK for Java, and allows you to create DB2 UDB instances, create new users and groups, configure existing users, configure communications, create the tools catalog database, and set up notification. The DB2 UDB Setup wizard also allows you to create a response file.

Using **db2setup** is the best method for less experienced users, because most of the configuration is automatically performed during the installation. Again, you require X Window System software capable of rendering a graphical user interface, to use db2setup successfully.

3.1.2. db2_install installation

The **db2_install** command starts the **db2_install** script. This script installs all DB2 UDB packages on your Linux system using the RPM installation utility. This method is reliable and commonly used by expert users for installing DB2 UDB on larger, more complex multiple−partition systems. However, tasks such as setting up users and groups, creating instances, tools catalog database creation, and notification setup need to be performed manually after the installation is completed.

A limitation to **db2_install** is that it only installs user assistance (like help, messages, and tool interfaces) in English, whereas **db2setup** installs support for one or more different languages. Also, a **db2_install** installation can take longer (considering the higher number of manual configuration tasks), it requires a higher level of knowledge and skill, and it cannot create response files.

3.1.3. Response file installation

A response file can be created using the DB2 UDB Setup wizard or by editing a sample response file. It allows you to install DB2 UDB across multiple machines with consistent installation and configuration settings. A response file installation is fast, because it bypasses the graphical wizard and does the configuration for you. Another advantage of using a response file is that it creates a Database Administration Server (DAS) on each machine, while with **db2_install** the DAS must be created manually after installation.

The sample response file can be used to install DB2 UDB, configure users, create instances, set up notification, create the tools catalog, and configure a large number of DBM parameters. This is the quickest installation method if you already have all the information you need. Unlike **db2setup**, the response file installation is not interactive, and it takes longer to prepare the response file.

3.1.4. RPM command installation

The Red Hat Package Manager (RPM) is a Linux software installation command. The various DB2 UDB installation methods use RPM to lay down the appropriate files on the system. Using the RPM command to install DB2 UDB allows you to select specific DB2 UDB files. However, the RPM installation method will only install the DB2 UDB code. It will not create instances, user IDs, or response files during DB2 installation. Please note that this installation method is not officially supported or recommended.

3.2. Storage planning

In this section these storage considerations are discussed:

- Log storage
- Temporary table space storage

3.2.1. Log storage

By default, DB2 UDB sets the log path to the default database path during database creation.

You should store both the primary copy of the logs and the mirror logs each on a physically separate disk, preferably one that is also on a different disk controller.

Mirror logs are created using the MIRRORLOGPATH configuration parameter. Log mirroring allows the database to write an identical second copy of log files to a different path. For example, you can change the primary log path from the default to `/db2log1`, and set the mirror log path to `/db2log2` by issuing these commands:

```
db2 update db cfg for db_name using NEWLOGPATH /db2log1
db2 update db cfg for db_name using MIRRORLOGPATH /db2log2
```

These changes will only take place after you deactivate and activate your database, using the **db2stop** and **db2start** commands.

3.2.2. Temporary table space storage

DB2 UDB uses system temporary table spaces for many SQL operations, such as JOIN and SORT. The temporary table space, TEMPSPACE1, is one of the three default table spaces (SYSCATSPACE, TEMPSPACE1, and USERSPACE1) that gets created during database creation. By default, TEMPSPACE1 gets placed in the database path. For larger systems, we recommend that your temporary tables paces are located on a separate file system and disk.

In a multiple−partition database environment, the catalog node should contain all three default table spaces, and the other database partitions should each contain only TEMPSPACE1 and USERSPACE1.

The following example shows how to create a system temporary table space on multiple nodes in the `/db2temp` file system:

```
connect to db_name;
create temporary tablespace TEMPSPACE01 in IBMTEMPGROUP
managed by SYSTEM
using ('/db2temp/$INSTANCE/db_name/n001tmp/tempspace01') on \
    dbpartitionnum (1)
using ('/db2temp/$INSTANCE/db_name/n002tmp/tempspace01') on \
    dbpartitionnum (2)
using ('/db2temp/$INSTANCE/db_name/n003tmp/tempspace01') on \
    dbpartitionnum (3)
using ('/db2temp/$INSTANCE/db_name/n004tmp/tempspace01') on \
    dbpartitionnum (4)
extentsize 32
prefetchsize 128
bufferpool IBMDEFAULTBP
overhead 24.1
transferrate 0.9;
drop tablespace
tempspace1;
```

3.3. User and group setup

DB2 UDB requires user IDs and groups. This section discusses these IDs and groups and how to create them.

3.3.1. User and group requirements

DB2 UDB requires a minimum of three users and groups: the instance owning user, the fenced user, and the Database Administration Server (DAS) user. You may use the default names provided by the setup utilities, or specify your own user and group names. For example, consider the following default user IDs and group names produced by the DB2 UDB Setup wizard:

Table 6. Default user IDs and group names

Required user	User name	Group name	Description
Instance owner	db2inst1	db2iadm1	Administers the instance
Fenced user	db2fenc1	db2fadm1	Responsible for executing fenced user defined functions, such as JDFs and stored procedures.
DAS user	dasusr1	db2asgrp	Administers the DB2 UDB Administration Server

3.3.2. Creating users − single−partition

You must have root authority to create users and groups. There are three ways in which you can create a DB2 UDB user ID:

- DB2 UDB Setup Wizard. The DB2 UDB Setup Wizard creates all of the required users and groups for you during installation. The default users and groups that get created are displayed in the previous table. DB2 UDB Setup also gives you an option to specify your own user and group names.
- Response file. Users can also be created during a response file installation if you specify user and group information in the response file. For example, consider the following entries in a response file that create the three required users and groups for DB2 UDB:

```
* DAS user
DAS_USERNAME = dasusr1
DAS_GROUP_NAME = dasadm1
DAS_HOME_DIRECTORY = /home/dasusr1
DAS_PASSWORD = ******
* Instance-owning user
inst1.NAME = db2inst1
inst1.GROUP_NAME = db2grp1
inst1.HOME_DIRECTORY = /db2home/db2inst1
inst1.PASSWORD = ******
inst1.AUTOSTART = YES
inst1.AUTHENTICATION = SERVER
inst1.SVCENAME = db2c_db2inst1
inst1.PORT_NUMBER = 50001
inst1.FCM_PORT_NUMBER = 60000
inst1.MAX_LOGICAL_NODES = 4
```

```
* Fenced user
nst1.FENCED_USERNAME = db2fenc1
inst1.FENCED_GROUP_NAME = db2fgrp1
inst1.FENCED_HOME_DIRECTORY = /db2home/db2fenc1
inst1.FENCED_PASSWORD = ******
```

- Manually using command line. To use this method, follow these steps:

  1. Log on to your machine as root.
  2. Create groups for the instance owner, the fenced user and the DAS user by using these commands:
     ```
     groupadd db2iadm1
     groupadd db2fadm1
     groupadd db2asgrp
     ```
     Linux will create the groups using the first available group ID (GID) numbers available over 500. You can specify GID numbers if you prefer.
  3. Create a user that belongs to each group and specify the home directory. For example, you can choose to place all home directories in `/db2home` by entering these commands:
     ```
     useradd -g db2iadm1 -m -d /db2home/db2inst1 db2inst1 -p password1
     useradd -g db2fadm1 -m -d /db2home/db2fenc1 db2fenc1 -p password2
     useradd -g db2asgrp -m -d /db2home/dasusr1 dasusr1 -p password3
     ```

3.3.3. Creating users − multiple−partition

In a multiple−partition database environment, you only need to create one shared home directory for the instance owner and fenced user (but remember to create users on each computer). When creating users in a multiple−partition environment, make sure that the user and group IDs are the same on each machine. Consider the following setup:

- A shared home directory, `/db2home`, on the instance−owning machine, which is NFS−mounted on the remaining machines in the cluster. In this directory are found the home directories for the instance−owning user and fenced user: db2inst1 and db2fenc1.
- A local home directory for the DAS user on each computer, dasusr1, which is stored in the `/home` directory.

The DB2 Administration Server (DAS) has changed significantly in DB2 UDB Version 8. Therefore, you should take note of the following DAS user considerations.

3.3.3.1. DAS user considerations for a multiple−partition environment

- A DAS must be running on each physical machine in the multiple−partition database for the graphical administration tools (for example, Control Center) to work.
- You can only have one Version 8 DAS on each machine (although a V7 and V8 DAS can co−exist on one machine).
- Just like an instance, each DAS must be created under a user ID. It does not matter whether a different user ID is used for each DAS in the environment, or whether the same user ID is used and that the user ID's home directory is not shared.
- If the same user ID is to be used on each machine, then that user ID's home directory cannot be shared with the other machines.
- If a different user ID is used for each DAS, then the home directories of the user IDs that are used can be shared.
- If an existing user is used as the DAS user, this user must also exist on all the participating computers

before installation.

- For response file installations: If your response file specifies to create a new DAS user on each machine in the cluster, and that user already exists on any of the participating computers, then that user must have the same primary group as the new DAS user.

3.3.4. Adding existing users as DB2 users

To enable existing system users to act as DB2 administrators, you must make a few modifications to the user accounts:

- Add your user to the db2grp1 group. The database manager SYSADM_GROUP configuration parameter controls who has that level of access for all databases in the that instance. You can determine what group is set as the SYSADM group by entering:
  ```
  db2 get dbm cfg | grep SYSADM_GROUP
  ```
  By default, DB2 sets this as the db2grp1 group. To add the user to the SYSADM group, modify the /etc/groups file (as root), adding the user account ID to the line for the DB2 system admin group. For example:
  ```
  db2grp1:x:102:jackc,ellenp,frankj
  ```
- Update the user's profile by adding the following lines to the .bashrc file in the home directory
  ```
  # The following three lines have been added by the DB2 SYSADM.
  if [ -f /home/db2inst1/sqllib/db2profile ]; then
      . /home/db2inst1/sqllib/db2profile
  fi
  ```
- If you want to add desktop icons for the user account, run the command:
  ```
  /opt/IBM/db2/V8.1/bin/db2icons username
  ```
  where username is replaced by the user account you wish to enable.
- After installation, if you wish to grant or revoke access to databases (or other objects), use the "User and Groups" section of the Control Center or the GRANT and REVOKE SQL statements.) For example, by default, the sample database grants a number of privileges to "PUBLIC".

4. Installing DB2 UDB Version 8.2

4.1. db2setup installation

Prerequisites.

- Ensure that your system meets the installation, memory, and disk requirements outlined in Section 3.
- Review the installation considerations for IBM Developer Kit for Java 1.4.1.
- You require root authority to perform the installation.
- The DB2 UDB Setup wizard is a graphical installer. You must have X Window System software capable of rendering a graphical user interface for the DB2 UDB Setup wizard to run on your machine.
- For 2.6 kernels, ensure that Asynchronous I/O (AIO) has been installed. To use AIO you must install libaio−0.3.96 or later, and have a kernel that supports AIO. AIO can be enabled and disabled at run time by issuing the db2set command. To enable AIO for DB2 UDB after installation, execute the command **db2set DB2LINUXAIO=true** and restart DB2 UDB.

  For additional 2.6 kernel performance enhancements, see *DB2 UDB and the 2.6 kernel*.

Procedure. To install DB2 UDB ESE (single−partition):

1. Log on to the system as a user with root authority.
2. Enter the command to mount the DB2 UDB installation CD:
   ```
   mount /mnt/cdrom
   ```
3. Change to the `/mnt/cdrom` directory.
4. Run the **db2setup** command.
5. When the IBM DB2 UDB Launchpad opens, select Install Products.
6. Select the DB2 UDB Enterprise Server Edition button on the next window. Click Next. The DB2 Information Management Software splash screen will appear as the DB2 UDB installer prepares your system for the installation of DB2 UDB.
7. When the Welcome to the DB2 Setup Wizard window appears, click Next.
8. On the Software License Agreement window, read the license agreement and select the Accept button if you accept the terms (the Decline button is selected by default). Click Next.
9. On the Select the installation type window, select the kind of installation you want to perform. You have three installation options:

   - Typical: 370−480 MB. This option installs most of the DB2 UDB components, using a typical configuration with default values. This option includes all of the required features, ODBC support, and commonly used tools such as the Control Center. This option does not install the Application Development and Business Intelligence tools. If you want to install these tools, then you should choose the Custom installation type. You can select the View Features button to see what components get installed with this option.
   - Compact: 320−390 MB. This option installs only the basic DB2 UDB components, with minimal configuration performed on your computer. With this option you get server support, client support, and getting started functionality. The DB2 UDB Setup wizard creates an instance and customizes it to use the communication protocols that are detected on your system. Click the View Features button for more information.
   - Custom: 320−890 MB. This option allows you to install only those components that you select, and to create and customize an instance to use communication protocols that are

detected on your system.

Click Next. (The rest of these instructions are based on the Custom installation option).

10. On the Select the installation action window, select the Install DB2 UDB Enterprise Server Edition on this computer check box. You can also choose to save your installations settings in a response file by selecting the Save your settings in a response file check box. You should choose this option if you plan to install DB2 UDB, with these same settings, on other computers in your environment.
11. Select the features you want installed on the Select the features to install window. You can preview each component by selecting it and reading the corresponding description in the Description section on the window. Once you have decided what features you want to install, click Next.
12. Select what languages you want installed on your computer on the Languages window. English is selected by default in the Selected languages section of the window. Click Next.
13. In the Specify the location of the DB2 Information Center window, select how you want to access the DB2 UDB Information Center. The DB2 UDB Information Center contains documentation for DB2 UDB and other related DB2 products. The On the IBM Web site button is selected by default. By choosing this option DB2 UDB is configured to access documentation on the Web using your browser. However, if you want to install a copy of the DB2 UDB Information Center on your local computer, then select the Install the DB2 Information Center separately after this DB2 product installation button (refer to the section on Installing the DB2 UDB Information Center in this HOWTO for more information). If you have a copy of the DB2 UDB Information Center installed elsewhere on an intranet server, you can choose the On the following intranet server button to gain access to DB2 UDB documentation. Specify the values in the Host name and Port number text fields (the default value for the port number is `51000`). Click the Help button for more information. Click Next.
14. Fill out the information for the user that will administer the Database Administration Server on the Set user information for the DB2 Administration Server window. Select either the New user button, or the Existing user button. The Password and Confirm password text boxes are outlined in red indicating required fields that must be filled in for the installation to continue. Click Next.
15. If you want to create a new DB2 UDB instance, then select the Create a DB2 instance button on the Set up a DB2 instance window. Otherwise, select the Do not create a DB2 instance button. If you select this option, you will need to create a new instance after the installation by using the **db2icrt** command. Click Help for more information. Click Next.
16. If your installation is a single−partition installation, as in this example, select the Single−partition instance button on the Select how the instance will be used window. Choosing this option means that this instance will only reside on your local computer and will not be used in a partitioned environment. Click Next.
17. Fill out the information for the DB2 UDB instance owner on the Set user information for the DB2 instance owner window. Select either the New user button, or the Existing user button. The Password and Confirm password text boxes are outlined in red indicating required fields that must be filled in for the installation to continue. Click Next.
18. Fill out the information for the fenced user on the Set user information for the fenced user window. Select either the New user button, or the Existing user button. The Password and Confirm password text boxes are outlined in red indicating required fields that must be filled in for the installation to continue. For more information on fenced users, and how fenced users relate to fenced and non−fenced user−defined functions (UDFs) and stored procedures, click Help. Click Next.
19. Configure the DB2 UDB instance for TCP/IP communications on the Configure DB2 instance TCP/IP communication window. Either accept the default values found in the Service name and Port number fields, or select your own. If you choose to configure the instance at a later time, select the Do not configure at this time button. Click Next.
20. In the Set instance properties window, select an instance authentication type that will apply to all of the databases that are owned by the instance. There are three authentication types:

- ♦ Server: this type specifies that authentication occurs on the server using the local operating system security.
- ♦ Server Encrypt: this type specifies that the server accepts encrypted server authentication schemes.
- ♦ Client: this type specifies that authentication occurs on the database partition where the application is invoked using operating system security.

Select the Autostart the instance of system startup button if you want the instance to start when you start up your system. Click Next.

21. On the Prepare the DB2 tools catalog window select the Use a local database button if you plan on using the Task Center or scheduler. This option creates a database on your local computer that stores task metadata. If you do not have this repository, the scheduler will not function. If you want to create a DB2 UDB tools catalog after you finish the installation, select the Do not prepare the DB2 tools catalog on this computer button. You can create the DB2 UDB tools catalog manually by using the **CREATE TOOLS CATALOG** command. Click Next.
22. If you decide to create a DB2 UDB tools catalog, you can specify the required information for the local database on the Specify a local database to store the DB2 tools catalog window. You can specify which instance, which database, and which schema in which to store the DB2 UDB tools catalog. The DB2 UDB tools catalog will be placed in the home directory of the instance owner. Click Next.
23. Specify the administration contact list on the Set up the administration contact list window. This list stores administrator contact information that is used for notifying administrators if and when a database requires attention. Either create a new contact list that is stored locally, or use an existing global contact list that resides on a remote DB2 UDB server. If you check the Enable notification button, your system will search for an available SMTP server and set it to be used for e−mail notifications. Click Next.
24. On the Specify a contact for health monitor notification window select the health monitor that will run on the DB2 UDB instance you are setting up. The DB2 UDB health monitor will send a notification e−mail to the person you specify in the Name and Email address fields, if and when a health indicator threshold is breached. You can choose to have this notified by a pager by selecting the Address is for a pager check box. Select the Defer this task until after installation is complete button if you want to specify health notification contacts later on. You can use the Task Center or the Control Center to create these contacts. Click Next.
25. Click Finish on the Start copying files window to complete the installation. The Current settings section provides a summary of your installation and configuration settings.

While the files are being copied to your computer you will see two status bars on the Installing DB2 UDB Enterprise Server Edition window, indicating the installation task progress and the overall installation progress. When the installation is finished the Setup Complete window indicates what post−installation steps need to be taken, along with a status report of the installation. Click Finish to exit the DB2 UDB Setup wizard.

5. Installing DB2 Version 8.2 on OpenSuse 10.0

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM officially supports this distribution for DB2 UDB V8.2 only for application development or client purposes.

5.1. Installation

Installing with the DB2 Express Discovery Disc, the installation went smoothly.

5.2. Post−installation notes

If you have enabled your firewall (the **iptables** service), and you want to enable external connectivity to your DB2 server, you must open up ports 553 and **db2 get dbm cfg | grep SVCENAME** in your firewall.

6. Installing DB2 Version 8.2 on Red Hat Enterprise Linux (RHEL) 3

Before proceeding, please ensure that you have read Section 3 and Section 4.

These installation notes are based on a single−partition installation using the 2.4.21−20 kernel (September 20, 2004).

6.1. Pre−installation notes

The pdksh package required for a multiple−partition environment setup is installed by default on Red Hat Enterprise Linux 3.

RHEL 3 has backported the asynchronous I/O kernel capabilities of the 2.6 kernel into their 2.4 kernel, but you will need to download and install the libaio RPM package to take advantage of this.

To check for the pdksh or libaio RPM packages, run the commands:

```
rpm −qi pdksh
rpm −qi libaio
```

If either package needs to be installed, you can add it to your system with the Red Hat up2date utility (this requires registration for a Red Hat Network account). Just run:

```
up2date pdksh
up2date libaio
```

No Java package is installed by default, so there are no conflicts with the Java 2 SDK from IBM that DB2 UDB installs.

6.2. Installation notes

The installation on Red Hat Enterprise Linux 3 completed successfully, with no problems getting the GUI tools or installer working, and no errors during the install. This smooth installation isn't surprising, given that RHEL 3 is an officially supported platform for Version 8.2

6.3. Post−installation notes

Initialize the Asynchronous I/O so DB2 UDB can take advantage of this new kernel feature. As the instance owner:

```
db2set DB2LINUXAIO=TRUE
```

7. Installing DB2 Version 8.2 on Red Hat Enterprise Linux (RHEL) 4

Before proceeding, please ensure that you have read Section 3 and Section 4.

These installation notes are based on a single−partition installation using the 2.6.9−5.0.5 EL kernel (April 20, 2005).

7.1. Pre−installation notes

The pdksh package required for a multiple−partition environment setup is installed by default on Red Hat Enterprise Linux 4.

The java−1.4.2−gcj−compat Java package is installed by default, but there are no conflicts with the Java 2 SDK from IBM that DB2 UDB installs.

7.2. Installation notes

The installation on Red Hat Enterprise Linux 4 completed successfully, with no problems getting the GUI installer or tools working, and no errors during the install. This smooth installation isn't surprising, given that RHEL 4 is an officially supported platform for Version 8.2

7.3. Post−installation notes

Initialize the Asynchronous I/O so DB2 UDB can take advantage of this new kernel feature. As the instance owner:

```
db2set DB2LINUXAIO=TRUE
```

8. Installing DB2 Version 8.2 on SuSE Linux Enterprise (SLES) 9

Before proceeding, please ensure that you have read Section 3 and Section 4.

SuSE Linux Enterprise Server 9 was released in August 2004 to great anticipation. It is currently the only Enterprise level Linux distribution certified by DB2 UDB that runs on the 2.6 kernel by default.

These notes are based on the 2.6.5−7 kernel (the SuSE Linux kernel update as of September 10, 2004), and installing a single−partition environment.

8.1. Pre−installation notes

1. The pdksh package is installed in a default setup (required for multiple−partition and high availability installations)
2. The IBMJava2−JRE−1.4.2 rpm package is also installed by default
3. The kernel parameters (ipcs −l) before DB2 UDB installation:
```
        max number of segments:         4096
        max segsize:                    32768
        max number of arrays:           128
        max queues system wide:         16
        default max size of queue:      16384
```
4. Asynchronous I/O package (libaio) is also installed by default.

8.2. Installation Notes

I received an error during start up of the db2setup utility:

```
/db2/ese/db2/linux/install/db2jinst:
line 131: 3474 Segmentation fault
$JAVA_PATH/$JAVA_INTERPRETER $JAVA_OPTIONS \
    −cp $JAVA_CLASSPATH $DB2SetupRun "$@" 2>/tmp/db2setup.err.running
```

The solution here is to run

```
export LD_ASSUME_KERNEL="2.4.19"
```

from the terminal where you launched the **db2setup** utility, then run **db2setup** again. This change will only affect the Java installer, not the kernel performance.

8.3. Post−installation notes

The kernel parameters (ipcs −l) after installation look like:

```
        max number of segments:         4096
        max segsize:                    262144 (updated)
        max number of arrays:           1024 (updated)
        max queues system wide:         1024 (updated)
        default max size of queue:      16384
```

Initialize the Asynchronous I/O so DB2 UDB can take advantage of this new kernel feature. As the instance owner:

```
db2set DB2LINUXAIO=TRUE
```

Optional: update the DB2 UDB JDK_PATH parameters so DB2 UDB uses SLES's installed Java 2 package. As the instance owner:

```
db2 UPDATE DBM CFG USING JDK_PATH /opt/IBMJava2-142
                db2 UPDATE ADMIN CFG USING JDK_PATH /opt/IBMJava2-142
```

Surprisingly, my installation didn't have Mozilla installed by default. To have the DB2 UDB help display properly, Mozilla should be installed using YaST.

---

9. Installing DB2 Version 8.2 on SuSE Linux Enterprise (SLES) 8

Before proceeding, please ensure that you have read Section 3 and Section 4.

These installation notes are based on the 2.4.21 kernel (SLES kernel on September 20, 2004)

9.1. Pre−installation notes

The pdksh package is installed in the default setup (required for multiple−partition and high availability installations).

Also, SLES 8 installs the IBMJava2−SDK−1.3.1 rpm package in a typical installation.

9.2. Installation notes

SuSE Linux Professional installs its own 1.3.1 Java JRE by default (the "IBMJava2−SDK−1.3.1" rpm package). During the DB2 UDB installation, this will cause a warning to be generated:

```
Command to be run: "/bin/rpm -ivh \
    '/mnt/cdrom/db2/linux/Java-1.4/IBMJava2-SDK-1.4.1-2.0'.i386.rpm"
Warning: there may be a version of this package already installed.
```

Since this package was supplied by SuSE Linux, it will have been installed into a different directory tree than the JDK from IBM that DB2 UDB installs. Therefore, there are 3 options you can choose to deal with this issue:

1. Uninstall the Java2−1.3.1 package before installing, using either the YaST tool, or with the command line:
   ```
   rpm −e IBMJava2-SDK-1.3.1
   ```
   After installation, this will leave the IBMJava2−SDK−1.4.1 package as the only Java installation on the system.
2. Leave both Java packages installed. By default, DB2 UDB will use the 1.4.1 JRE that it installs. The fact that SuSE Linux's Java package installs to a different directory means that the two packages will not conflict.
3. Remove the IBM package (either with YaST or with rpm −e IBMJava2−SDK−1.4.1) after the installation completes. Update the DB2 UDB JDK_PATH configuration parameters to use the package installed by SuSE Linux:
   ```
   db2 UPDATE DBM CFG USING JDK_PATH /usr/lib/java
   db2 UPDATE ADMIN CFG USING JDK_PATH /usr/lib/java
   ```

Other than that one minor issue, DB2 UDB installed cleanly and setup was problem free.

10. Installing DB2 Version 8.2 on SuSE Linux Professional 9.1

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM does not officially support this distribution for DB2 UDB V8.2. Any problems you encounter using this distribution must be replicated on a supported distribution before IBM Support will help you resolve the issue.

These notes are based on a single−partition installation on the 2.4.21 kernel (September 20, 2004).

10.1. Pre−installation notes

If you want to run a multiple−partition environment, you will need to install the pdksh package found on CD 4.

10.2. Installation notes

SuSE Linux Professional installs its own 1.4.2 Java JRE by default (the "java2−jre−1.4.2" package). During the DB2 UDB installation, this will cause a warning to be generated:

```
Command to be run: "/bin/rpm -ivh \
    '/mnt/cdrom/db2/linux/Java-1.4/IBMJava2-SDK-1.4.1-2.0'.i386.rpm"
Warning: there may be a version of this package already installed.
```

Since this package was supplied by SuSE Linux, it will have been installed into a different directory tree than the JDK from IBM that DB2 UDB installs. Therefore, there are 3 options you can choose to deal with this issue:

1. Uninstall the java2−jre−1.4.2 package before installing, using either the YaST tool, or with command line **rpm erase** command. This will leave the IBMJava2−SDK package as the only Java installation on the system.
2. Leave both Java packages installed. By default, DB2 UDB will use the 1.4.1 JRE that it installs. The fact that SuSE Linux's Java package installs to a different directory means that the two packages will not conflict.
3. Remove the IBM package (either with YaST or with rpm −e IBMJava2−SDK) after the installation completes. Update the DB2 UDB JDK_PATH configuration parameters to use the package installed by SuSE Linux:
   ```
   db2 UPDATE DBM CFG USING JDK_PATH /usr/lib/java
   db2 UPDATE ADMIN CFG USING JDK_PATH /usr/lib/java
   ```

Other than that one minor issue, DB2 UDB installed cleanly and setup was problem free.

11. Installing DB2 Version 8.2 on SuSE Linux Personal 9.1

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM does not officially support this distribution for DB2 UDB V8.2. Any problems you encounter using this distribution must be replicated on a supported distribution before IBM Support will help you resolve the issue.

These notes are based on the 2.6.4−52−default kernel and installing to a single−partition environment.

11.1. Pre−installation notes

The pdksh package is not installed by default. If you want to install a multiple−partition environment on SuSE Linux Personal, you will have to install pdksh.

The installer may complain about being unable to use the SMTP server even though **postfix** is installed and **sendmail** is available in the path. Make sure that **postfix** has been started before installation.

11.2. Installation notes

The usual trick of adding an xhost entry for localhost and then exporting the DISPLAY=localhost:0 when installing (so that the graphical utilities are displayed on a non−root user's desktop) didn't seem to work here. Instead, execute the following command to log in as root and forward the X server:

```
ssh −X root@localhost
```

12. Installing DB2 Version 8.2 on Fedora Core 2

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM does not officially support this distribution for DB2 UDB V8.2. Any problems you encounter using this distribution must be replicated on a supported distribution before IBM Support will help you resolve the issue.

Fedora Core 2 is a release of the Fedora Project (the Red−Hat−sponsored and community−supported open source project). The Fedora Project is a source for new technologies and enhancements that may be incorporated into Red Hat Enterprise Linux (RHEL) in the future.

These notes are based on both the default 2.6.5−1 and the update 2.6.8−1 kernels (the Fedora Core 2 kernel update on September 13, 2004 was 2.6.8−1), and installing to a single−partition environment.

12.1. Pre−installation notes

The kernel parameters (ipcs −l) before DB2 UDB installation:

```
        max number of segments:             4096
        max segsize:                        32768
        max number of arrays:               128
        max queues system wide:             16
        default max size of queue:          16384
```

The asynchronous I/O package (libaio) is also required to take advantage of the 2.6 kernel improvements. Install the libaio package from the Fedora Core CD.

12.2. Install Notes

The default 2.6.5 kernel had no installation issues with the db2setup GUI installer.

On the 2.6.8 kernel, however, I received an error when running the db2setup utility (the error may not display to the terminal; instead, if the db2setup command seems to fail without error, check the /tmp/db2setup.err file) :

```
JVMDG080: Cannot find class com/ibm/jvm/Trace
JVMXM012: Error occurred in diagnostics initialization(2)
Could not create the Java virtual machine.
```

The workaround for this error is to install Sun's Java 2 Platform Standard Edition 5.0 RC (Fedora Core 2 doesn't install a Java package by default). The rpm file is available from http://java.sun.com/j2se/1.5.0/download.jsp. After installing Java, set the environment variables in the terminal where you execute the **db2setup** utility:

```
export DB2USELOCALJRE=TRUE
export JAVA_HOME=/usr/java/jre1.5.0/
```

/usr/java/jre1.5.0/ is the location where Java is installed by default in the rpm from Sun.

The last step to making this work − and this is very important − is to add in a link named "jre" in Sun's Java folder so that db2setup works properly with the directory:

```
cd /usr/java/jre1.5.0/
mkdir jre
cd jre
ln -s ../bin bin
```

After making these changes, the **db2setup** GUI installer works properly.

Alternatively, you can avoid the Java−based problems by installing on Fedora Core 2 using

```
./db2setup -r db2ese.rsp
```

where db2ese.rsp is a slightly modified response file (the sample response files are found in the `db2/linux/samples/` directory)

12.3. Post−installation notes

The kernel parameters (ipcs −l) after installation look like:

```
        max number of segments:         4096
        max segsize:                    262144 (updated)
        max number of arrays:           1024 (updated)
        max queues system wide:         1024 (updated)
        default max size of queue:      16384
```

Initialize the Asynchronous I/O so DB2 UDB can take advantage of this new kernel feature. As the instance owner, run:

```
db2set DB2LINUXAIO=TRUE
```

For the GUI tools (DB2 UDB Control Center, DB2 UDB Configuration Assistant, DB2 UDB First Steps) to work properly under the 2.6.8−1 kernel, you'll need to tell DB2 UDB to use the Sun JRE. As the instance owner, execute these commands:

```
db2 UPDATE DBM CFG USING JDK_PATH /usr/java/jre1.5.0/
db2 UPDATE ADMIN CFG USING JDK_PATH /usr/java/jre1.5.0/
```

With that update, the DB2 UDB Java−based tools will work properly.

13. Installing DB2 Version 8.2 on Fedora Core 4

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM officially supports this distribution for DB2 UDB V8.2 only for application development or client purposes.

13.1. Pre−installation notes

Before installing DB2 UDB V8.2 on Fedora Core 4, you must install the following compatibility libraries:

- `compat−libgcc`
- `compat−libstdc++`

13.2. Installation

Installing with the DB2 Express Discovery Disc, the installation went smoothly.

13.3. Post−installation notes

To inherit the DB2 instance environment from another home directory, create databases, dynamically load the DB2 libraries, and other tasks, you must disable SELinux or modify the default SELinux policies. You can disable SELinux at run−time by running the following command as root:

```
# /usr/sbin/setenforce 0
```

Note that when you reboot SELinux will be enabled again. To disable SELinux permanently, edit `/etc/selinux/config` and set the following value:

```
SELINUX=disabled
```

If you have enabled your firewall (the **iptables** service), and you want to enable external connectivity to your DB2 server, you must open up ports 553 and **db2 get dbm cfg | grep SVCENAME** in your firewall.

14. Installing DB2 Version 8.2 on Mandriva Linux 10.1 (Community)

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM does not officially support this distribution for DB2 UDB V8.2. Any problems you encounter using this distribution must be replicated on a supported distribution before IBM Support will help you resolve the issue.

These notes are based on the 2.6.8.1−16mdk kernel and installing to a single−partition environment.

14.1. Pre−installation notes

The pdksh package is not installed by default. If you want to install a multiple−partition environment on Mandriva Linux, you will have to install it using **drakconf** (it's on CD3). Also, no Java package is installed by default, so we don't have to worry about any conflicts arising from that.

14.2. Installation notes

The GUI based installation went smoothly, with no reported problems or issues. If you do experience a hang with the Java installer, try the Sun Java workaround described in the Fedora Core 2 section.

14.3. Post−installation notes

Be careful of the Mandriva Linux security **msec** utility − it can prevent DB2 from functioning correctly by changing the permissions on directories like `/home/db2inst1/sqllib/` that must be world−readable for DB2 UDB to non−world readable. If you find that DB2 UDB is not working properly, investigate your security settings.

15. Installing DB2 Version 8.2 on Gentoo (2005−01−17)

Before proceeding, please ensure that you have read Section 3 and Section 4.

Note that IBM does not officially support this distribution for DB2 UDB V8.2. Any problems you encounter using this distribution must be replicated on a supported distribution before IBM Support will help you resolve the issue.

These notes are based on a single−partition installation on the 2.6.9−gentoo−r13 kernel (2005−01−17). Following the manual install instructions in the DB2 Information Center, we use the RPM method of installing DB2 UDB V8.2, through a modified **db2_install** script. This is just one of many ways you could accomplish this task, but this method has been tested (at least once) and it works.

15.1. Pre−installation notes

Emerge the rpm package and any of its pre−requisites:

```
emerge rpm
```

If you want to run a multiple−partition environment, you will need to emerge either the pdksh or ksh package.

15.2. Installation notes

Gentoo uses the latest version of the GNU coreutils package of utilities, which includes the **tail** command. The DB2 install script (**db2_install**) and several instance administration shell scripts rely on deprecated command line arguments to **tail** that have been removed entirely from the latest GNU `coreutils` version of **tail**. Therefore, you must copy the contents of the install CD and modify the **db2_install** script before you can successfully run the install.

The RPM packages supplied with the DB2 install CD contain dependency listings that cannot be satisfied inside a Gentoo environment, because Gentoo does not use RPM as a native packaging method. To overcome these dependencies, you also need to modify the **db2_install** script to ignore dependencies and force the installation of DB2 onto your system.

Perform the following steps to modify the **db2_install** script for the required **tail** and **rpm** changes and install the product:

1. (Skip this step if you downloaded a tarball containing the DB2 install code.) Assuming you want to copy the contents of the DB2 install CD to a subdirectory named "cdrom" in your home directory, issue the following command:
   ```
   cp −r /mnt/cdrom/* ~/cdrom
   ```
2. Open the `~/cdrom/db2/linux/db2_install` script in your text editor and modify the following lines

   From:

   ```
   output_df="`df -k ${dirname2?} | tail -1`"
   ```
   To:

```
output_df="`df -k ${dirname2?} | tail -n -1`"
```

From:

```
( rpm -ivh  ${INSTDEV?}/${pkg?}${RPMEXTN?} 2>&1; \
   echo $? > ${TMPRC?} ) | tee -a ${DB2INSTLOG?}
```

To:

```
( rpm --force --nodeps -ivh  ${INSTDEV?}/${pkg?}${RPMEXTN?} 2>&1; \
   echo $? > ${TMPRC?} ) | tee -a ${DB2INSTLOG?}
```

3. Change to the `~/cdrom` directory and issue the following command as root:

```
./db2_install
```

4. The DB2 installer prompts you to enter the keyword of the product you want to install. for example, issuing the **db2_inst** command for the DB2 UDB Express Edition for Linux install displays the following screen:

```
Specify one or more of the following keywords,
separated by spaces, to install DB2 products.

 Keyword        Product Description
 DB2.EXP        DB2 UDB Express Edition for LINUX

Enter "help" to redisplay product names.

Enter "quit" to exit.

***********************************************************
DB2.EXP
```

The DB2 installer will now install many RPM packages into the /opt/IBM/db2/V8.1/ directory.

15.2.1. Installing DB2 UDB FixPaks on a non−RPM distribution

When installing a DB2 FixPak onto an existing DB2 installation, the RPM utility is used, which will cause problems on a distribution that does not use the RPM installation method by default. Since Gentoo is one of those distributions, you must make the following changes to the **installFixPak** script.

In the **installFixPak** script, change this line:

```
                echo "Updating to ${fullfsname?} ... "
                rpm -ivh ${fullfsname?}
```

to read:

```
                echo "Updating to ${fullfsname?} ... "
                rpm --force --nodeps --ivh ${fullfsname?}
```

15.3. Post−installation notes

Now that you have installed DB2, you need to create the groups, users, DB2 Administration Server instance, and database instance for your database. The following instructions create users and groups with the default names used in the DB2 documentation (dasadm1, db2inst1, db2fenc1).

1. As root, create the required groups:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```

2. As root, create the required users and assign them to their corresponding groups:

```
useradd -u 1004 -g db2iadm1 -m -d /home/db2inst1 db2inst1
useradd -u 1003 -g db2fadm1 -m -d /home/db2fenc1 db2fenc1
useradd -u 1002 -g dasadm1 -m -d /home/dasusr1 dasusr1
```

3. Before you can create the DB2 Administration Server and database instance, you have to modify the db2iutil library to update calls to the **tail** command to the current GNU syntax. As root, open the `/opt/IBM/db2/V8.1/instance/db2iutil` file in an editor and make the following changes (or apply the changes to the file using the patch utility):

```
--- db2iutil.orig       2005-01-17 07:05:58.000000000 -0500
+++ db2iutil    2005-01-17 07:06:32.000000000 -0500
@@ -327,7 +327,7 @@
     fi

     # Get output of the "df" command
-    output_df="`df -k  ${dirname2?} | tail -1`"
+    output_df="`df -k  ${dirname2?} | tail -n 1`"
     # On some platforms, the filesystem is on the previous line causing us
     # one less token for awk to find.
     free_space_in_fs=`echo ${output_df?} | awk '{if ($4 !~ /%/) {print $4} else
 {print $3}}'`
@@ -382,7 +382,7 @@
     rm -f ${TMPFILE3?}

     # Get the name of the filesystem where dir $dirname2 resides.
-    df -k ${dirname2?} | tail +2 > ${TMPFILE3?}
+    df -k ${dirname2?} | tail -n +2 > ${TMPFILE3?}

     # There must be only one line in TMPFILE3 file
     lcount=`wc -l ${TMPFILE3?} | awk '{print $1}'`
@@ -394,7 +394,7 @@
             chk_fsystype 22
         stop_prog 1
     fi
-    fsname=`awk '{print $NF}' ${TMPFILE3?} | tail -1`
+    fsname=`awk '{print $NF}' ${TMPFILE3?} | tail -n 1`

     rm -f ${TMPFILE3?}
     foundit=${FALSE?}
@@ -879,7 +879,7 @@
         ${DB2VER_V2?})
             # Dir of DB2 V2 instances
             if [ -f ${DB2V2ILIST?} ]; then
-                tail +2 ${DB2V2ILIST?} > ${TMPFILE3?}
+                tail -n +2 ${DB2V2ILIST?} > ${TMPFILE3?}
                 if [ -s ${TMPFILE3?} ]; then
                     for iname in `cat ${TMPFILE3?}`; do
                         db2ilist=${db2ilist?}" ${iname?}"
@@ -891,7 +891,7 @@
         ${DB2VER_DJ?})
             # Dir of DB2 V2 instances
             if [ -f ${DJV211ILIST?} ]; then
-                tail +2 ${DJV211ILIST?} > ${TMPFILE3?}
+                tail -n +2 ${DJV211ILIST?} > ${TMPFILE3?}
                 if [ -s ${TMPFILE3?} ]; then
                     for iname in `cat ${TMPFILE3?}`; do
                         db2ilist=${db2ilist?}" ${iname?}"
```

4. Create the DB2 Administration Server:

```
/opt/IBM/db2/V8.1/instance/dascrt -u dasusr1
```

Create the DB2 database instance:

```
/opt/IBM/db2/V8.1/instance/db2icrt -a server -u db2fenc1 db2inst1
```

15.3.1. Enabling remote connections

When you perform a manual install of DB2, the installer does not automatically set up the communication protocols for your DB2 server. To enable connections to your DB2 server from remote clients, perform the following steps:

1. Set the port on which DB2 should communicate. Ensure that you select a port that is not blocked by a firewall or used by another service defined in the `/etc/services` file. To set the communications port, update the DB2 database manager configuration variable SVCENAME. For example, to set the communications port to 50055, issue the following command as the instance user ("db2inst1", if you've been using the defaults):
   ```
   db2 update dbm cfg using svcename 50055
   ```
2. Set the DB2 communications protocol registry variable to define the protocol supported by the server. Typically the only protocol you would use is TCP/IP, so issue the following command as the instance user ("db2inst1" if you've been using the defaults):
   ```
   db2set DB2COMM=tcpip
   ```
3. Restart the database instance to enable the settings to take effect.
   ```
   db2stop
   db2start
   ```

You should now be able to catalog and connect to your Gentoo server from a remote client.

15.3.2. Running the DB2 Control Center and other DB2 GUI tools

To run the DB2 Control Center or other DB2 GUI tools, you must install the IBM Software Developers' Kit for Java. This RPM is included in the DB2 install CD–ROM in the `/db2/linux/Java-1.4/` directory. To install the IBM Software Developers' Kit for Java, change to the root directory of the CD–ROM and issue the following command as root:

```
rpm --force --nodeps -ivh db2/linux/Java-1.4/IBMJava2-SDK-1.4.1-2.0.i386.rpm
```

Log in as the db2inst1 user and invoke the DB2 Control Center:

```
db2cc
```

16. Installing DB2 Version 8.2 on Ubuntu 5.04 'Hoary Hedgehog'

Please follow the instructions in Section 18, which were based on an Ubuntu 5.04 install.

17. Installing DB2 Version 8.2 on Ubuntu 5.10 'The Breezy Badger'

Please follow the instructions in Section 18, which were based on an Ubuntu 5.04 install and confirmed with an Ubuntu 5.10 'default' installation.

18. Installing DB2 Version 8.2 on systems that use .deb packages

Before proceeding, please ensure that you have read Section 3 and Section 4.

These notes are based on a single−partition installation on Ubuntu 5.04. As Ubuntu uses `.deb` packages, we use the **alien** package conversion utility to convert and install the DB2 RPM packages directly.

18.1. Pre−installation notes

Ensure that you have installed **alien**:

```
sudo apt−get install alien
```

If you want to run a multiple−partition environment, you will need to install the pdksh package.

18.2. Installation notes

The RPM packages supplied with the DB2 install CD contain dependency listings that cannot be satisfied inside a Ubuntu environment, because Ubuntu does not use RPM as a native packaging method. To overcome these dependencies, you need to modify the **db2_install** script to ignore dependencies and force the installation of DB2 onto your system using **alien**.

Perform the following steps to modify the **db2_install** script and install the product:

1. Open the ~/db2/linux/db2_install script in your text editor and modify the following lines

   From:

   ```
   ( rpm −ivh  ${INSTDEV?}/${pkg?}${RPMEXTN?} 2>&1; \
      echo $? > ${TMPRC?} ) | tee −a ${DB2INSTLOG?}
   ```

   To:

   ```
   ( alien −dic  ${INSTDEV?}/${pkg?}${RPMEXTN?} 2>&1; \
      echo $? > ${TMPRC?} ) | tee −a ${DB2INSTLOG?}
   ```
2. Issue the following command to begin installing DB2:

   ```
   sudo ./db2_install
   ```
3. The DB2 installer prompts you to enter the keyword of the product you want to install. for example, issuing the **db2_install** command for the DB2 UDB Express Edition for Linux install displays the following screen:

   ```
   Specify one or more of the following keywords,
   separated by spaces, to install DB2 products.

    Keyword        Product Description
    DB2.EXP        DB2 UDB Express Edition for LINUX

   Enter "help" to redisplay product names.

   Enter "quit" to exit.

   ***********************************************************
   ```

```
DB2.EXP
```

The DB2 installer will now install many RPM packages into the `/opt/IBM/db2/V8.1/` directory.

18.3. Post−installation notes

Now that you have installed DB2, you need to create the groups, users, DB2 Administration Server instance, and database instance for your database. The following instructions create users and groups with the default names used in the DB2 documentation (dasadm1, db2inst1, db2fenc1).

1. Create the required groups:
   ```
   sudo groupadd −g 999 db2iadm1
   sudo groupadd −g 998 db2fadm1
   sudo groupadd −g 997 dasadm1
   ```
2. Create the required users and assign them to their corresponding groups:
   ```
   sudo useradd −u 1004 −g db2iadm1 −m −d /home/db2inst1 db2inst1
   sudo useradd −u 1003 −g db2fadm1 −m −d /home/db2fenc1 db2fenc1
   sudo useradd −u 1002 −g dasadm1 −m −d /home/dasusr1 dasusr1
   ```
3. Create the DB2 Administration Server:
   ```
   sudo /opt/IBM/db2/V8.1/instance/dascrt −u dasusr1
   ```
   Create the DB2 database instance:

   ```
   sudo /opt/IBM/db2/V8.1/instance/db2icrt −a server −u db2fenc1 db2inst1
   ```

18.3.1. Enabling remote connections

When you perform a manual install of DB2, the installer does not automatically set up the communication protocols for your DB2 server. To enable connections to your DB2 server from remote clients, perform the following steps:

1. Set the port on which DB2 should communicate. Ensure that you select a port that is not blocked by a firewall or used by another service defined in the `/etc/services` file. To set the communications port, update the DB2 database manager configuration variable SVCENAME. For example, to set the communications port to 50055, issue the following command as the instance user ("db2inst1", if you've been using the defaults):
   ```
   db2 update dbm cfg using svcename 50055
   ```
2. Set the DB2 communications protocol registry variable to define the protocol supported by the server. Typically the only protocol you would use is TCP/IP, so issue the following command as the instance user ("db2inst1" if you've been using the defaults):
   ```
   db2set DB2COMM=tcpip
   ```
3. Restart the database instance to enable the settings to take effect.
   ```
   db2stop
   db2start
   ```

You should now be able to catalog and connect to your DB2 server from a remote client.

18.3.2. Running the DB2 Control Center and other DB2 GUI tools

To run the DB2 Control Center or other DB2 GUI tools, you must install the IBM Software Developers' Kit for Java. This RPM is included in the DB2 install CD−ROM in the `/db2/linux/Java−1.4/` directory. To install the IBM Software Developers' Kit for Java, change to the root directory of the CD−ROM and issue the following command:

```
sudo alien −dic db2/linux/Java−1.4/IBMJava2−SDK−1.4.1−2.0.i386.rpm
```

Log in as the db2inst1 user and invoke the DB2 Control Center:

```
db2cc
```

19. Installing the DB2 UDB Information Center

The new DB2 UDB Information Center is a new initiative from IBM to improve the documentation delivery system in Version 8.2. It's based on the open−source Eclipse architecture help system, and provides a better way to serve up the DB2 UDB documentation.

19.1. Installing the Information Center locally

In this example, we'll install the Information Center on Red Hat Enterprise Linux (RHEL) 3, just because the DB2 UDB installation was so smooth.

To begin, mount the CD in the drive, or expand the tarball containing the installation utilities. Since the DB2 UDB Information Center has it's own Java based installer, any Java issues that you experienced above while installing DB2 UDB will likely recur here. However, the fixes we've provided for the DB2 UDB install should work for the DB2 UDB Information Center install too.

After mounting, run the **db2setup** utility from the base directory. This will launch the Java−based setup wizard. You can also use the other setup methods we've talked about, but since the GUI installation sets up and configures everything for you, that's the method we recommend.

Installation Steps:

1. Select the "Install Products" option from the launcher, and then the "DB2 UDB Information Center" from the next panel, and click "Next".
2. Read the description and click "Next".
3. Read the license agreement, and if you agree, select "Accept" and click "Next".
4. If you want to keep your installation settings in a response file (for installing on a different machine with identical settings), check the empty box, otherwise click "Next".
5. Add any additional languages that you want to have installed. English is installed by default and cannot be deselected. Each additional language will add anywhere from 5 MB to 50 MB to the installation size. Click "Next".
6. The next panel is where things get a little interesting. Here you can specify the service name that the Information Center will run under. An entry will be added to the `/etc/services` file for this service. You probably don't need to change this, unless you have another service running as "db2icserver".

   Also, this panel is where you specify the system port that the Information Center will use to serve up the documentation. By default, this is port 51000.

   The check box on this page will change the configuration for all DB2 UDB products installed on the local system, so that their help system will use the parameters you define here. Click "Next".
7. Read the summary of what will be done, then click "Finish" to start the file installation

19.1.1. Troubleshooting

After installation, there are a couple of minor gotchas that you might run into.

First of all, there is a bug in the CD installation that won't unlock the CD mount. When you try to unmount the CD, you'll get the dreaded "device is busy" message from the OS. The solution here is to restart the

Information Center daemon with the commands:

```
cd /tmp
/etc/init.d/db2icd restart
```

After the restart is complete, you can unmount the CD cleanly. This information is covered in an ibm.com technote (http://www.ibm.com/support/docview.wss?uid=swg21179518)

Another configuration change that you can make is to change what directory the Information Center uses as a temporary working directory. To do this, modify the DB2_ECLIPSEIC_DATA parameter in the db2ic.conf file found in the `/var/db2/v81/` directory. This temporary directory can be anywhere where the "bin" user can write to.

Also in this file, you can change the port used to serve up the Information Center, in case you change your mind later on about the one you set up during the install. Just change the DB2_ECLIPSEIC_PORT parameter to the new number (though it must be under 65535).

After making any changes, do a

```
/etc/init.d/db2icd restart
```

to have your changes picked up by the Information Center (see below for more details on the db2icd daemon).

19.2. The Information Center daemon

The Information Center daemon is set up during installation and is used to launch the background process that runs the Information Center. The daemon script, named **db2icd** is installed to the `/etc/init.d/` directory, and start up symbolic links are added to the `/etc/init.d/rc.X/` run−level folders. By default, the daemon is turned on for run−levels 2, 3 and 5. You can modify these run−levels with any system init utility (**chkconfig**, **ntsysv**, or any of the X−based variants).

There are 5 commands you can pass to the daemon:

1. start   Starts the Information Center, using the configuration parameters set in the /var/db2/v81/db2ic.conf file.
2. stop   Brings down the Information Center.
3. restartor reload A combination command that stops, then starts the daemon again.
4. reload   Same as restart.
5. status   Provides information on whether the Information Center is running, and if it is running, what the process IDs (PID) are.

19.3. Accessing an installed Information Center

There are three places where you can access the DB2 UDB Information Center from your client:

1. Access the publib.boulder.ibm.com Web site (this is the default set during a client installation)
2. Access an intranet machine that is hosting the DB2 UDB Information Center within your own organization
3. Access an Information Center that is set up on your local machine

You can configure the access points during the installation (only by specifying a "Custom" installation) or afterwards either in the tools configuration panel (in the GUI tools), or from the command line.

For the GUI tools (Control Center), open the Tools menu, select Tools Settings, and then the Documentation tab. On this panel, you can specify a host URL and port where an Information Center can be found. When putting a hostname in, do not add the "http://" prefix, or any subdirectories – just type in the hostname, and the internal code will do the rest.

To change the access points from the command line, execute these commands:

```
db2set DB2_DOCHOST=host_URL
db2set DB2_DOCPORT=host_port
```

where host_URL is the location hostname of the Information Center. This could be: mydocserver.myorg.net or localhost (for a local installation), and host_port is the port number where the Information Center was installed.

By default, these DB2 UDB environment settings are blank, which means that DB2 UDB defaults to the http://publib.boulder.ibm.com/infocenter/db2help/ Web site.

To invoke the DB2 UDB Information Center from the GUI tools, click on the "?" icon in the toolbar, any of the "Help" hyperlinks or buttons, or by selecting any of the menu items in the "Help" menu.

To invoke the DB2 UDB Information Center from the command line, issue the db2icdocs or db2helpcommands.

20. Removing DB2 UDB Version 8.2

1. Remove the DAS by issuing the /opt/IBM/db2/V8.1/instance/dasdrop command as root.
2. Remove the instance by issuing the /opt/IBM/db2/V8.1/instance/db2idrop `db2 instance`command as root, where `db2 instance` refers to the actual instance you want to drop (for example, /opt/IBM/db2/V8.1/instance/db2idrop db2inst1).
3. Mount the DB2 UDB installation CD and run the /mnt/cdrom/db2_deinstall command as root. This will remove DB2 UDB from your system.
4. Remove the users that were created during the installation of DB2 UDB ESE by entering these commands as root:

   ```
   userdel -r db2inst1
   userdel -r dasusr1
   userdel -r db2fenc1
   ```

   (If you created different user names during installation, substitute as necessary.)

Optionally, you can remove IBM's version of the Java 2 package that DB2 UDB installs:

```
rpm -e IBMJava2-SDK
```

21. DB2 UDB and the 2.6 kernel

DB2 UDB Version 8.2 can take advantage of a number of performance enhancements introduced in the 2.6 kernel. There is a good paper by Rav Ahuja, Dan Behman and John Keenleyside that goes in depth on the specific enhancements of DB2 UDB V8.2 (see *Resources*); we'll cover a few of the highlights and how to implement them.

Note that these enhancements are for advanced users only. Implement each with care and be sure to thorougly test each change individually before moving on to the next enhancement.

- Asynchronous I/O handling

  Asynchronous I/O is a kernel processing enhancement that allows applications to submit I/O requests, and then continue working without having to wait for the I/O request to be completed.

  As mentioned in a few of the distribution−specific instructions, this requires the installation of the libaio package, and the setting of the DB2LINUXAIO DB2 UDB configuration parameter.
- Direct I/O caching

  Direct I/O is a different cache writing policy. Instead of using buffered I/O writes, direct I/O caching eliminates the file copying from the file cache to user buffer, thus reducing CPU utilization.

  To enable Direct I/O cache on a database, you specify the NO FILE SYSTEM CACHING parameter for the CREATE, TABLESPACE, ALTER TABLESPACE or CREATE DATABASE commands.
- Vector I/O

  Vector (or scatter/gather) I/O allows a number of I/O buffers to be used instead of one contiguous buffer, thus avoiding additional copying during reads.

  To enable Vector I/O, set the DB2_SCATTERED_IO configuration parameter to "ON".
- Large page support

  Large page support allows the application to force the kernel into using larger memory pages, thus reducing the memory overhead required to track the pages, boosting the overall available memory.

  Enabling large page support requires both a 2.6 kernel (or a backported 2.4 kernel like the one in RHEL), and installation of the libcap RPM. Turning on large page support is distribution dependent.

  For SuSE Linux Enterprise 8:

  - During boot−up, add the parameter bigpages=xxxM to the kernel boot command, where xxx is the desired large page size in MB
  - After boot, log in as root and run
    ```
    echo 1 > /proc/sys/kernel/shm-use-bigpages
    ```
    This will dynamically enable the large page support in the kernel
  - Before running db2start, execute this command to tell DB2 UDB to use the large page configuration:
    ```
    db2set DB2_LGPAGE_BP=YES
    ```

  For Red Hat Enterprise Linux 3

♦ As root, run:

```
echo XXX > /proc/sys/vm/hugetlb_pool
```

where XXX is a numerical value equivalent to the number of "huge pages" desired. (Red Hat Linux uses "huge" pages instead of large pages). Note also that the size you enter must be available in a contiguous block of memory, otherwise the large page limit will be restricted by the largest block that is available.

♦ Check the large page usage by running the command:

```
cat /proc/meminfo | grep -i huge
```

♦ Before running db2start, execute this command to tell DB2 UDB to use the large page configuration:

```
db2set DB2_LGPAGE_BP=YES
```

For Distributions using 2.6 kernel

♦ As root, run:

```
echo XXX > /proc/sys/vm/nr_hugepages
```

where XXX is a numerical value equivalent to the number of "huge pages" desired

♦ Check the large page usage by running the command:

```
cat /proc/meminfo | grep -i huge
```

♦ Before running db2start, execute this command to tell DB2 UDB to use the large page configuration:

```
db2set DB2_LGPAGE_BP=YES
```

22. Using DB2 UDB

This section gives you the basic information you need to start working with DB2 UDB on Linux. It includes instructions on issuing DB2 UDB commands and SQL statements from the command line, as well as the commands you need to start the DB2 UDB Control Center and the DB2 UDB Information Center.

22.1. DB2 UDB Control Center

The left side of the DB2 UDB Control Center provides an object−oriented view of the database objects that you have catalogued, including DB2 UDB instances and databases on other DB2 UDB servers. One way to add, edit, or drop database objects is to right−click on an object to bring up a menu.

To start the DB2 UDB Control Center

1. Ensure you are logged on to your Linux workstation using either the DB2 UDB Administration Server user ID or the DB2 UDB instance user ID. If you use the su command to become the user, ensure you include the −l parameter to initialize the environment for the user.
2. Start the X server, if it is not already started.
3. Issue the db2cc command to start the DB2 UDB Control Center.

22.2. DB2 UDB Information Center

As mentioned above, the new DB2 UDB Information Center is a revamped documentation display system that now encapsulates the task, concept and reference information required to properly run DB2 UDB. It can be installed locally, on an intranet server, or accessed directly from an ibm.com hosting service

To start the DB2 UDB Information Center:

1. Ensure you are logged on to your Linux workstation using either the DB2 UDB Administration Server user ID or the DB2 UDB instance user ID. If you use the su command to become the user, ensure you include the −l parameter to initialize the environment for the user.
2. Ensure that the X server is started, because the Information Center requires a Web browser.
3. Issue one of these launch commands:

    - db2help
    - db2ic
    - db2cc −ic

Alternatively, you can launch the Information Center from within the Control Center (or any of the other GUI tools), by clicking an item under the "Help" menu, or by clicking on the "?" icon.

If you did not install the DB2 UDB Information Center, or have not configured your tools to use a different server, these commands will launch a detected browser, and take you directly to the IBM hosted DB2 UDB Information Center at http://publib.boulder.ibm.com/infocenter/db2help/.

22.3. DB2 UDB CLP

If you are logged on to your Linux workstation using either the DB2 UDB Administration Server user ID or the DB2 UDB instance user ID, you can issue DB2 UDB commands and SQL statements from the command line.

If this is your first time using DB2 UDB, I would suggest creating the sample database that ships with DB2 UDB. The sample database is used throughout the DB2 UDB documentation and is required by most of the sample applications.

To create the sample database, you can either select the Create the sample database option from the DB2 UDB First Steps launchpad, or by issuing the db2sampl command.

Before you can issue an SQL statement, you have to connect to a database. To connect to a database, enter the command:

```
db2 CONNECT TO database USER userID USING password
```

To connect to a database, and have DB2 UDB prompt you for the password, issue the command:

```
db2 CONNECT TO database USER userID
```

To connect to a database using the default user ID, issue the command:

```
db2 CONNECT TO database
```

Once you have connected to a database, you can then issue SQL statements or DB2 UDB commands against that database. For example, to select all of the columns from the EMPLOYEE table in the SAMPLE database, issue this command:

```
db2 "SELECT * FROM employee"
```

You can avoid typing db2 as the prefix for every SQL statement and DB2 UDB command by issuing DB2 UDB commands using the Command Line Processor (CLP). To start the CLP, issue the db2 command by itself. DB2 UDB provides the following prompt:

```
You can issue database manager commands and SQL statements from the
command prompt. For example:
    db2 => connect to sample
    db2 => bind sample.bnd

For general help, type: ?.
For command help, type: ? command, where command can be
the first few keywords of a database manager command. For example:
 ? CATALOG DATABASE for help on the CATALOG DATABASE command
 ? CATALOG          for help on all of the CATALOG commands.

To exit db2 interactive mode, type QUIT at the command prompt. Outside
interactive mode, all commands must be prefixed with 'db2'.
To list the current command option settings, type LIST COMMAND OPTIONS.

For more detailed help, refer to the Online Reference Manual.

db2 =>
```

Some users prefer to avoid the CLP because it prevents using the command line history features of the shell.

Some users claim that running the CLP within an Emacs shell gives them the best of both worlds: they get command line history, and they don't have to escape commands that contain lots of quotation marks and brackets.

---

23. Resources

This section lists available additional resources for information on using DB2 UDB 8.2 for Linux:

23.1. Newsgroups

- comp.databases.ibm−db2

  Covers the use of DB2 on all platforms, including Linux and UNIX, Windows, z/OS, iSeries, and pSeries. This forum is quite active.
- ibm.software.db2.udb

  Similar to comp.databases.ibm−db2, but handles questions specific to the DB2 Universal Database product only.
- ibm.software.db2.udb.beta

  Although this newsgroup was meant to serve users trying out the DB2 UDB V8.2 "Stinger" beta code, there are a lot of general questions about DB2 UDB on Linux that might help you. Now that the final release is available, activity on this forum has slowed down significantly.

23.2. Web sites

- DB2 Universal Database for Linux

  http://www.ibm.com/software/data/db2/linux/

  The home portal for DB2 UDB information on Linux.
- DB2 UDB Version 8.2 Information Center

  http://publib.boulder.ibm.com/infocenter/db2help/index.jsp

  The DB2 UDB Information Center is the most current source of official DB2 UDB documentation in 29 different languages. Set your browser language preference and the site will return the translated version of a given page if it is available, falling back to English content if not. Access to the information is offered via search, navigation tree, or a comprehensive master index.
- DB2 UDB Version 8.2 Linux validation site

  http://www.ibm.com/software/data/db2/linux/validate/

  Information regarding IBM recommended and validated Linux distributions for the current version of DB2 UDB.
- DB2 UDB Online Support

  http://www.ibm.com/software/data/db2/udb/support.html

  The official support channel for DB2 Universal Database for Linux, UNIX and Windows. Contains information and downloads regarding Fixpaks, APARs, Technotes, DB2 UDB Flashes, and more.
- DB2 UDB Version 8.2 PDF Manuals

http://www.ibm.com/software/data/db2/udb/support/manualsv8.html

The place to get the official DB2 UDB manuals in PDF format.

- DB2 UDB for Linux white papers

  http://www.ibm.com/software/data/db2/linux/papers.html

  White papers typically provide overviews or introductions to new technology or new releases of a product. Includes the DB2 UDB "Stinger" Enhancements on Linux white paper.
- DB2 Magazine

  http://www.db2mag.com

  This Web site is the online version of DB2 Magazine, which publishes articles about using DB2 UDB on Linux and other platforms. You can also sign up for a free subscription to the print version of the magazine.
- IBM DeveloperWorks for DB2

  http://www.ibm.com/developerworks/db2/

  A great resource for tutorials, learning resources, help and tips for improving DB2 UDB and developing DB2 UDB−based applications.
- International DB2 Users Group (IDUG)

  http://www.idug.org/html/home.asp

  IDUG holds a number of international conferences on DB2 and has regional user groups. While IDUG has traditionally focused on DB2 for z/OS and iSeries, they have increasingly included information on DB2 UDB for Linux, UNIX, and Windows. The IDUG Web site includes online discussion forums and links to other DB2 resources.
- Red Hat Linux

  http://www.redhat.com

  Home to the Red Hat Linux distribution.

  Also see http://www.redhat.com/software/rhn/ (the Red Hat Network) for more information on system updates and kernel information.
- SuSE Linux

  http://www.novell.com/linux/suse/index.html

  A Novell business, this site is the home for the SuSE Linux distribution.

  Also see http://support.novell.com/filefinder/ for more information on system updates and kernel information.
- Fedora Project

  http://fedora.redhat.com/

  Home to the Fedora Project Linux distribution.

The Fedora Project is a Red−Hat−sponsored and community−supported open source project. It is not a supported product of Red Hat, Inc.

- Mandriva Linux

  http://www.mandriva.com/

  Home to the Mandriva Linux distribution.

  Their current slogan is "A Linux for everyone".

---

24. GNU Free Documentation License

24.1. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

24.2. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world−wide, royalty−free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front−matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front−Cover Texts or Back−Cover Texts, in the notice that says that the Document is released under this License. A Front−Cover Text may be at most 5 words, and a Back−Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine−readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard−conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine−generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

24.3. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

24.4. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front−Cover Texts on the front cover, and Back−Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine−readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer−network location from which the general network−using public has access to download using public−standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

24.5. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

1. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
2. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
3. State on the Title page the name of the publisher of the Modified Version, as the publisher.
4. Preserve all the copyright notices of the Document.
5. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
6. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.

7. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
8. Include an unaltered copy of this License.
9. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
10. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
11. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
12. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
13. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
14. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
15. Preserve any Warranty Disclaimers.

If the Modified Version includes new front−matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties−−for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front−Cover Text, and a passage of up to 25 words as a Back−Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front−Cover Text and one of Back−Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

24.6. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

24.7. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

24.8. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

24.9. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

24.10. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

24.11. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See http://www.gnu.org/copyleft/.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

25. Trademarks

DB2, DB2 UDB,, DB2 Universal Database, IBM, iSeries, pSeries, zSeries, and z/OS are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Java and all Java−based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Windows is a trademark of Microsoft Corporation in the United States, other countries, or both.

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.