

DB2 Universal Database Express Edition Version 8.2 basics:

Installation, configuration, and maintenance

Peter Kohlmann and Madhava Enros

Table of Contents

1. Introduction	4
DB2 UDB Express product overview	4
2. Getting up and running	6
Installation	6
Overview of installation methods	6
Installation requirements	7
Preinstallation tasks	9
Installation	10
Postinstallation tasks	12
Basic Administration and maintenance operations	14
Database creation	14
Creating your own database using the Create Database with Automatic Maintenance wizard	14
Reviewing your new database	16
3. Frequently asked questions	18
Index of questions	18
Answers	20
Appendix A: DB2 administration tools	68
Appendix B: Automatic maintenance	82
Notices	84

1. Introduction

This document provides information on how to install DB2 Universal Database (UDB) Express. It also introduces DB2 UDB maintenance, administration, and development.

For a complete look at DB2 UDB, see the DB2 UDB Information Center at <http://publib.boulder.ibm.com/infocenter/db2help>

DB2 UDB Express product overview

IBM® DB2 Universal Database™ Express (DB2 UDB Express) is the newest member of the DB2 Universal Database Version 8 product family. It combines the power, function, and reliability of IBM's award winning DB2 UDB relational database with simplicity in packaging, installation, and deployment at a minimal investment cost to meet the data management needs of small and medium businesses.

DB2 UDB Express is designed for customers with minimal in-house database skills who need an easy-to-install database integrated into their application software solutions. It is a multi-user version of DB2 that supports local and remote applications in stand-alone and local area network (LAN) environments.

DB2 UDB Express delivers the following key functions:

Simplified deployment

The response file installation capability allows an application or solution provider to incorporate the installation of DB2 UDB Express as a transparent part of a solution. Combining response file installation with the ability to create databases and to configure DB2 UDB Express using scripts, DB2 UDB Express supports the creation of up-and-running application solutions.

DB2 UDB Express can be deployed on Linux or Windows operating systems, offering you flexibility in the choice of hardware and operating system for deployment.

Autonomic management capabilities

DB2 UDB Express incorporates a suite of advanced management and autonomic computing capabilities. These capabilities support the simplification and automation of administration tasks, eliminating many of the costs traditionally associated with managing relational databases.

The Configuration Advisor provides facilities to configure the DB2 UDB Express environment to optimize performance. The Health Center monitors your DB2 UDB system, alerting you to potential operational issues and providing advice about how to resolve them.

Autonomic maintenance can also be set up out of the box with a new Wizard. It creates a new database on the disk or directory of your choice, assigns disk space for your data, configures the new database for performance, and turns on automatic maintenance and configures notification by e-mail or pager if the database needs attention. See Appendix B, Automatic maintenance.

Designed for 24-hour availability

DB2 UDB Express is designed to support 24x7 operation; you do not need to shut down the database to perform maintenance. Database backups can be performed online or offline.

Application development support

DB2 UDB Express supports most industry standard application interfaces including SQL, X/Open, ODBC, .NET, JDBC, Java, XML and Web services. DB2 UDB Express provides add-ins for popular IDEs such as WebSphere Studio Application Developer, Microsoft Visual Studio .NET, Microsoft Visual Basic, Microsoft Visual C++, and Microsoft Visual InterDev.

Application developers can leverage the DB2 UDB Express advanced functions in the programming model they are familiar with, leverage their existing skills, reduce time to market, cut overall application deployment costs, and accelerate time to value.

Investment protection

DB2 UDB Express offers the same core capabilities across the entire line of IBM database servers. You can benefit from the same reliability, security, application development interfaces and optimization as your data volumes, number of users, or utilization increase. You can leverage your investment with DB2 UDB Express through upgrades to other scalable editions of DB2 UDB to support more operating systems and higher throughput on larger hardware installations.

2. Getting up and running

Installation

Overview of installation methods

DB2 Setup wizard

The DB2 Setup wizard is an installer that is available on both Linux and Windows platforms. The DB2 Setup wizard provides an easy-to-use interface for installing DB2 UDB and performing initial setup and configuration tasks.

Response file installation

A response file is an ASCII file that contains setup and configuration values. The file is passed to the db2setup program and the installation is performed according to the values that have been specified in the response file. There are a number of ways to create a response file:

- Using the response file generator (Windows).
- Using the DB2 Setup wizard (Linux and Windows).
- By customizing sample response files that are provided for each DB2 product (Linux and Windows).

Installation overview (Windows)

This topic outlines the steps for installing a DB2 UDB Express server on Windows operating systems.

Procedure

To install a DB2 UDB Express server on Windows:

1. Review the DB2 product prerequisites.
2. Install the DB2 product using the DB2 Setup wizard. DB2 Setup wizard features include:
 - A launchpad from which you can view installation notes and release notes, and learn about DB2 UDB Version 8 features.
 - Typical, Compact, and Custom installation types.
 - Multiple languages installation support.
 - DB2 Administration Server setup (including DAS user setup).
 - Instance setup and configuration (including instance user setup).
 - Response file creation. You can save your installation choices in a response file for later installation or to duplicate the installation on another computer.

Installation overview (Linux)

This topic outlines the steps for installing a DB2 UDB Express server on Linux.

Procedure

To install a DB2 UDB Express server on Linux:

1. Review the DB2 product prerequisites.
2. Mount the DB2 product CD-ROM.
3. Modify kernel parameters.
4. Install the DB2 product using the DB2 Setup wizard. DB2 Setup wizard features include:
 - A launchpad from which you can view installation notes and release notes, and learn about DB2 Version 8 features
 - Typical, Compact, and Custom installation types
 - Multiple languages installation support
 - DB2 Administration Server setup (including DAS user setup)
 - Instance setup and configuration (including instance user setup)
 - Response file creation. You can save your installation choices in a response file for later installation or to duplicate the installation on another computer.

Installation requirements

Disk and memory requirements (Windows and Linux)

Disk requirements

The disk space required for your product depends on the type of installation you choose and the type of file system you have. The DB2 Setup wizard provides dynamic size estimates based on the components selected during a typical, compact, or custom installation.

On Windows, you might require significantly more space on FAT (File Allocation Table) drives than with NTFS (New Technology File System) drives.

Remember to include disk space for required software, communication products, and documentation.

Memory requirements

At a minimum, DB2 UDB requires 256 MB of RAM. 512MB of RAM memory is recommended if you use the DB2 tools. When determining memory requirements, be aware of the following:

- For DB2 client support, these memory requirements are for a base of 5 concurrent client connections. You will need an additional 16 MB of RAM per 5 client connections.
- Additional memory is required for other software that is running on your system.
- Additional memory might be required to improve the performance of the DB2 tools.

- Specific performance requirements can determine the amount of memory needed.
- Memory requirements are affected by the size and complexity of your database system.
- Memory requirements are affected by the extent of database activity and the number of clients accessing your system.
- On Linux, ensure that your SWAP space is at least twice as large as your RAM.

Windows installation requirements

To install DB2 UDB Express, the following operating system, software, and communications requirements must be met:

Operating system requirements

DB2 UDB Express runs on:

- Windows NT Version 4 with Service Pack 6a or higher
- Windows 2000. Service Pack 2 is required for Windows Terminal Server.
- Windows XP Professional
- Windows Server 2003

Hardware requirements

An Intel x86 compatible processor.

Software requirements

A browser is required to view online help.

Communication requirements

You can use TCP/IP, Named Pipes, and NetBIOS. To remotely administer a DB2 Version 8 database, you must connect using TCP/IP.

For TCP/IP, Named Pipes, and NetBIOS connectivity, no additional software is required.

If you plan to use LDAP (Lightweight Directory Access Protocol), you require either a Microsoft LDAP client or an IBM SecureWay LDAP client V3.1.1.

Linux installation requirements

Installation requirements (Linux)

This topic lists the hardware, distribution, software, and communication requirements for DB2 UDB Express.

Hardware requirements

An Intel x86 compatible processor.

Distribution requirements

For the latest information on supported distribution and kernel levels, point your browser to: <http://www.ibm.com/db2/linux/validate>

Software requirements

The IBM Developer Kit for Java 1.4.1 (SDK) is required for DB2 servers, for using the DB2 Control Center, and for creating and running Java applications, including stored procedures and user-defined functions. During the installation process, if the correct level of the SDK is not already installed, it will be installed.

A browser is required to view online help.

Communication requirements

TCP/IP is required to access remote databases.

Installing DB2 products or sharing instance directory on NFS

Currently, we do not support the installation of DB2 products on NFS. Installing DB2 on NFS (for example, NFS mounting /usr/opt/db2_08_01 or /opt/IBM/db2/V8.1) can be error prone and these errors can be difficult to diagnose.

The following configuration is not supported:

- Setting up an instance on a file system
- NFS mounting a file system from multiple computers, and then running DB2 on these machines using that same instance.

This configuration can cause file locking and performance problems.

Preinstallation tasks

Extending the directory schema (Windows)

If you plan to use LDAP with Windows 2000 or Windows Server 2003, you must extend the directory schema to contain DB2 object classes and attribute definitions. You must do this before you install any DB2 products. Refer to the *DB2 UDB Express Quick Beginnings Guide* for more information.

Installing the IBM Software Developer Kit for Java (Linux)

The IBM Software Developer Kit for Java 1.4.1 (SDK) is required for DB2 servers to use the DB2 Control Center or to create and run Java applications, including stored procedures and user-defined functions. Only the IBM Software Developer's Kit (SDK) is supported. The IBM Developer Kit for Java 1.4.1 (SDK) for 32-bit platforms is shipped on the product CD.

The installation of the IBM Developer Kit for Java 1.4.1 is attempted automatically during a response file or a GUI installation. This task is only required if the automatic installation fails or if you did not install the GUI tools initially with the product.

The IBM Software Developer Kit for Java 1.4.1 is installed whenever a component that requires Java is being installed. However, if the installer detects that the SDK is already installed, it will not install it again. The SDK is installed in its own directory and does not overwrite any previous levels of the SDK .

Installation

Windows installation

Installing DB2 UDB Express using the DB2 Setup wizard (Windows)

This task describes how to install DB2 UDB Express using the DB2 Setup wizard on Windows. You will use the DB2 Setup wizard to define your installation and install DB2 UDB Express to your system.

Prerequisites

Before you start the DB2 Setup wizard:

- Ensure that your system meets installation, memory, and disk requirements.
- You must have a local administrator user account with the recommended user rights to perform the installation.

Procedure

To start the DB2 Setup wizard:

1. Log on to the system with the administrator account that you have defined for the DB2 UDB Express installation.
2. Close all programs so the installation program can update files as required.
3. Insert the CD-ROM into the drive. If enabled, the auto-run feature automatically starts the DB2 Setup launchpad. From this window, you can view installation prerequisites and the release notes. You might want to review the installation prerequisites and release notes for late-breaking information. Select **Install Products** and select the DB2 product that you want to install.

The DB2 Setup wizard determines the system language and launches the setup program for that language. If you want to run the setup program in a different language, or the setup program failed to auto-start, you can start the DB2 Setup wizard manually.

To start the DB2 Setup wizard manually:

1. Click Start and select the Run option.
2. In the Open field, enter the following command:

x:\setup /i language

where:

x: represents your CD-ROM drive and **language** is the territory identifier for your language (for example, EN for English). If the **/i** flag is not specified, the installation program runs in the default language of the operating system.

3. Click **OK**.

Once you have initiated the installation, proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help** or press F1. You can click **Cancel** at any time to end the installation.

If you want to verify your installation using the sample database, be sure to install the sample database component found under the Getting Started component group. The sample database is included as part of a typical installation.

For information on errors encountered during installation, see the db2.log file. The db2.log file stores general information and error messages resulting from the installation and uninstallation activities. By default, the db2.log file is located in the `My Documents\DB2LOG\` directory. The location of the `My Documents` directory depends on the settings on your computer.

Linux installation

Installing DB2 UDB Express using the DB2 Setup wizard (Linux)

This task describes how to install DB2 UDB Express using the DB2 Setup wizard on Linux systems. The DB2 Setup wizard is used to define your installation preferences and install DB2 UDB Express to your system.

Prerequisites

Before you start the DB2 Setup wizard:

- Ensure that your system meets installation, memory, and disk requirements.
- You require root authority to perform the installation.
- The DB2 UDB Express product CD-ROM must be mounted on your system.
- You must have X/Windows software capable of rendering a graphical user interface for First Steps to run on your computer. Ensure that you have properly exported your display. For example, export `DISPLAY=9.26.163.144:0`.

Procedure

To start the DB2 Setup wizard:

1. Log on to the system as a user with root authority.
2. Insert the DB2 UDB Express product CD-ROM.
3. Change to the directory where the CD-ROM is mounted by entering the following command:
cd /mnt/cdrom
where `/cdrom` represents the mount point of the CD-ROM.
4. Enter the **./db2setup** command to start the DB2 Setup wizard.
5. After a moment, the IBM DB2 Setup launchpad opens.

From this window, you can view installation prerequisites and the release notes, you can take a Quick Tour to explore the features of DB2 UDB, or you can proceed directly to the installation. You might want to review the installation prerequisites and release notes for late-breaking information.

Once you have initiated the installation, proceed through the DB2 Setup wizard installation pages and make your selections. Installation help is available to guide you through the remaining steps. To invoke the installation help, click **Help** or press F1. You can click **Cancel** at any time to end the installation.

When you have completed your installation, DB2 UDB Express will be installed in the `/opt/IBM/db2/V8.2` directory.

Postinstallation tasks

Verifying the installation using the command line processor (CLP)

Once you have completed installing DB2 UDB Express, you can verify the installation by creating a sample database and running SQL commands to retrieve sample data.

Prerequisites

The sample database component must be installed on your system; this is included in a typical installation.

Procedure

To verify the installation:

1. Log on to the system. On Linux, use the `db2inst1` ID. On Windows, use either a Windows administrator ID or the ID you specified during installation.
2. Enter the **db2sampl** command to create the SAMPLE database. This command may take a few minutes to process. There is no completion message; when the command prompt returns, the process is complete. The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.
3. Start the database manager by entering the **db2start** command.
4. Enter the following DB2 commands from a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees that work in department 20, and reset the database connection:

```
db2 connect to sample  
db2 "select * from staff where dept = 20"  
db2 connect reset
```

After you have verified the installation, you can remove the SAMPLE database to free up disk space. Enter the **db2 drop database sample** command to drop the SAMPLE database.

Verifying the installation of DB2 UDB Express using First Steps

You should verify that the DB2 UDB Express installation has completed successfully by accessing data from the SAMPLE database.

Prerequisites

You must have the Control Center and the First Steps component installed to perform this task. First Steps is part of the Getting started component grouping in the DB2 Setup wizard. It is installed as part of a typical installation or may be selected when performing a custom installation.

You must have SYSADM or SYSCTRL authority to perform this task (for more information on security, see section 5. Security).

For Linux, you must have Xwindow software capable of rendering a graphical user interface for First Steps to run on your computer. Ensure that you have properly exported your display. For example, export DISPLAY=9.26.163.144:0.

On Linux, ensure that you are logged on as the instance owner (db2inst1 by default). An instance is a logical database manager environment where you catalog databases and set configuration parameters.

Procedure

1. Log on to the system with the user account that you want to use to verify the installation. Ensure that the domain user account you use to create the sample database has SYSADM or SYSCTRL authority.
2. Start First Steps by running the **db2fs** command. This command is found in the DB2 folder on Linux (/opt/IBM/db2/V8.2/bin/) and in the Start menu on Windows.
3. Select Create SAMPLE databases on the First Steps launchpad to open the Create SAMPLE databases window.
4. In the Create SAMPLE databases window, select **DB2 UDB Sample database**.
5. Click **OK**. This command may take a few minutes to process. When the SAMPLE database has been created, you will receive a completion message. Click **OK**.
6. Once the database is created, select Work with Databases on the First Steps launchpad to start the Control Center. The Control Center allows you to perform administration tasks on different instance and database objects.
7. Select **Basic** in the Control Center View window and select **OK**. Your new sample database should appear under the **All Databases** folder. Expand the SAMPLE database folder and select the **Tables** folder. You can see the contents of any table in the SAMPLE database by double clicking on it in the list of tables on the right.

After you have verified the installation, you can remove the SAMPLE database to free up disk space. To remove the SAMPLE database, right-click on the SAMPLE database object and select **Drop**.

Basic administration and maintenance operations

Database creation

Now that you have installed DB2 UDB Express and verified its successful installation with the SAMPLE database, you are ready to create your own database. The following sections guide you through the steps of creating a new database that will take advantage of DB2 automatic maintenance features.

Creating your own database using the Create Database with Automatic Maintenance wizard

The Create Database with Automatic Maintenance wizard is the quickest way to set up your own database. Based on information that you provide, the wizard:

- Creates a new database on the disk or directory of your choice.
- Assigns disk space for your data.
- Configures the new database for optimal performance.
- Turns on automatic maintenance.
- Configures health notification by e-mail or pager.

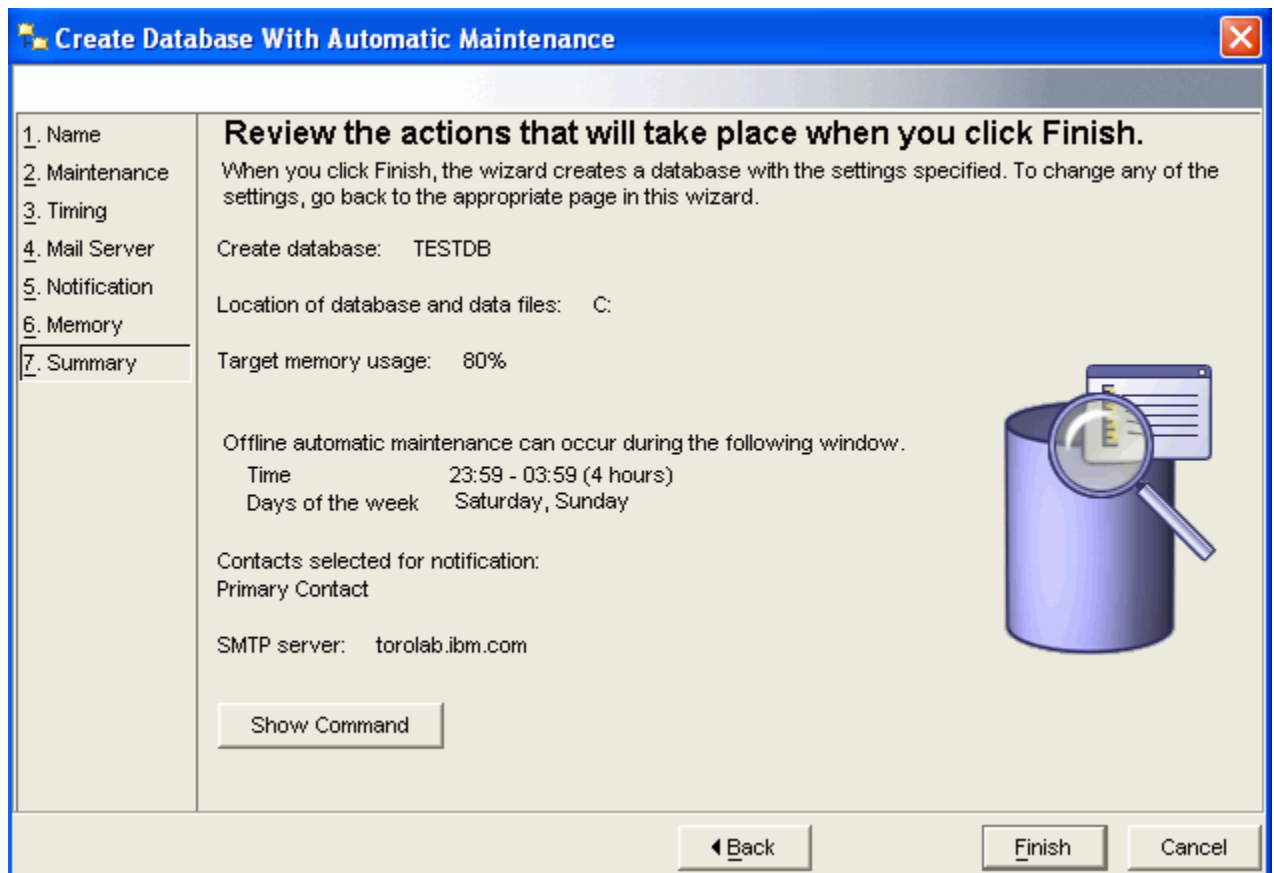


Figure 1. Create Database with Automatic Maintenance wizard

The Create Database with Automatic Maintenance wizard can be started from First Steps or from directly within the Control Center.

Prerequisites

If you wish to start the wizard from First Steps, you must also have the First Steps component installed. First Steps is part of the Getting started component grouping in the DB2 Setup wizard. It is installed as part of a typical installation or may be selected when performing a custom installation.

You must have SYSADM or SYSCTRL authority to perform this task (for more information on security, see section 5. Security).

For Linux, you must have X Windows software capable of rendering a graphical user interface for First Steps and the Control Center to run on your computer. Ensure that you have properly exported your display. For example, export DISPLAY=9.26.163.144:0.

On Linux, ensure that you are logged on as the instance owner (db2inst1 by default). An instance is a logical database manager environment where you catalog databases and set configuration parameters.

Procedure

- Log on to the system as the instance owner to create the database.
- Start the Create Database with Automatic Maintenance wizard in the one of the following ways:
 - From First Steps: Click on the button marked **Create Your Own Database**
 - From the Control Center: Click on the **All Databases** folder. Click on **Create New Database** in the All Databases pane. Follow the steps of the wizard. You need to provide some details and preferences about your new database, including such things as its name and the contacts that can be notified under various conditions. You also need to select one of two database maintenance strategies.
- If you created your own database from First Steps, you must start the Control Center to begin exploring and working with your database. To start exploring your database, click on the button marked **Work With Databases in First Steps**.

Reviewing your new database

Once you have started the Control Center, you can verify that your database was created by finding it in the Object Tree that fills the left side of the Control Center window.

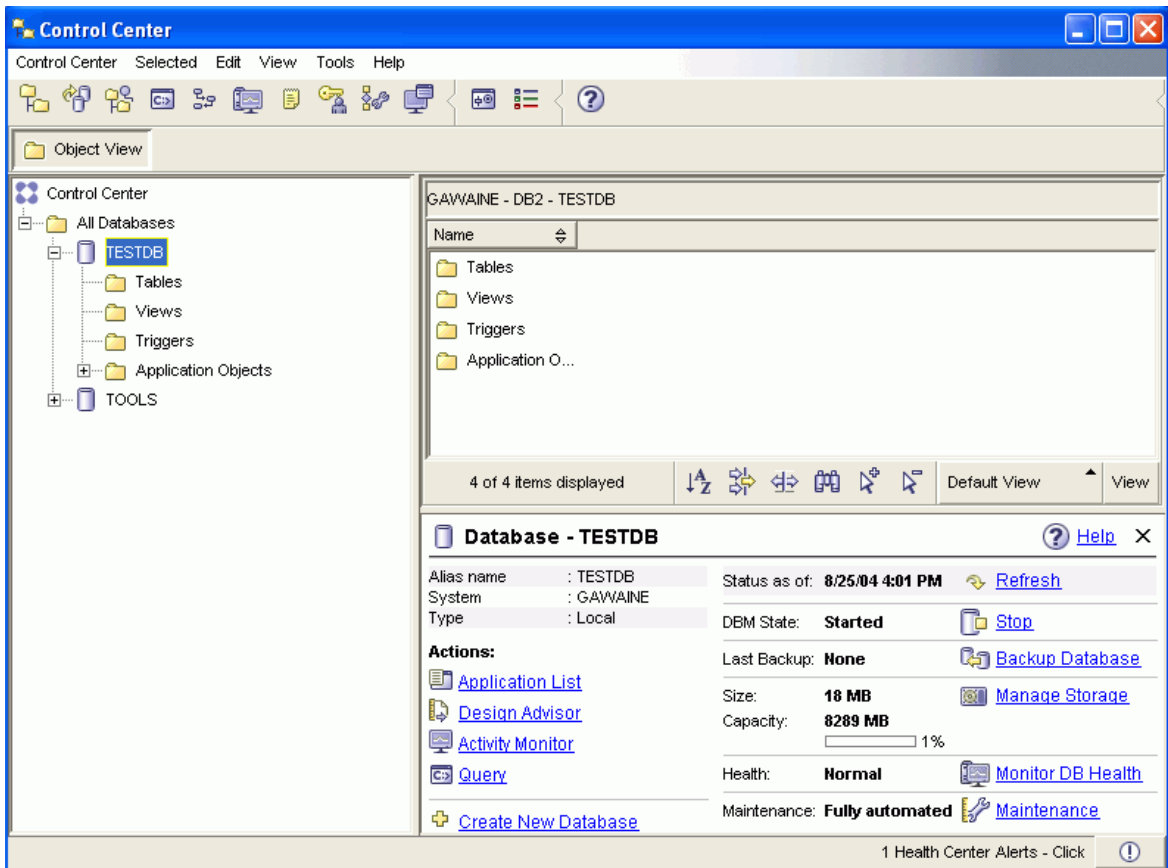


Figure 2. Control Center with a database selected

Expand the Object Tree until you can see your database and select it. The right side of the Control Center will fill with details about the database's current operation and links to relevant actions (see Appendix A, DB2 administration tools).

When you created your database with the Create Database with Automatic Maintenance wizard, you enabled certain Automatic features. The features are set up and do not necessarily require any further attention. You might, however, find it useful to know where their settings can be found and further configured.

The first tool is the Automatic Maintenance wizard, which you can start by clicking on the Maintenance link in the Object Details Pane (see Figure 2). As the wizard's name suggests, you can use it to change the way in which automatic maintenance is carried out for your database.

The second tool is the Configuration Advisor. It can be started by right-clicking on the database to be configured, and selecting Configuration Advisor. This tool helps you configure your database for optimal performance easily.

The third and last of these tools is the Health Center. You can open it by clicking on Monitor DB Health in the Object Details Pane or by selecting Health Center from the Tools menu. Under normal circumstances, you should never have to look at the Health Center. If the health of your database needs attention, you will be notified by e-mail. You can then find out more about what needs attention.

3. Frequently asked questions

If you have followed the steps outlined in section 2, Getting up and running, you should have DB2 UDB installed and a database created. You might have questions regarding using, maintaining, and tuning your database. Many of the most common of these questions are listed below with answers.

A. Index of questions

1. General

- 1.1. How do I issue DB2 commands and SQL statements?
- 1.2. How do I interpret error codes?
- 1.3. How do I automate the set up of a DB2 client?
- 1.4. Where can I get more information about DB2 UDB?
- 1.5. How can I access the DB2 Information Center without Internet access?
- 1.6. How can I schedule and script tasks?
- 1.7. How do I create a tools catalog when I can't schedule tasks?
- 1.8. Where can I find a log file that I can send to my application vendor when something seems wrong?
- 1.9. What are FixPaks? Why and how do I apply them?

2. DB2 UDB performance

- 2.1. I'm experiencing a DB2 UDB performance problem. What should I do?
 - 2.1.1. Is DB2 UDB the problem?
 - 2.1.2. Are you receiving health notifications?
 - 2.1.3. Has your hardware configuration changed?
 - 2.1.4. Is there a configuration conflict?
 - 2.1.5. Is automatic maintenance turned on and working?
 - 2.1.6. Have you checked statistics?
 - 2.1.7. How can I confirm that my table statistics are up-to-date?
 - 2.1.8. Have you done a reorg check?
 - 2.1.9. Is there a locking problem?
 - 2.1.10. Are there database design issues?
- 2.2. Can performance problems be caused by infrequent database use?
- 2.3. Is there a simple way to see if scheduled tasks are taking longer to complete than normal?
- 2.4. How can I confirm that reorganizations are taking place?

3. Health

- 3.1. How do I get recommendations about database health?
- 3.2. I haven't been notified of any problems...
 - 3.2.1. How can I confirm that nothing is wrong?
 - 3.2.2. How can I confirm that database health is being checked?
 - 3.2.3. How can I confirm that health notification is working?

4. Storage and recoverability

- 4.1 What kinds of failures do I need to protect my database from?
- 4.2 What kind of protection do I get with basic automatic backup?
- 4.3 Where can I find a record of my backups?
- 4.4 Where does automatic maintenance place my database backup images by default?
- 4.5 What exactly is included and not included in a DB2 UDB backup?
- 4.6 What other options do I have when it comes to backing up and restoring my database?
- 4.7 Why should I care about logging and log files? With my old database, I didn't even have log files.
- 4.8 What is the difference between circular (default) and archive (advanced) logging?
- 4.9 How do I switch from the circular (default) to archive (advanced) logging?
- 4.10 I am processing a large amount of data in my database and want to take a backup in case anything goes wrong. How can I take a manual backup before automatic maintenance runs?
- 4.11 How can I take advantage of online backups and point-in-time recovery even after a disk crash when my log files are taking up a lot of room on my primary disk?
- 4.12 How do I make sure that there are no users accessing the database before I back it up or perform manual maintenance?
- 4.13 What is the best way to ensure that I never lose my log files?
- 4.14 How can I recover my database on a different computer?

5. Security

- 5.1 What is the difference between authentication and authorization?
- 5.2 How do I add a user?
- 5.3 How is my user ID authenticated in Windows and in Linux?
- 5.4 How can I change the authorization of my DB2 user?

B. Answers

1. General

1.1 How do I issue DB2 commands and SQL statements?

DB2 commands and SQL statements can be issued from the Command Line Processor (CLP) text-only interface or through the Command Editor. Also, each Wizard or dialog in the Control Center includes a Show Command or a Show SQL button that displays the command or the SQL statement about to be executed. You have the option of running the command, saving it to a file or the Task Center, or just learning about DB2 UDB SQL and command structure (see Figure 4).

You can start a stand-alone Command Editor in the following ways:

1. Select Command Editor from the **Tools** menu of another tool.
2. Click the **Command Editor** icon from the toolbar of another tool.
3. On Windows operating systems, click **Start > Programs > IBM DB2 > Command Line Tools > Command Editor**.

Additionally, you can open the Command Editor as a new tab in the Control Center by clicking on the **Query** action in the Object Details pane for any database or table object, as shown in Figure 3.




Figure 3. Query action for a given database

This action opens the integrated Command Editor, as shown in Figure 4.

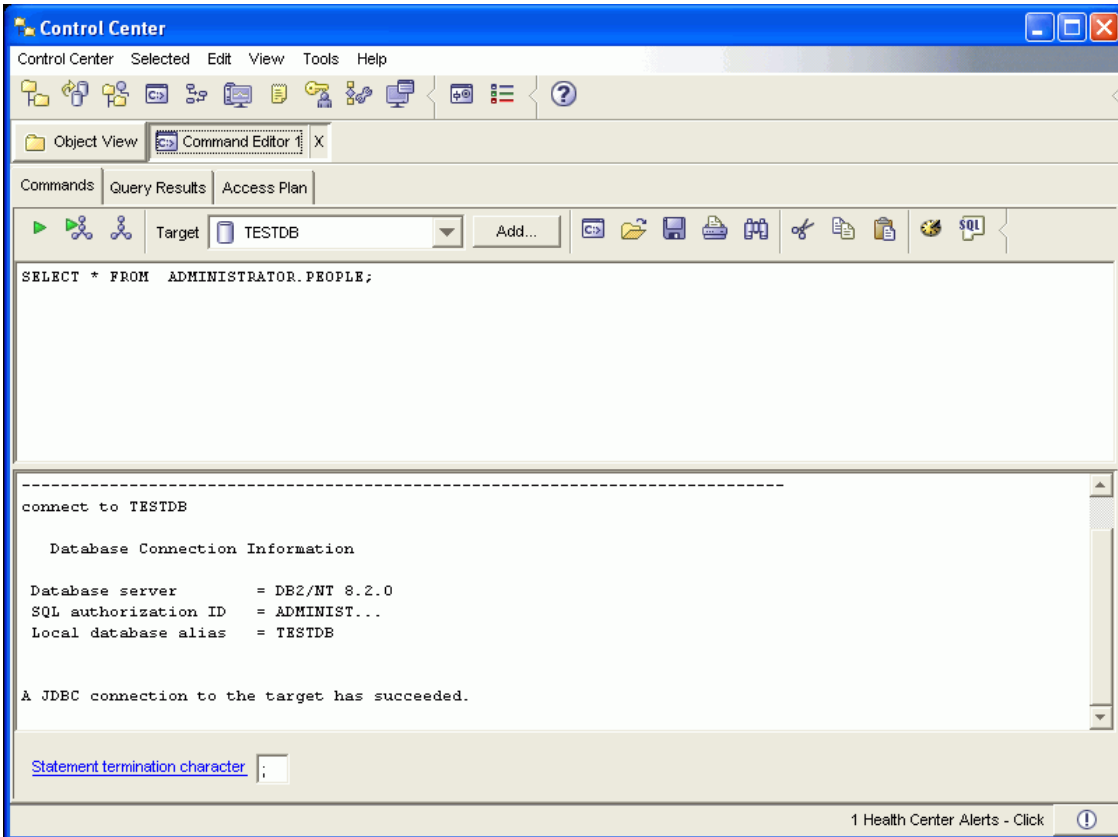


Figure 4. Command Editor integrated into the Control Center

You can use the Command Editor to accomplish the following tasks:

- Run DB2 commands and SQL statements.
- Work with query results.
- View a graphical representation of the access plan for explained SQL statements.
- Run commands in sequence.
- Create and save scripts. You can optionally store a saved script in the Task Center, where you can schedule the script to run at a specific time.
- Look up error codes.

For more detail about the Command Editor see Appendix A, DB2 administration tools.

1.2 How do I interpret error codes?

Every interaction with DB2 UDB, whether through SQL or other means (such as the various programming interfaces) is classified as successful, successful with warnings, or errors.

For every error returned by DB2 UDB, there is always an associated SQLCODE and SQLSTATE.

SQLCODE is an integer status code specific to DB2 UDB and is consistent across DB2 UDB for Linux and Windows operating systems. A few SQLCODEs are common to all database vendors,

such as +100 (“NOT FOUND”), but the vast majority of codes are unique to each vendor. In general, the rules for SQLCODEs are:

SQLCODE = 0	the command was successful
SQLCODE > 0	the command was successful, but returned a warning
SQLCODE < 0	the command was unsuccessful and returned an error

SQLSTATE is a five-digit numeric string that conforms to the ISO/ANSI SQL92 standard. The first two characters of SQLSTATE are known as the SQLSTATE class code.

- A class code of 00 means the command was successful.
- A class code of 01 implies a warning
- A class code of 02 implies a not found condition.
- All other class codes are considered errors

The following samples are examples of error messages:

SQL0911N The current transaction has been rolled back because of a deadlock or timeout. Reason code 2. SQLSTATE=40001

SQL0100W No row was found for FETCH, UPDATE or DELETE; or the result of a query is an empty table. SQLSTATE=02000

SQLCODE SQL0911N means -911 (error, with N implying negative) and SQL0100W means +100 (warning, with W implying warning). The SQLSTATEs associated with these errors are 40001 and 02000, respectively.

You can get more information about an error by using a question mark to prefix the SQLCODE in the Command Editor.

Example: ? SQL0911N

This information is also available through the Information Center. Search for the string that you would type into the Command Editor; be sure to surround the string with quotation marks.

Example: “SQL0911N”

1.3 How do I automate the setup of a DB2 client?

When you installed your DB2 server, you probably installed a client on the same computer. You might also want to set up a DB2 client on another computer so that DB2 UDB can be administered or accessed remotely.

The process of setting up a DB2 UDB Express client breaks down into four main steps:

1. Preparing to set up a DB2 UDB server profile.
2. Exporting a server profile (this includes everything you need to set up the connection from the remote client to the DB2 server).
3. Recreating the database connection on the remote computer.
4. Testing the connection.

Essentially, the process involves first creating a connection between a server and a client that resides on the same computer, but creating this connection *as if the database were remote*. Once this connection has been created, you can export a profile that describes the way in which the

client connects to the server; this profile can then be used for the new, remote client. Figure 5 explains this process.

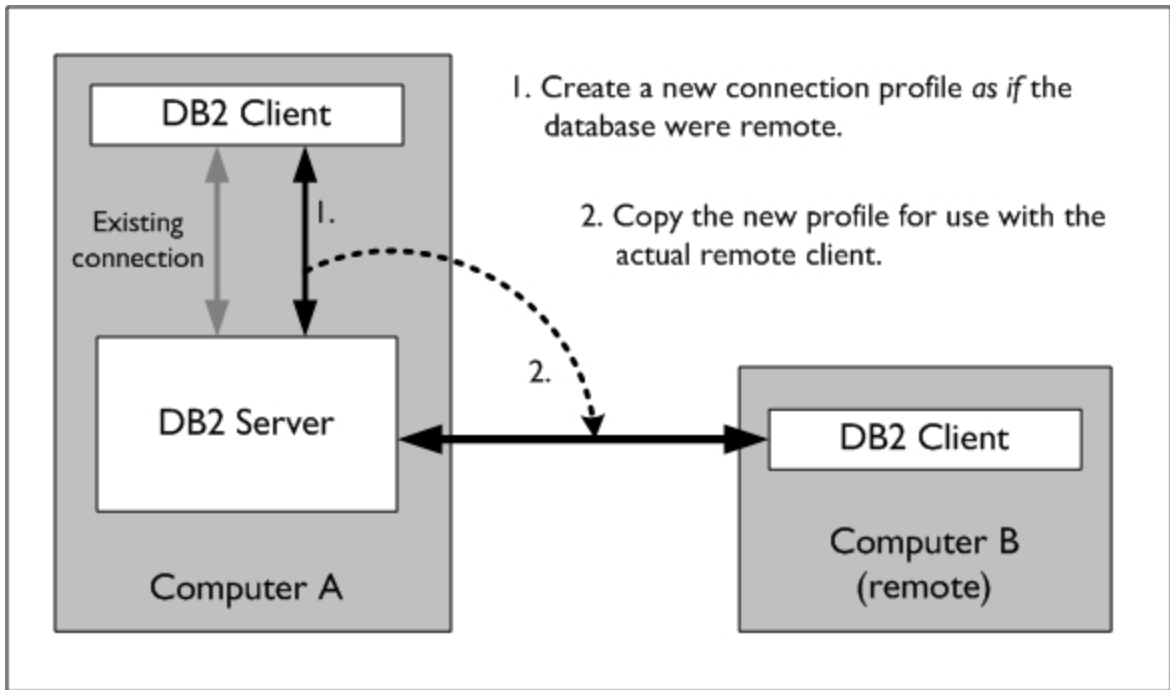


Figure 5. Creating a remote client connection

To automate the setup of a DB2 client:

1. Open the DB2 UDB Express Configuration Assistant.

Click **Start > IBM DB2 > Set-up Tools > Configuration Assistant**

The configuration assistant lists all the local and remote databases that are catalogued on this computer; in addition, it lets you catalog new systems and databases, configure them, and package everything you need to reproduce your environment on another computer. The Configuration Assistant is shown in Figure 6.

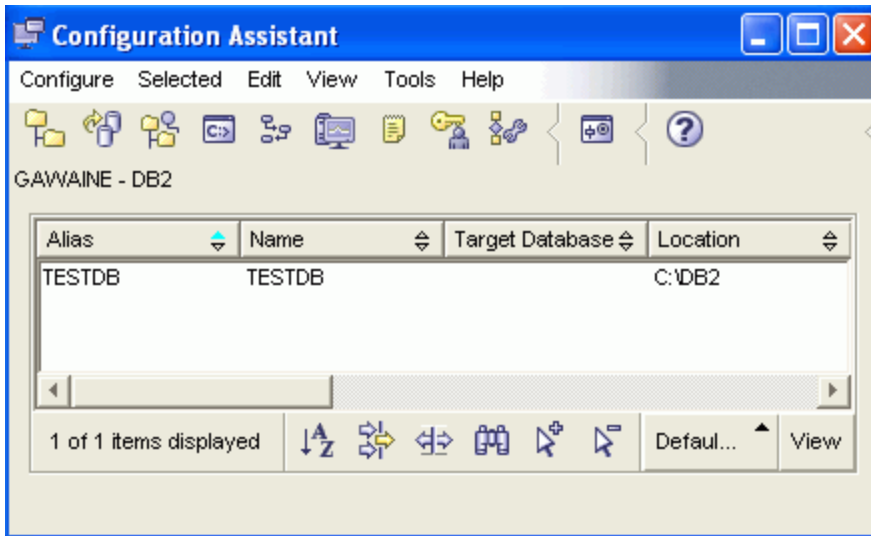


Figure 6. Configuration Assistant

2. Create a remote connection to your local database by selecting the Add Database wizard

Click **Selected > Add Database**

On page 1 (Source) of the wizard, use the DB2 UDB database discovery tool by selecting **Search the network**. It will find your database and retrieve everything you need to set up a successful connection.

On page 2 (Network), find your database through the **Other systems** folder; select it and click **Next**.

On page 3 (Alias), give your database a unique alias such as Remote.

On page 4 (Data Source), register the database as an ODBC source if this is required. Click **Finish**. You will be presented with the opportunity to test the connection (see Figure 7). You will need to provide a password, because this test is acting as a remote connection.

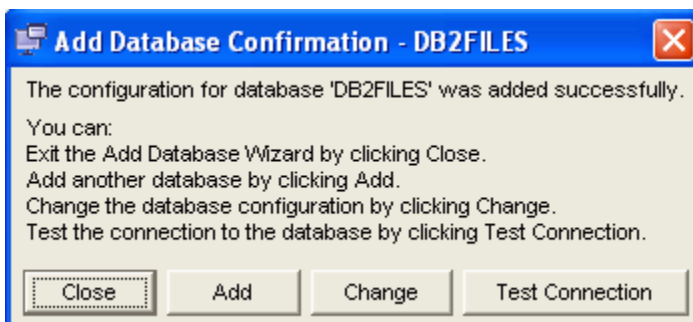


Figure 7. Test Connection button, Add Database Confirmation

Once this test is done, click **Close**. You should see your new remote connection pointing to your database. Notice the database's location, in the **Location** column of Figure 8.

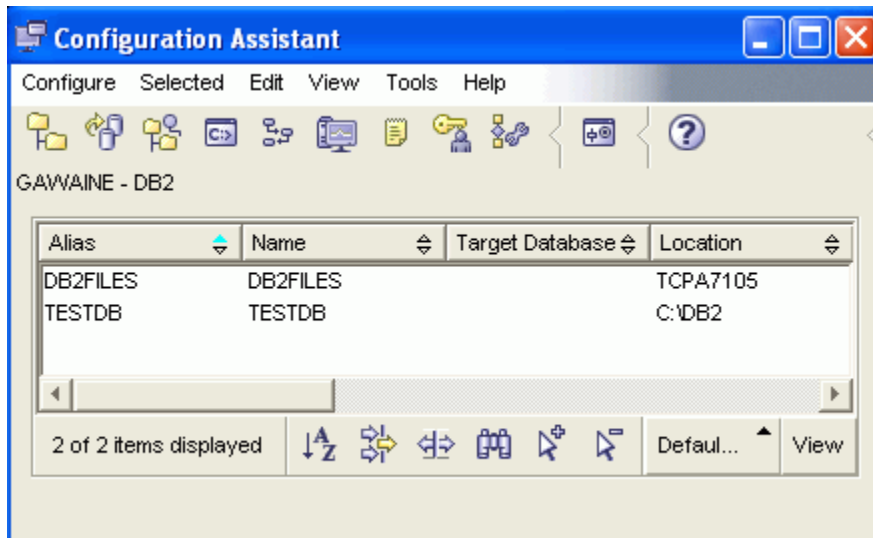


Figure 8. Configuration Assistant with new connection

3. Now you can export the profile for this remote connection so that it can be used on a remote computer.

Click **Configure > Export Profile > Customize**

The Customize Export Profile window opens.

Give the profile a file name (eg., C:\testremoteaccpac.db2).

Select the Database connections check box so that you can export only the remote connection.

Select your new remote connection from the list of available database aliases.

Click **Export**.

Upon successful export, you should see the message “Exported database aliases successfully.”

4. If you are curious, you can open the profile file in a text editor to see the details about the format of the profile, communication setup, and ODBC configuration.

5. Install a DB2 client on the remote computer. DB2 clients are available online through the DB2 Web site. The Runtime Client is a minimal installation that does not include the administration tools or more than is necessary to connect to a DB2 database. If you want the administration tools, install the Administration Client.

6. Create a connection from the remote computer. You can create this connection using the Configuration Assistant or from the command line.

Configuration Assistant

Click **Selected > Add Database**

On page 1 (Source) of the Add Database wizard, select **Use a profile**.

Type the profile file name, and continue through the remainder of the wizard.

Command Line

Issue the **db2cfimp** command, as follows:

```
C:\db2cfimp [profile file name]
```

1.4 Where can I get more information about DB2 UDB?

There is DB2 product information available online in the following forms:

- DB2 Information Center
<http://publib.boulder.ibm.com/infocenter/db2v8luw/index.jsp>
- DB2 manuals in PDF format
<http://www.ibm.com/software/data/db2/udb/support/manualsv8.html>
- DB2 newsgroups can often provide answers to common problems
 - comp.databases.db2
 - ibm.software.db2.udb
- developerWorks provides technical resources for DB2 information management software
<http://www.ibm.com/developerworks/db2/>

1.5 How can I access the DB2 Information Center without Internet access?

The DB2 Information Center, the main DB2 product documentation repository, can be installed from the DB2 product CDs as well as being available online. It is installed separately from other DB2 products and from its own CD.

To start the DB2 Setup wizard (Windows):

Insert the DB2 Information Center CD into the drive. The auto-run feature starts the DB2 Setup wizard. The DB2 Setup wizard determines the system language and starts the setup program for that language. If you want to run the setup program in a different language, or the setup program failed to auto-start, you can start the DB2 Setup wizard manually.

The DB2 Setup launchpad opens. Select **Install Products** from the menu along the left side of the window.

Once you have initiated the installation, proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. To invoke the online help, click Help or press F1. You can click Cancel at any time to end the installation.

To start the DB2 Setup wizard manually (Windows):

1. Click **Start > Run**.
2. In the **Open** field, type the following command:

`x:\setup`

where `x:` represents your CD-ROM drive

3. Click **OK**.

To start the DB2 Setup wizard (Linux):

Log on to the system as a user with root authority. Insert the DB2 Information Center CD-ROM into the drive. Change to the directory where the CD-ROM is mounted by entering the **cd /mnt/cdrom** command (where **/cdrom** represents the mount point of the CD-ROM). Enter the **./db2setup** command to start the DB2 Setup wizard. After a moment, the IBM DB2 Setup launchpad opens. Select **Install Products** from the menu along the left side of the window.

Once you have initiated the installation, proceed by following the setup program's prompts. Online help is available to guide you through the remaining steps. To invoke the online help, click Help or press F1. You can click Cancel at any time to end the installation.

1.6 How can I schedule and script tasks?

The DB2 UDB automatic maintenance features can take care of most of the regularly required maintenance tasks for you. If you want to learn more about these features, see Appendix B, Automatic maintenance.

You may want to script and schedule tasks, maintenance or otherwise, to supplement the DB2 UDB automatically scheduled task set. To that end, DB2 UDB provides two administrative tools, the Task Center and the Journal, to help you schedule and script your own database tasks. In addition, each utility available through the Control Center for example, the Backup wizard, includes a schedule page.

Note: If you are having difficulties scheduling tasks, see question 1.7: "How do I create a tools catalog when I can't schedule tasks?"

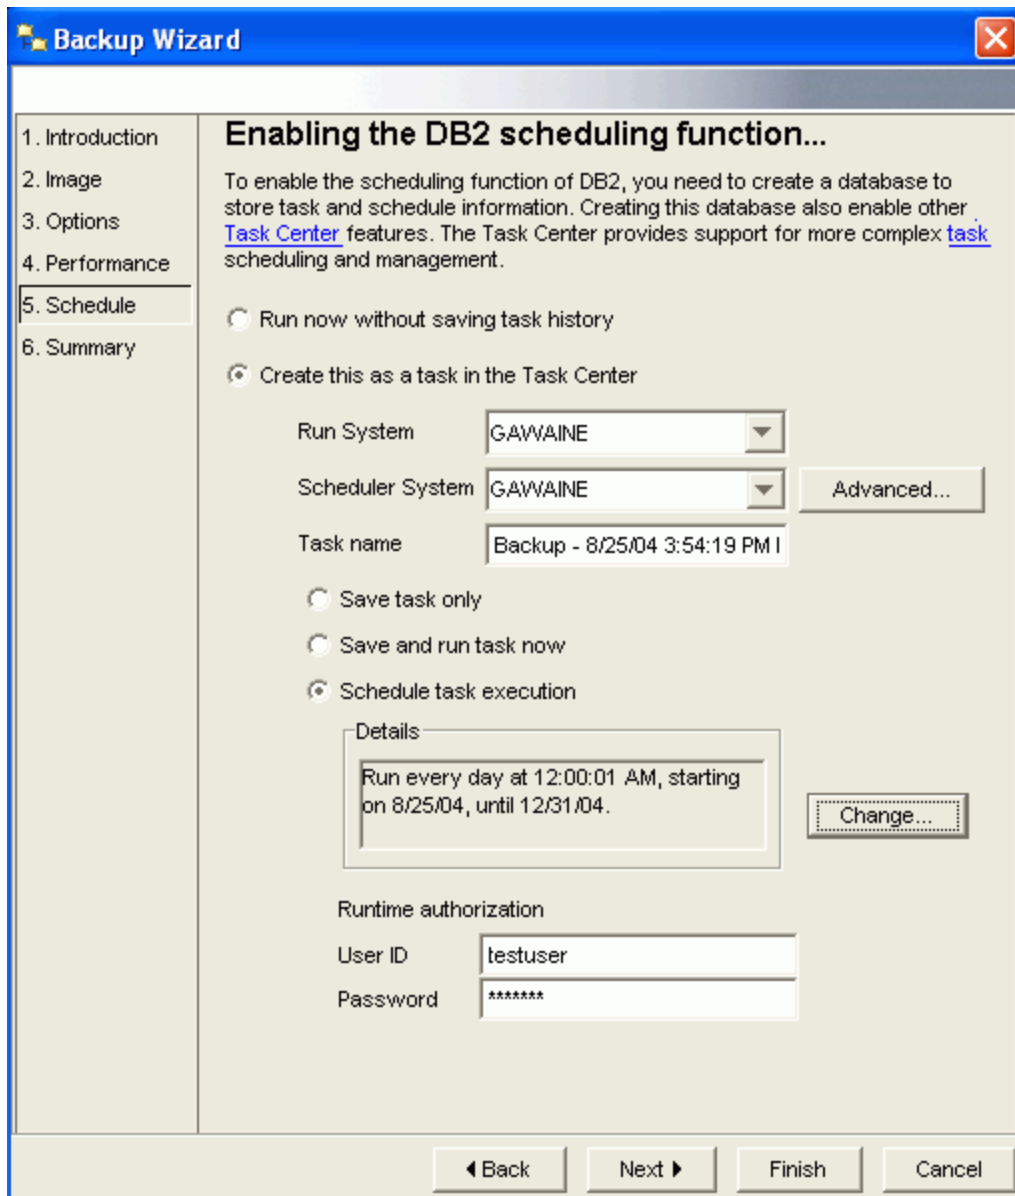


Figure 9. Task scheduling through wizards and notebooks

Use the Schedule page of various wizards and notebooks to run a task now, or to schedule a task to run later. You can also create a task and save the script in the Task Center. The task history is automatically saved in the Journal. Do one of the following actions:

1. To run the task now, without creating a task in the Task Center or saving the task history to the Journal, click **Run now without saving task history**.
2. To create a task for generating the script and saving it in the Task Center, click **Create this as a task in the Task Center**. Then, specify the task information and options.
3. Specify the name of the system on which you want to run the task, in the **Run System** field. This system must be online at the time the task is scheduled to run.
4. In the **Scheduler System** field, specify the system where you want to store the task and the schedule information.

This system will store the task and notify the run system when it is time to run the task. The drop-down list contains any system that is cataloged and has the scheduler enabled. The scheduler system must be online so that it can notify the run system.

5. To save the task in the Task Center, but not actually run the task, click **Save task only**.
6. To save the task in the Task Center and run the task now, click **Save and run task now**.
7. To save the task to the Task Center, and schedule a date and time to run the task later, click **Schedule task execution**. The **Change** push button is enabled.
8. Click **Change**. A window opens where you can type the date and time at which to run the task.
9. The **Details** group box on the scheduling page displays the schedule information you selected.
10. To run a task in the Task Center you must supply a user ID and password. Type the **User ID** and **Password** that you want to use to run the task.
11. You can use the Scheduler Settings page of the Tools Settings notebook to set the default scheduling scheme. If you set a default scheduling scheme, you can still override it at the task level.

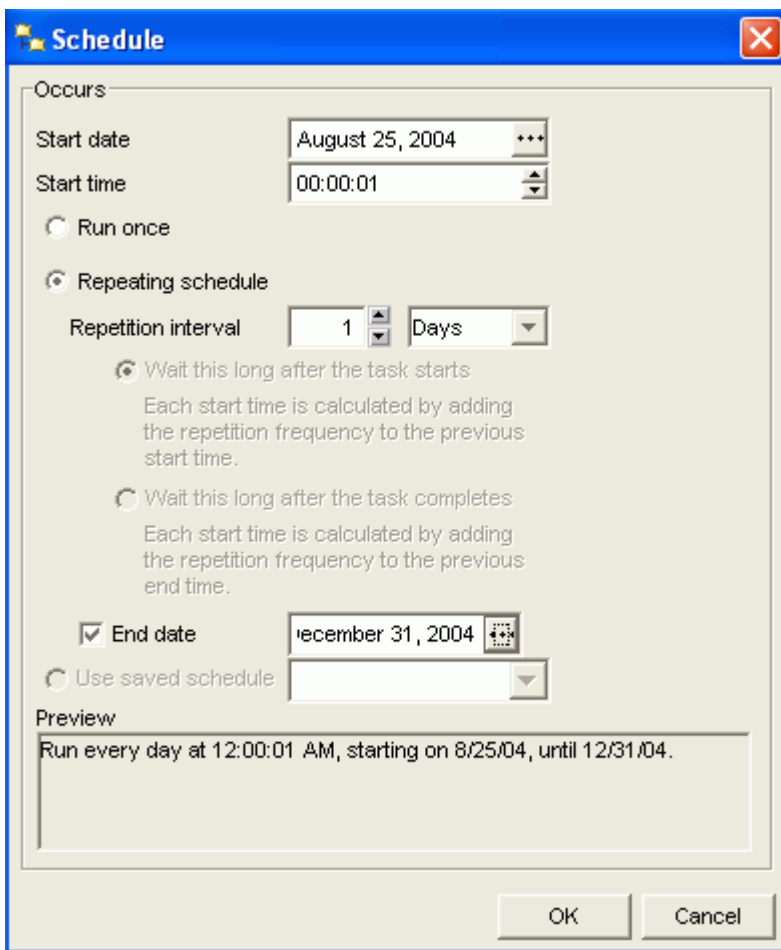


Figure 10. Custom schedule creation

1.7 How do I create a tools catalog when I can't schedule tasks?

When scheduling tasks for the first time, you might run into the following DB2 message:

```
DBA1177N  A database connection to the Tools Catalog
Database "" cannot be made.  SQLCODE "1013"
```

To use the Task Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2 UDB, you can do so from within the Task Center.

Click **Tools > Tools Settings**

When the Tools Settings window opens, select the Scheduler Settings tab.

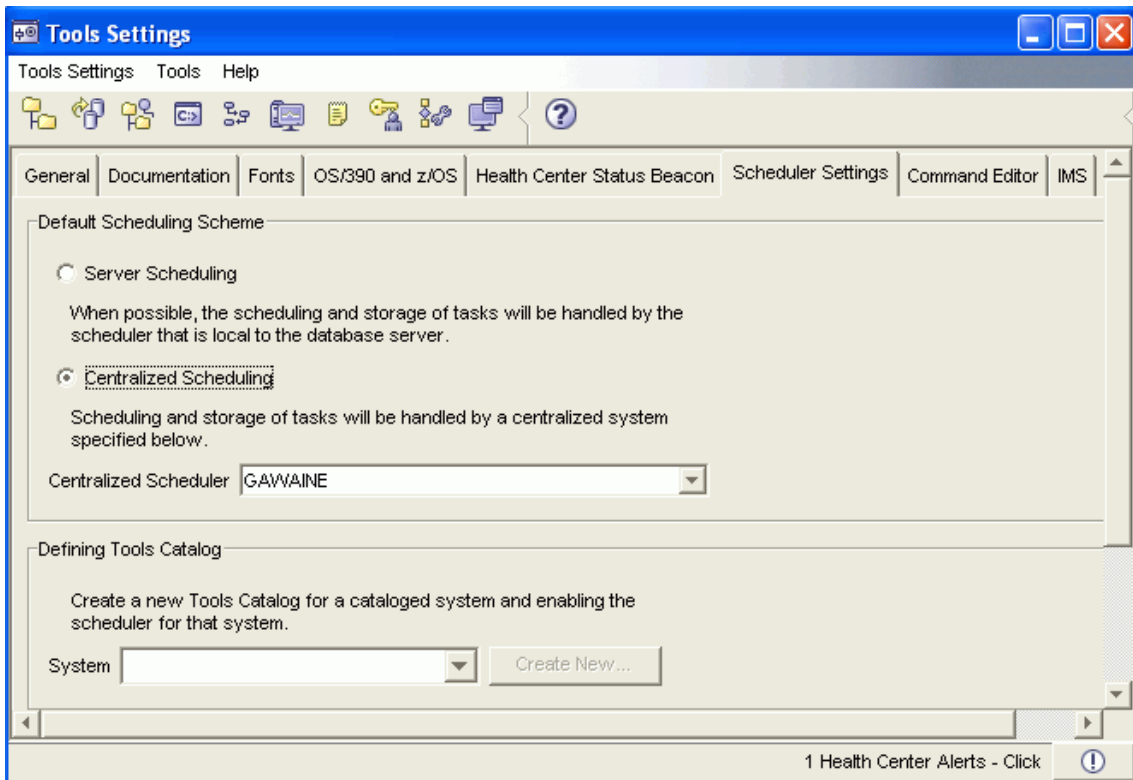


Figure 11. Scheduler Settings tab

Click the **Create New...** push button under **Defining Tools Catalog**. A new window opens that allows you to create a new tools catalog in either a new database or as part of an existing one. Figure 12 illustrates the window through which the new tools catalog is created.

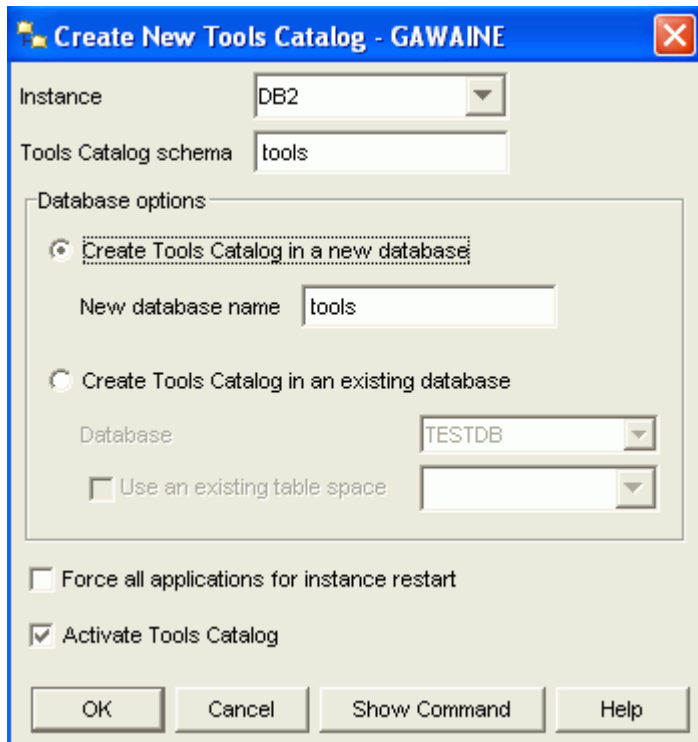


Figure 12. Create New Tools Catalog

When the tools catalog is created, you should be able to schedule tasks through the Task Center and through the scheduling pages of wizards.

1.8 Where can I find a log file that I can send to my application vender when something seems wrong?

The most useful log file for problem determination is the DB2 diagnostic log file (`db2diag.log`). The `db2diag.log` file contains significant diagnostic information, and it should be collected on *all* computers that were involved in processing when the problem occurred.

The location of the `db2diag.log` file is controlled by the DB2 server configuration parameter `DIAGPATH`. The `DIAGPATH` database configuration parameter points to the directory location for placing diagnostic data. Therefore, the diagnostic data directory on your system might be different from the default diagnostic data directory.

- **On Linux operating systems:** The `db2diag.log` file is located, by default, in the `INSTHOME/sqlllib/db2dump` directory, where `INSTHOME` is the home directory of the instance owner.
- **On Windows NT and Windows 2000 operating systems:** The `db2diag.log` file is located, by default, in the `x:\sqlllib\instance` directory, where `x:` is the drive where you installed DB2 UDB or the DB2 Data Links Manager, and `instance` is the name of the instance that is owned by the DB2 database owner or the Data Links Manager Administrator (`dlfm`, by default).

In addition to sending the diag.log file to the application vendor, it is also worthwhile to collect any dump files mentioned in the db2diag.log file. Dump files are named x.nnn, where x identifies the process that produced the dump file, and nnn identifies the database partition number (000 on single-partition databases).

1.9 What are FixPaks? Why and how do I apply them?

DB2 UDB problems you encounter may be due to defects in DB2 UDB itself. To rectify these issues, IBM regularly releases software patches called “FixPaks.” In some cases, FixPaks also include new features.

Installation of FixPaks changes the version number of the installed product. For example, the initial release of DB2 version 8 is 8.1.0; when the first FixPak was applied, the version number incremented one minor number to 8.1.1.

A total of six FixPaks were released between DB2 versions 8.1 and 8.2. Rather than installing version 8.2 from scratch, users of version 8.1 can upgrade to version 8.2 by applying FixPak 7.

FixPaks are available at the following Web site:

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

Click on the link: **Download FixPaks and Clients**

On Windows operating systems, you can check for updates at any time by clicking
Start > Programs > IBM DB2 > Information > Check for DB2 Updates

Recommendation: If you are developing new applications, it is advisable to use the most recently available FixPak.

2. Performance

2.1 I'm experiencing a DB2 UDB performance problem. What should I do?

If you are encountering performance issues with DB2 UDB or with your application that uses DB2 UDB, there are a variety of possible problems and resolutions. The following questions and answers provide a logical diagnostic progression for troubleshooting performance problems.

2.1.1 Is DB2 UDB the problem?

It is always possible that performance degradation is due to the application querying the database rather than because of a DB2 problem. Try running a query against DB2 UDB from outside your application. You can perform this task from the Command Editor, or from the command line.

If the query takes longer than you expect to complete, the performance problem might be due to a problem with DB2 UDB.

2.1.2 Are you receiving health notifications?

DB2 UDB performs health self-monitoring and should notify you if it finds any health issues.

First, you should check to see if there are any unresolved database health problems. Start the Health Center; if there are any health problems, they will be listed there. If there are none listed, then outstanding health problems are not the cause of the database performance degradation.

If there are database health alerts in the Health Center but you were not notified of them, you might want to set up notification. The process for setting up notification is explained in question 3.2.3: "How can I confirm that health notification is working?"

The Health Center can provide you with recommendations about how to solve the problems it has detected. For an explanation, see question 3.1: "How do I get recommendations about database health?"

2.1.3 Has your hardware configuration changed?

If you have added new hardware, such as memory or CPUs, to your system, DB2 UDB may not be taking advantage of the new resources available to it. In order for the new resources to be used, you must alter the DB2 UDB configuration. Doing this manually can require a sophisticated understanding of DB2 parameters, so it is advisable to take advantage of the DB2 Configuration Advisor.

The DB2 Configuration Advisor is a graphical tool that simplifies tuning DBM and DB parameters by identifying the available resources, by asking a few questions, and by making a recommendation on how to tune itself. It is also able to review the hardware available. The Configuration Advisor makes its decision based on heuristics used by expert database administrators.

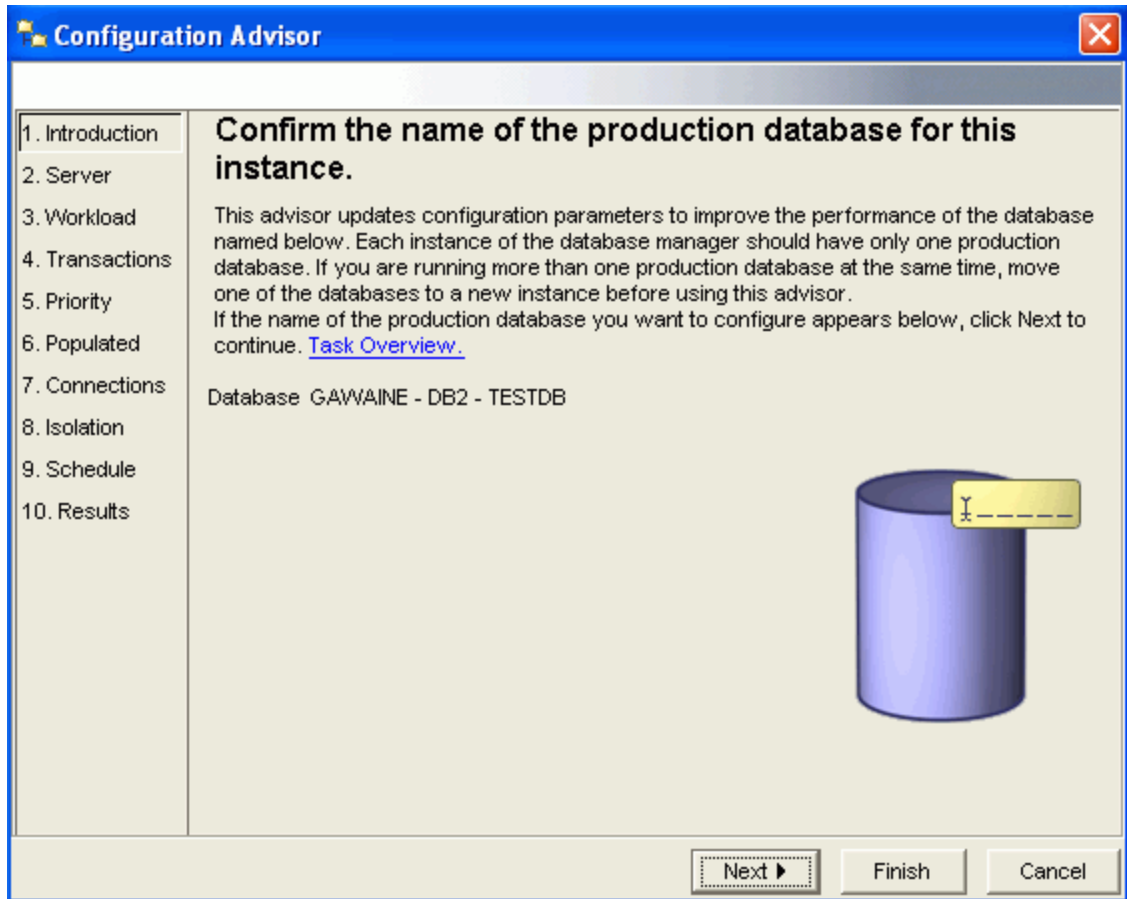


Figure 13. Configuration Advisor

To tune your database configuration for performance:

1. From the Control Center expand the object tree until you find the **Database** object.
2. Right-click the **Database** object and select **Configure Using wizard** from the pop-up menu. The Configuration Advisor opens.
3. Complete each of the applicable wizard pages. Each page is discussed below. The **Finish** push button is available when you complete enough information for the wizard to configure performance parameters for your database.

The wizard pages have the following uses:

Introduction

This page lists the database for which you want to configure performance parameters. Verify that the correct database is shown.

Server

Use this page to specify what percentage of the server’s memory is to be used by the database manager. If other applications are running on the server, set the value to less than 80%.

To avoid conflict with other databases or applications, start with 30%. If DB2 finds that your system has spare resources, it will let you know through the Health Center.

Workload

On this page, indicate the type of workload for which the database is used. You can indicate if your database is used mainly for queries (such as in a data warehousing environment), for transactions (such as order entry), or mixed (for a combination of queries and transactions).

Transactions

Use this page to describe a typical SQL transaction for the database. Indicate whether the average number of SQL statements per unit of work is less than 10 or more than 10. Also indicate approximately how many transactions per minute run on the database.

Priority

On this page, specify the administration priority for the database. If you optimize for faster transaction performance, the database will take longer to recover. If you optimize for faster data recovery time, transaction performance will be slower. If it is equally important to optimize both, you can balance the optimization of the two.

Populated

Use this page to indicate whether or not the database is already populated with data.

Connections

On this page, indicate the average number of local applications and the average number of remote applications that connect to the database. If you are not sure of the numbers, you can use the default values.

Isolation

Use this page to specify the isolation level the database will use when a user either reads or changes data. The isolation levels you can use are repeatable read, read stability, cursor stability, and uncommitted read.

Results

On this page, the wizard lists the recommended configuration parameter settings, based on the information you have provided. You can apply the new values immediately, or you can save them to be applied later.

4. Click **Finish** to make the recommended changes and tune database parameters for performance.

2.1.4 Is there a configuration conflict?

If you have multiple databases on the same server, each in its own instance, you might have conflicting configurations. For example, in Figure 14, a specific database has been allocated 80% of the server's memory. If this has been done for each of the databases on that server, these configurations are in conflict. You should reduce the percentage of memory allocated to each database.

Lowering memory allocation is also useful if you have memory intensive applications other than DB2 UDB running on the same server.

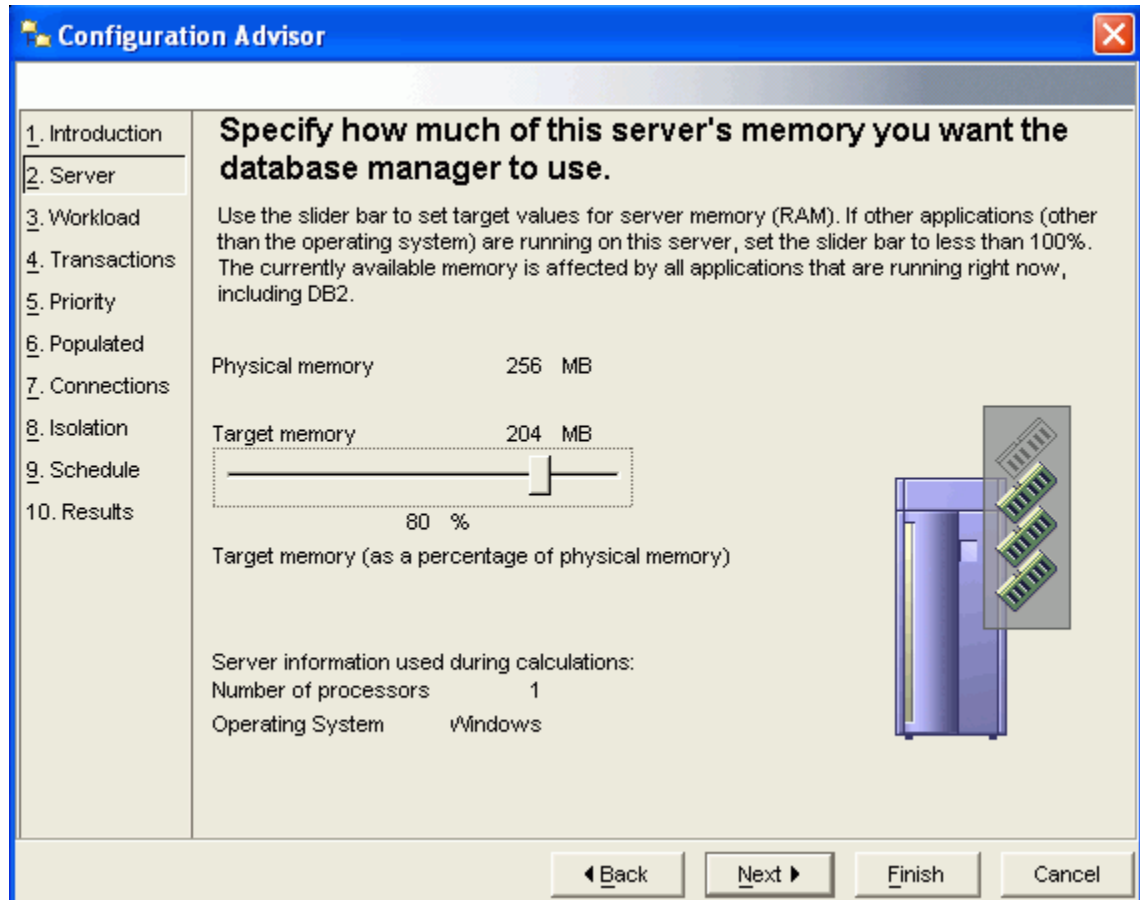


Figure 14. Target memory allocation

2.1.5 Is automatic maintenance turned on and working?

Has the performance degradation that you have noticed happened gradually? If so, there may be a problem with automatic maintenance (or, more simply, automatic maintenance may be turned off).

To see if automatic maintenance is enabled or disabled, select your database in the Control Center's tree, and look at the **Maintenance** line in the Object Details pane. If the maintenance status is **Disabled**, you can enable and set up automaintenance by clicking the **Maintenance** action located in the bottom right corner.

If, as in Figure 15, maintenance is **Fully Automated**, it is still worth confirming that automatic maintenance activities are actually taking place. In order to check automatic activities look in the db2diag.log file (see question 1.8: “Where can I find a log file that I can send to my application vendor when something seems wrong?” for instructions on finding the db2diag.log file). Open the file and look for an entry like the following.

```

2004-07-29-11.44.27.119000-240 I2428H327    LEVEL: Event
PID: 2012          TID: 676          PROD: db2fmp.exe
INSTANCE: DB2     NODE : 000
FUNCTION: DB2 UDB, Automatic Table Maintenance, db2HmonEvalStats, probe:100
START: Automatic Runstats: evaluation has started on database SAMPLE

```

If automaintenance is being carried out as it should be, you should be able to find records similar to the example. If not, there might be something wrong with automatic maintenance. One possibility is that the database manager has stopped due to database inactivity. See question 2.2: “Can performance problems be caused by infrequent database use?” for an explanation of how to activate the database.

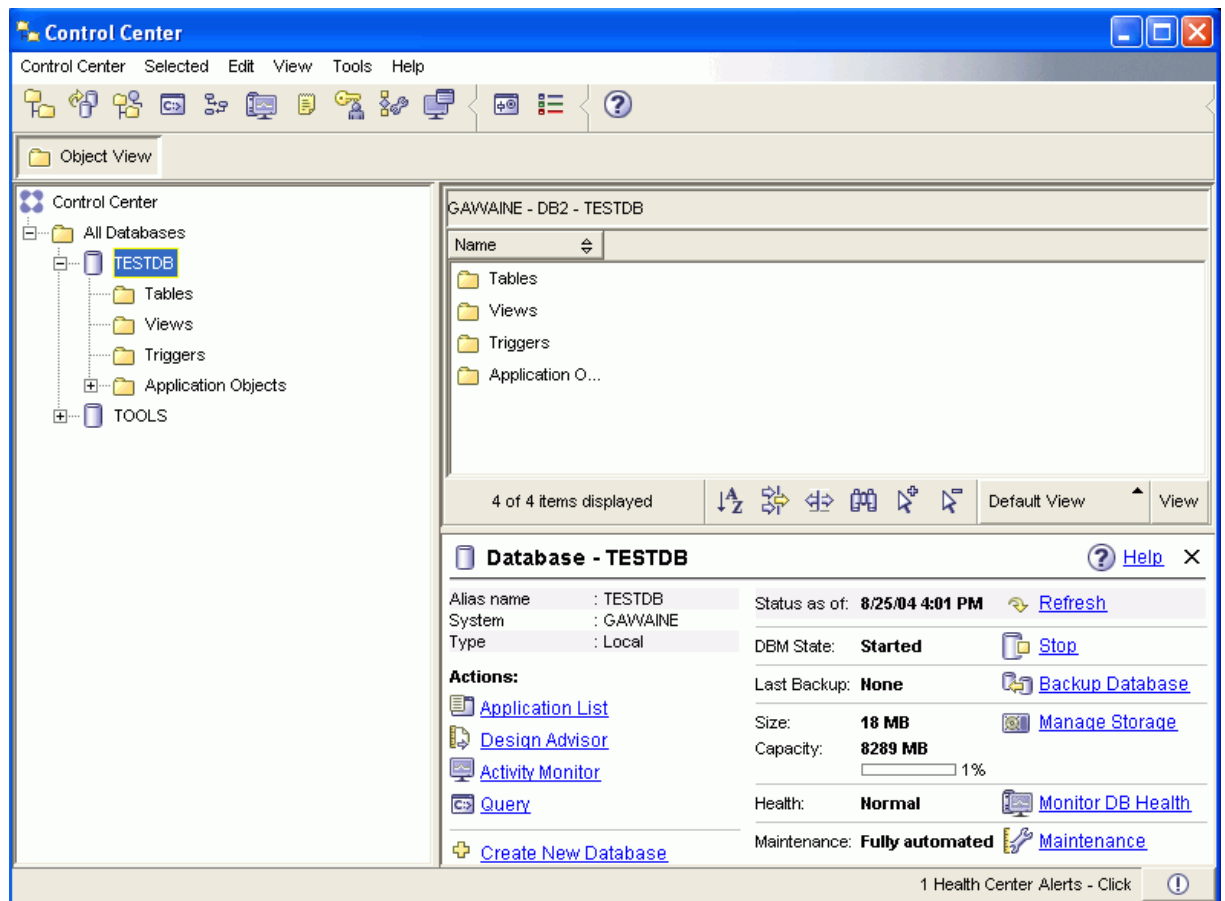


Figure 15. Maintenance: Fully automated

If for some unknown reason, automaintenance is not working, you may want to try to perform some of its typical functions manually. Two of these activities, checking statistics and performing reorganization checks are covered in the next section, 2.1.6.

2.1.6 Have you checked statistics?

Improving performance by keeping database statistics up-to-date

Automatic maintenance keeps up-to-date statistics on your database. You do not need to manually perform this task. However, if you want to manually update the statistics on your database, you can do so by using the **runstats** command.

When DB2 optimizes SQL query plans, its decisions are heavily influenced by statistical information about the size of the database tables and indexes. The optimizer uses this information to estimate the costs of alternative access plans for each query. When significant numbers of table rows are added or removed, or if data in columns for which you collect statistics is updated, the RUNSTATS utility needs to be run again to update the statistics. If current statistics are not available, the optimizer may choose an inefficient access plan based on inaccurate default statistics.

Recommendation: Use the **runstats** command to collect current statistics on tables and indexes, especially if significant update activity has occurred or new indexes have been created since the last time the **runstats** command was executed. This provides the optimizer with the most accurate information with which to determine the best access plan.

To improve performance by keeping database statistics up-to-date:

1. From the Control Center expand the object tree until you find the **Table** object.
2. Right-click the **Table** object and select **Run Statistics** from the pop-up menu. The Run Statistics notebook opens.
3. Complete each of the applicable pages. Each page is discussed below. The **OK** push button is available when you complete enough information to run statistics for the selected tables.
4. Click **OK**.

The notebook pages have the following uses:

Column

This page lets you select the columns to be included in the statistics collection. Select **Collect basic statistics on all columns**.

Index

This page lets you select the indexes to be included in the statistics collection. Select **Do not collect statistics for indexes**.

Options

This page lets you specify the level of access allowed, statistics profiling and throttling. Keep the default setting for Access and Statistics. If you wish to impact performance as little as possible, select **Throttle this utility...** and set the priority to below 50%.

Schedule

Use this page to indicate whether you want to run the backup operation immediately, or if you want to schedule the operation to run later. Recommendation: Schedule statistics collection on all tables at least once a month.

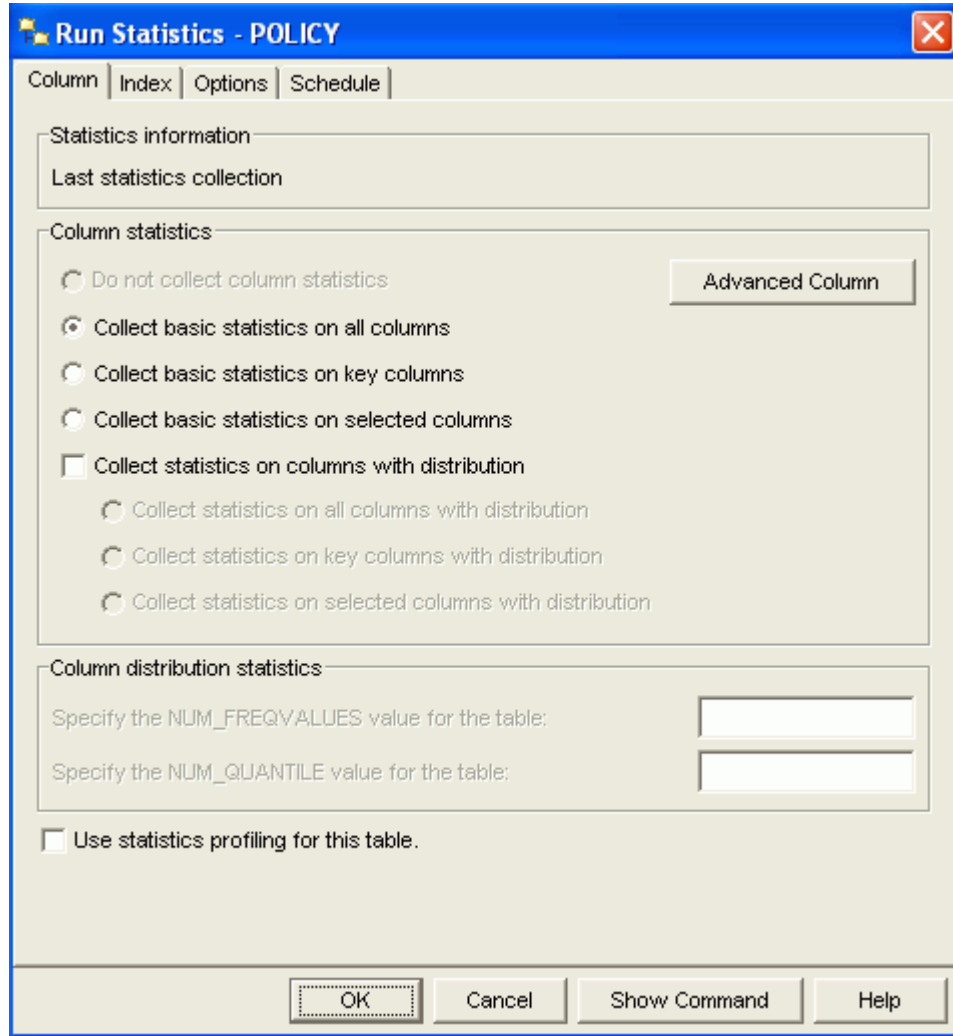


Figure 16. Run Statistics

2.1.7 How can I confirm that my table statistics are up-to-date?

There are several ways to check when statistics were last gathered for a table or index. The **Run Statistics** command includes the exact time the last statistics were collected.

You can also create a report of all the tables in your database, or a single schema, by issuing a simple SQL statement from the Command Center.

To check when statistics were last gathered for all your tables:

1. Open the Command Editor and connect to your database
2. Select a target database to connect to.
3. In the top pane, type the following SQL statement:


```
SELECT TABNAME, TABSCHEMA, STATS_TIME  
FROM SYSCAT.TABLES  
WHERE TABSCHEMA='your_schema_name';
```

4. Click  to execute the statement.

The Results page that is generated includes a report on each table and the time and date statistics were last collected. Figure 17 shows the generated report.

5. To create a similar report for indexes, use the following command:

```
SELECT INDNAME, INDSHEMA, STATS_TIME  
FROM SYSCAT.INDEXES  
WHERE TABSCHEMA='your_schema_name';
```



The screenshot shows the Command Editor interface with a table of results. The table has three columns: TABNAME, TABSCHEMA, and STATS_TIME. The data is as follows:

TABNAME	TABSCHEMA	STATS_TIME
CHECK_CONSTR...	SYSIBM	
COLUMNS	SYSIBM	
COLUMNS_S	SYSIBM	
PARAMETERS	SYSIBM	
PARAMETERS_S	SYSIBM	
REF_CONSTRAINTS	SYSIBM	
REFERENTIAL_CO...	SYSIBM	
ROUTINES	SYSIBM	
ROUTINES_S	SYSIBM	
SQLCOLPRIVILEG...	SYSIBM	
SQLCOLUMNS	SYSIBM	
SQLFOREIGNKEYS	SYSIBM	
SQLPRIMARYKEYS	SYSIBM	

At the bottom of the window, it indicates "129 row(s) in memory".

Figure 17. Table run statistics report

2.1.8 Have you done a reorg check?

Table and index reorganization (reorg)

Automatic maintenance performs table and index reorganization. You do not need to manually perform this task. However, if you want to manually reorganize a table or an index, you can do so.

While table and index reorganization can improve the performance of your database, it is usually not required unless:

- Your database has a very high volume of inserts, updates, and deletes.
- You see a decline in performance and updating table and index statistics does not improve performance.


After many changes to table data, the data may not be efficiently stored in sequential pages on the disk. This forces the database manager to read more pages than necessary to access data. Additional read operations are also required if a significant number of rows have been deleted. In such a case, you might consider reorganizing the table to put the data page in the same order as the index and to reclaim space. (You can reorganize the system catalog tables as well as database tables.)

Note: Because reorganizing a table usually takes more time than running statistics, you should always try to refresh the current statistics for your data before you try to reorganize it. If refreshed statistics do not improve performance, reorganization might help.

Checking the need to reorganize a table or index

Before you reorganize a table or index, you can create a detailed report on each table to see which tables really require reorganization. You can create a new task that runs a reorganization check. The command also updates your statistics, since the reorganization check requires up-to-date statistics.

To create a reorganization check task for your database:

1. From the Control Center window, select the **Task Center** icon 
2. Select **Task > New...** from the main menu bar. The new task window opens.
3. Type **Reorganization Check** into the **Name** field.
4. Type **DB2** into the **DB2 Instance or Partition** field.
5. Select the **Command Script** tab and enter the following command into the large entry field:
`CONNECT TO DATABASE;
REORGCHK UPDATE STATISTICS;`
6. Replace *DATABASE* with the name of your database.
7. Select **OK** to create the new task in the task center.
8. Right-click on the new task in the main Task Center details view.
9. You can choose to run the reorganization check immediately or schedule the task to recur on a regular basis. To run immediately, click **Run Now...** and provide a valid user id and password with sufficient privilege to run the task.
10. Once you run the task, refresh the Task Center details view. You should see a checkmark beside your new task, indicating that the reorganization check is complete.

11. Right-click on the task again and select **Results**.
12. Select the **Output** tab on the **Show Results** window.
13. The Output window contains a details report on each table and index in your database. Tables or indexes containing a * in the REORG column could benefit from a reorganization. An * will appear for each of the three tests conducted on the database that indicate a reorganization is required. More * s indicate a greater requirement for reorganization.
14. If you choose to schedule a reorganization check, you can track the results through the Journal. It keeps track of the results each time the task is run.

2.1.9 Is there a locking problem?

If queries are not returning results or are timing out, and there are multiple applications trying to access the same table, you may have a lock conflict.

Applications place locks on tables and individual table rows so that other applications cannot alter or delete data that the first application is currently accessing – this happens routinely and as a matter of course.

If an application requests a lock on an object that is already locked by another application, the new application will have to wait for the lock to be released before it can obtain a lock and complete its work on the object. The first application is said to be waiting on the lock held by the other application.

There is a limited amount of memory devoted to locking (controlled by the `locklist` database configuration parameter). A lock escalation occurs when the number of locks held on rows and tables in the database equals the percentage of the `locklist` specified by the `maxlocks` database configuration parameter. If the lock list fills, performance can degrade due to lock escalations and reduced concurrency on shared objects in the database.

You can see if there is a lock escalation by clicking on “Application List” in the Object Details pane for the database you are interested in (see Figure 18).

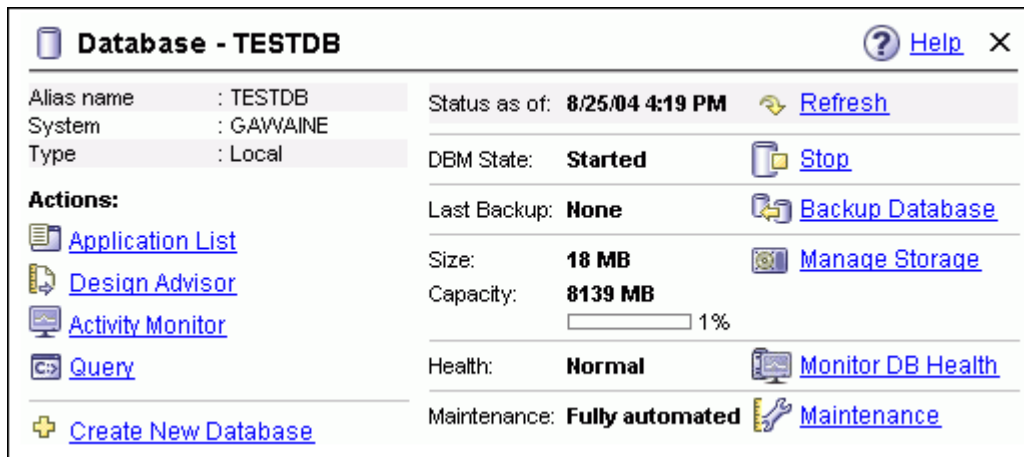


Figure 18. The Application List action in the Object Detail Pane

The Applications window opens. It lists all of the applications with processes involving the given database. If any of them are locked, the Applications window indicates that status in the **Status** column as shown in Figure 19.

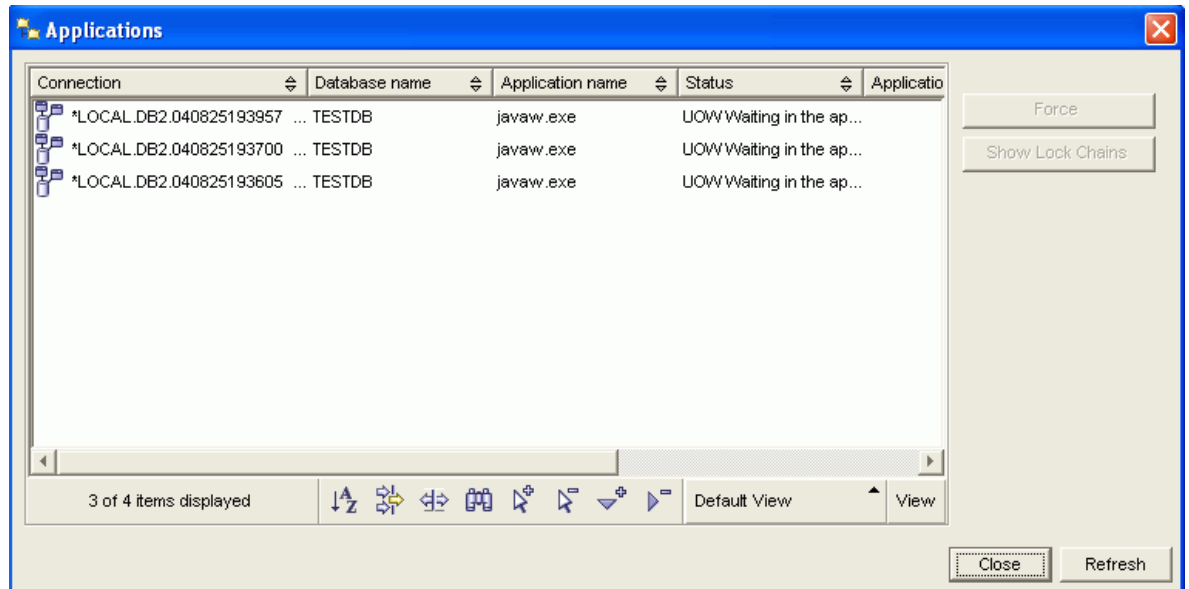


Figure 19. The Applications window

If there are many locks, you can click the **Show Lock Chains** button on the right side of the window to see the succession of locks and the tables involved. If there is, in fact, a locking problem, you can use the “Force” button to cancel the selected connection (application process).

2.1.10 Are there database design issues?

Finally, if there is no other cause for the performance problems that you are encountering, perhaps the design of your database or your use of indexes is not optimized for performance.

If no index exists on a table, a table scan must be performed for each table referenced in a database query. The larger the table, the longer a table scan takes because a table scan requires each table row to be accessed sequentially.

The Design Advisor can help you design and define suitable indexes for your data by helping you:

- Find the best indexes for a problem query.
- Find the best indexes for the set of queries that define a workload, subject to resource limits that are optionally applied.
- Test an index or materialized query table on a workload without having to create the index or materialized query table.

You can launch the Design Advisor from the Control Center by right-clicking on a database object and selecting **Design Advisor**. The Design Advisor is also available in the Actions section of the Object Details pane for any database object.

2.2 Can performance problems be caused by infrequent database use?

Yes. If a database is accessed infrequently, the database manager will likely be stopped. If a database has not been started, and a `CONNECT TO` (or an implicit connect) command is issued in an application, the application must wait while the database manager starts the required database, before it can do any work with that database. However, once the database is started, other applications can connect and use it without spending time on its start up.

You can use the `ACTIVATE DATABASE` command to start up selected databases. This action eliminates any application time spent on database initialization.

```
>>-ACTIVATE---+---DATABASE---+---database-alias----->
                \-DB-----`

>-----+-----+-----><
  \-USER-username---+-----+--`
                    \-USING-password-`
```

Command parameters

database-alias

Specifies the alias of the database to be started.

USER username

Specifies the user starting the database.

USING password

Specifies the password for the user name.

Automatic maintenance activities might not be able to take place if the database manager is stopped. Activating the database can ensure that this potential problem does not occur.

Once you have activated the database, it will stay active indefinitely.

2.3 Is there a simple way to see if scheduled tasks are taking longer to complete than normal?

You can monitor the amount of time it takes to run a task by using features built into the Task Center. You can do this by following these steps:

To create a schedule command:

1. Open the Task Center and choose the relevant Scheduler System.
2. Create a new task

Click **Task > New...**

3. Use the New Task window to create a new Command Script task.
4. Schedule the task to run periodically (every day or every week).

Now that this task is set up and executing periodically, you can refer to it from time to time to check on its performance.

To check on the task performance:

In the Task Center, right-click on this task and select “Show Statistics.” This can also be done from the Journal.

The Statistics window, as shown in Figure 20, opens. From it, you can compare this task’s most recent execution time to it’s fastest and slowest execution times. By checking these statistics regularly, you can get an estimate of the database’s performance fluctuations.

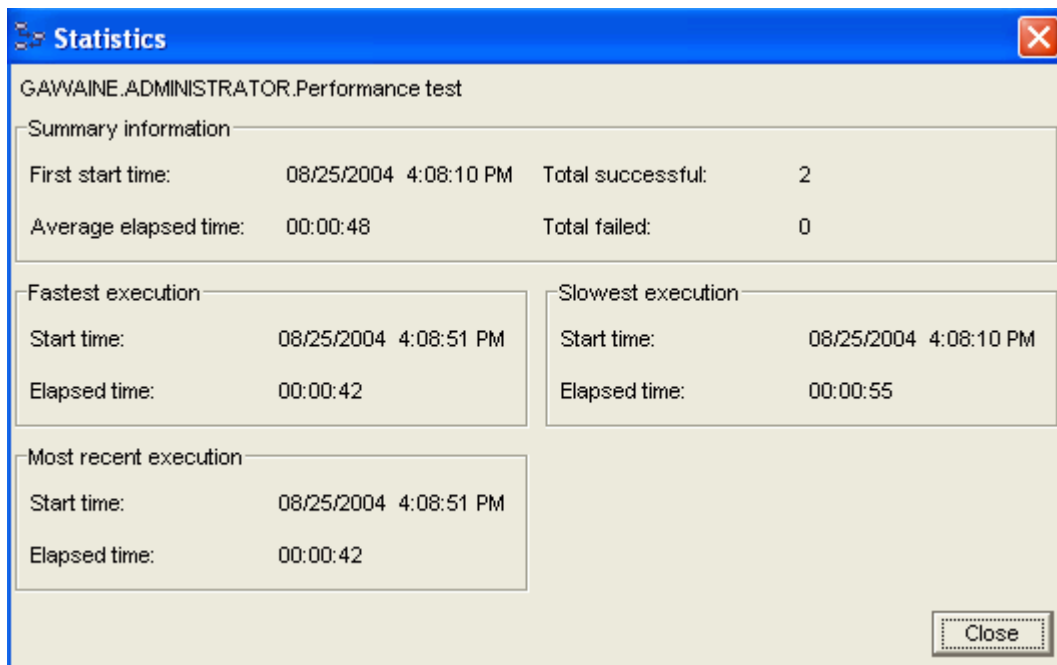


Figure 20. Statistics

2.4 How can I confirm that reorganizations are taking place?

If reorganizations are occurring, there will be a record of them in the Journal. Open the Journal, and select the **Database History** tab. Figure 21 shows the Database History page of the Journal.

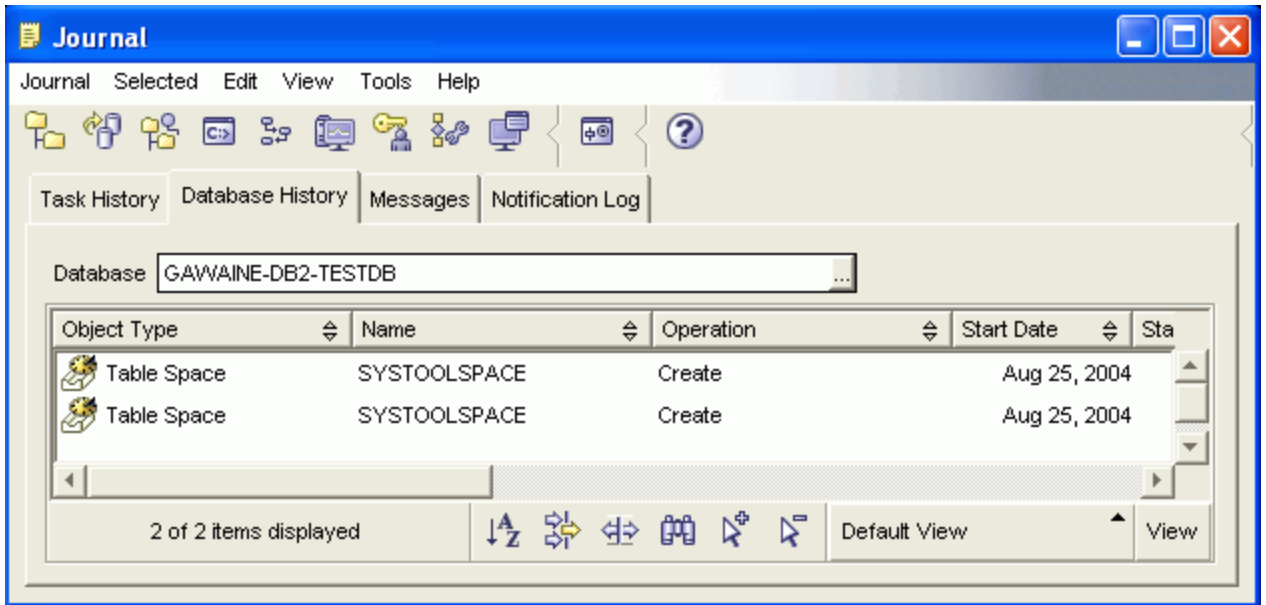


Figure 21. Database history in the Journal

3. Health

3.1 How do I get recommendations about database health?

Recommendations about your database's health are available in the Health Center when a problem has been detected. When this detection occurs, DB2 UDB attempts to notify you in three ways. Two of these methods are in the Control Center; as you can see in Figure 22, there is a health **Warning** flagged with a triangle in the Object Details view. There is also a push button in the lower right corner of the Control Center window frame with the label **1 Health Center Alert – Click** to its left. This button appears in other DB2 Administration tools as well. Clicking either the **Monitor DB Health** link or the alert button in the lower right corner launches the Health Center.

The third method of health alert notification is by Email. For notification to occur, you have to have set up Health Notification.

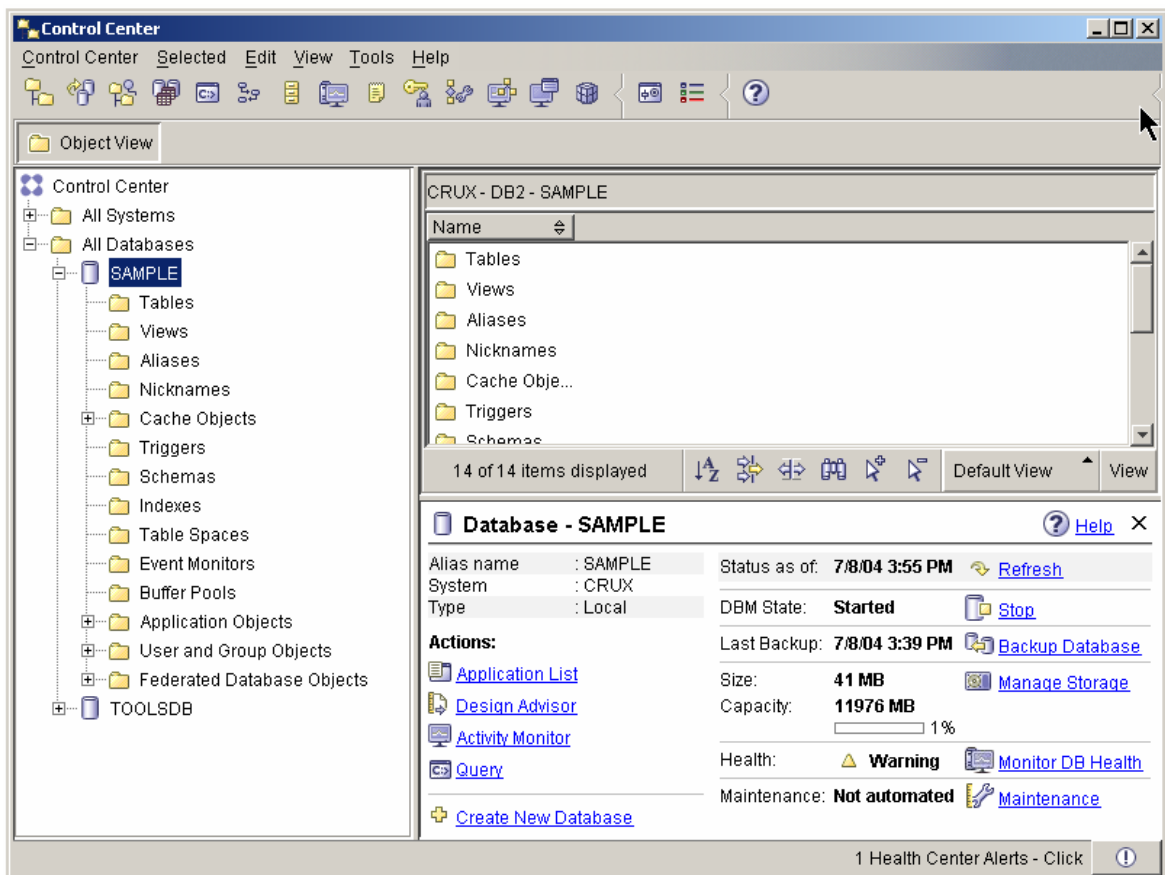


Figure 22. Database health in the Control Center

When you start the Health Center, the Instance object is selected by default, and you are shown a summary of health status for all of the objects within that instance. Figure 23 shows the summary for the DB2 UDB.

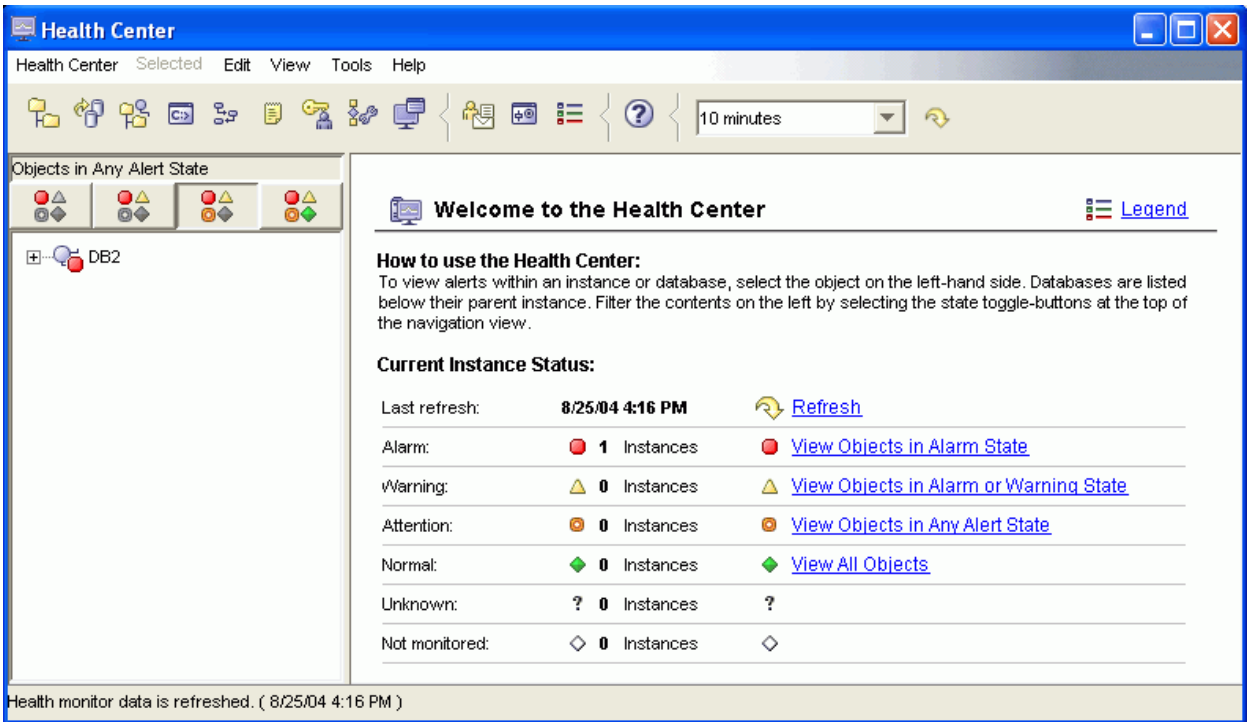


Figure 23. Health Center

Expanding the object tree allows you to expose health alerts for specific database objects. When you find one, right-click on it, and select the **Recommendation Advisor**.

The Recommendation Advisor explains the health alert, and guides you through steps that should help you remedy the problem.

3.2 I haven't been notified of any problems...

3.2.1 How can I confirm that nothing is wrong?

If you haven't been notified of problems, it is entirely possible that there is nothing wrong. However, to confirm that this is the case, you can check by looking at the state listed on the Health line in the Object Details pane (see Figure 24). You can get more detail by starting the Health Center: click **Monitor DB Health** next to the listed Health status.

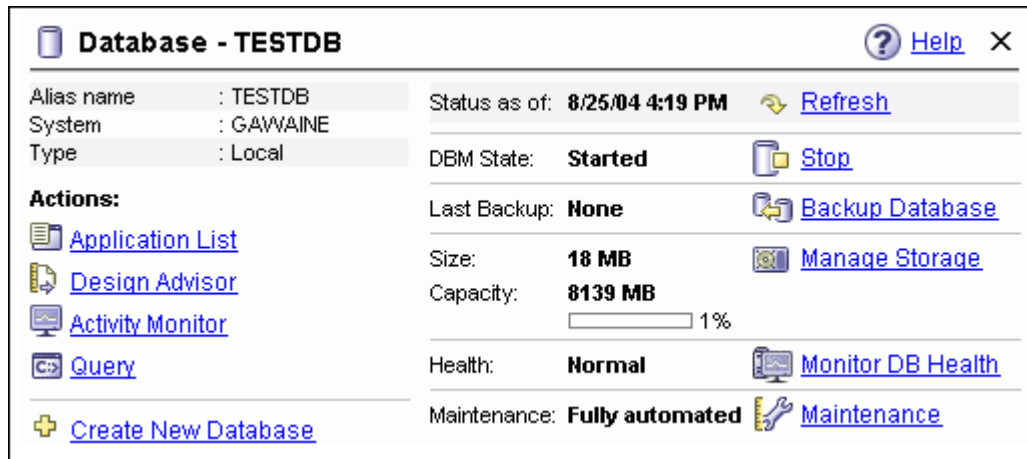


Figure 24. Normal health in the Object Details Pane

If there are problems listed, but you were not notified by Email, there may be a problem with your health notification setup. Refer to question 3.2.3: "How can I confirm that health notification is working?"

3.2.2 How can I confirm that database health is being checked?

If database health is being checked, you should be able to see confirmation of this in the db2diag.log file. Refer to question 1.8: "Where can I find a log file that I can send to my application vendor when something seems wrong?" for instructions on finding this file.

Once you have opened the file, search for an entry like the following:

```
2004-07-29-11.44.27.119000-240 I2428H327          LEVEL: Event
PID: 2012          TID: 676          PROC: db2fmp.exe
INSTANCE: DB2          NODE: 000
FUNCTION: DB2 UDB, Automatic Table Maintenance, db2HmonEvalStats,
probe:100
START: Automatic Runstats: evaluation has started on database SAMPLE
```

If an entry like this is present, database health is being checked.

3.2.3 How can I confirm that health notification is working?

If you are supposed to be receiving health alert notifications by email, and you are unsure of whether they are being sent properly, open the Health Center and click **Health Center > Configure > Alter Notification**.

The **Configure Health Alter Notification** window will open.

Click **Manage Contacts**. The Contacts window will open (see Figure 25).

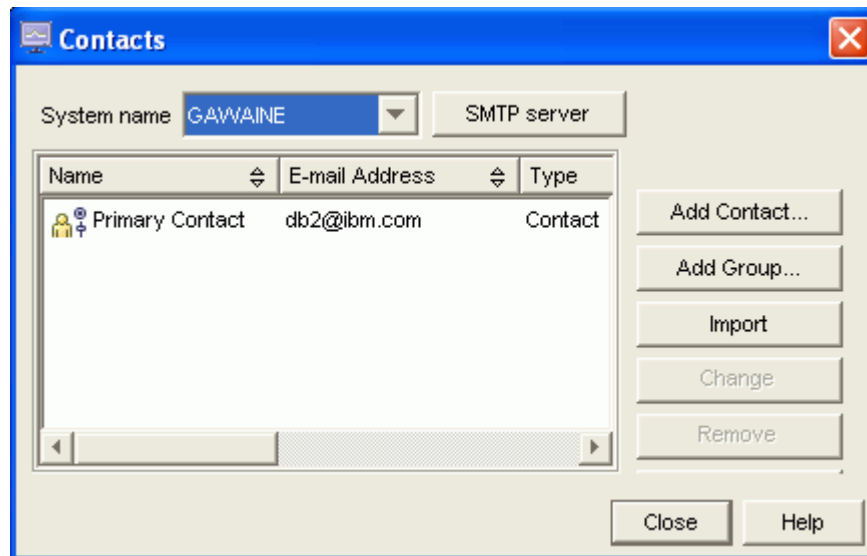


Figure 25. Contacts window

Select the contact about which you are concerned, and click the **Change** push button. The Change Contact window will open.

In this window, click **Test**. A test Email message will be sent to the Email address listed for the selected contact.

If the email does not arrive, return to the Configure Health Alert Notification window, select the problematic contact, and click **Troubleshoot**. This action will lead you to a tool that attempts to help diagnose and rectify the problem.

4. Storage and recoverability

4.1 What kinds of failures do I need to protect my database from?

A database can become unusable because of hardware or software failure, or both. You may, at one time or another, encounter storage problems, power interruptions, or application failures, and each type of failure requires a different recovery action. Protect your data against the possibility of loss by having a well-rehearsed recovery strategy in place. Some of the questions that you should answer when developing your recovery strategy are:

- Will the database be recoverable?
- How much time can be spent recovering the database?
- How much time will pass between backup operations?
- How much storage space can be allocated for backup copies and archived logs?

4.2 What kind of protection do I get with basic automatic backup?

If you create a database using automatic maintenance a backup is taken at least once a week. If you want to back up more often (up to once a day with automatic maintenance), use the Configure Automatic Maintenance wizard to change your backup settings. You can also change the location of your backup image and switch automatic backup to online backups using archive logging using the wizard. Simply select the Configure Settings button for Backup on the Configure Maintenance Activity page of the wizard (see Figure 26).

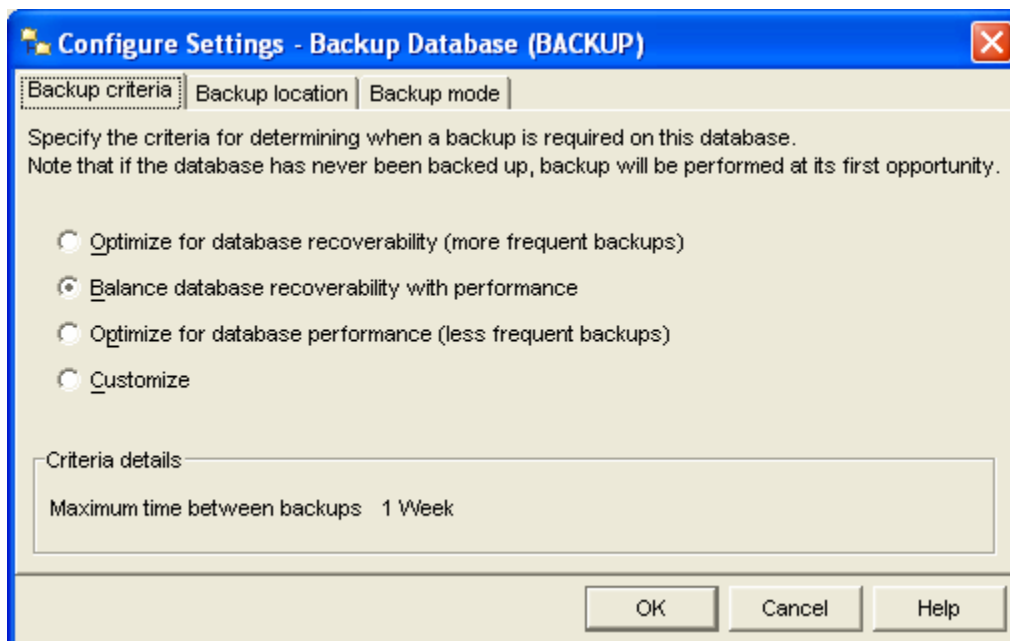


Figure 26. Backup criteria

4.3 Where can I find a record of my backups?

The *recovery history file* contains a summary of the backup information that can be used to determine recovery options, if all or part of the database must be recovered to a given point in time. This file is used to track recovery-related events such as backup and restore operations, among others. This file is located in the database directory.

You can access the recovery history file through the Journal.

4.4 Where does automatic maintenance place my database backup images by default?

Automatic backup places your backup files in the following directory:

```
<DFT_DB_PATH>\NODEXXXX\SQLXXXXX\DB2AUTOBACKUPS
```

where <DFT_DB_PATH> is your database directory. Also by default, automatic backup keeps only two backup files in this location. Older backups are automatically pruned. You can change the location of your backup image using the configure Automatic Maintenance wizard. Simply click the **Configure Settings** push button for Backup on the Configure Maintenance Activity page of the wizard.

Remember, this is most likely the same physical disk that is storing your database. If your hard drive crashes you may lose both your database and your backup. It is safest to override the default and specify a backup location on a separate physical drive.

4.5 What exactly is included and not included in a DB2 UDB backup?

The concept of a database *backup* is the same as any other data backup: taking a copy of the data and storing it on a different medium in case of failure or damage to the original. The simplest case of a backup involves shutting down DB2 UDB, to ensure that no further transactions occur, before you backup your entire system. Using the backup you can then rebuild your entire system if it becomes damaged or corrupted in some way. However, the recommended and supported approach is to perform an explicit backup of the database and then save the backup file either separately or include it in a full system backup. This approach ensures that the database is shut down properly first, all the files that are part of the database are included in the backup file, and the backup is recorded in the database history to simplify recovery later. Your overall strategy should also include procedures for recovering command scripts, applications, user-defined functions (UDFs), stored procedure code in operating system libraries, load copies and any other files that are not part of the database itself.

4.6 What other options do I have when it comes to backing up and restoring my database?

Backing up your database ensures that you can recover your critical data to the point in time of the last full backup. Log files record changes to the data in your database. Basic or circular logging allows you to recover your data to the last complete transaction in case of a system crash. Advanced or archive logging allows you to recover to the last complete transaction even in the case of a hard drive failure. It also allows you to perform your backups online without having to stop the database.

4.7 Why should I care about logging and log files? With my old database, I didn't even have logs.

Very simple database programs write immediately to disk whenever someone makes a change. This can be very slow and cause problems if there are a lot of users. Therefore, advanced database systems keep the changes in memory until it is appropriate to write them to disk. This improves performance and reduces contention between users. However, the database needs to record the changes in case there is a system crash before it can write everything out to disk. Basic or circular logging takes care of this.

By default, DB2 UDB takes care of basic log setup and will recover from a system crash, as long you did not experience a disk failure.

4.8 What is the difference between circular (default) and archive (advance) logging?

Each database includes *recovery logs*, which are used to recover from application or system errors. In combination with the database backups, they are used to recover the consistency of the database right up to the point in time when the error occurred.

Circular logging maintains a rotating set of log files. When the last log file is full, the first log file is reused and is overwritten. This provides sufficient protection against a system crash but not against a disk failure. In the case of a system failure, you can recover any changes to your data that were still in memory by replaying the changes recorded in your logs. However, since they are circular and can be overwritten, there is a limit to the number of changes that can be recorded.

To guard against a full data disk failure, archive logging writes new log files until you complete your next backup. In this way, you have a complete record of every change to the database made since your last backup. In the case of a disk failure, you can restore the last backup and then replay the logs right up until the last complete transaction before the failure. The additional cost comes in managing the log files. Ideally, they are continuously copied to a different location from your data.

4.9 How do I switch from the circular (default) to archive (advanced) logging?

You can use either the Configure Database Logging wizard or the Automated Maintenance wizard to configure data logging options for a database.

To open the Configure Database Logging wizard:

From the Control Center window, expand the object tree until you find the database for which you want to set up logging.

1. Right-click on the desired database and select Configure Database Logging wizard from the pop-up menu. The Configure Database Logging wizard opens.
2. The first page, as seen in Figure 27, allows you to switch between circular and archive logging. Subsequent pages allow you to specify the location of your log files, set up automatic archiving (moving to a safe location) of your log files, and scheduling the backups to go with log archiving. The **Finish** push button is available when you complete enough information for the wizard to configure logging options for your database.

Note: If you require an initial backup of your database, you can use the wizard to configure backup options as well.

To open the Configure Database Logging wizard:

From the Control Center window, expand the object tree until you find the database for which you want to set up logging.

1. Right-click on the desired database and select Configure Database Logging wizard from the pop-up menu. The Configure Database Logging wizard opens.
2. The first page allows you to switch between circular and archive logging. Subsequent pages allow you to specify the location of your log files, set up automatic archiving (moving to a safe location) of your log files, and scheduling the backups to go with log archiving. The **Finish** push button is available when you complete enough information for the wizard to configure logging options for your database.

Note: If you require an initial backup of your database, you can use the wizard to configure backup options as well.

3. Click **Finish** to configure logging options for your database.

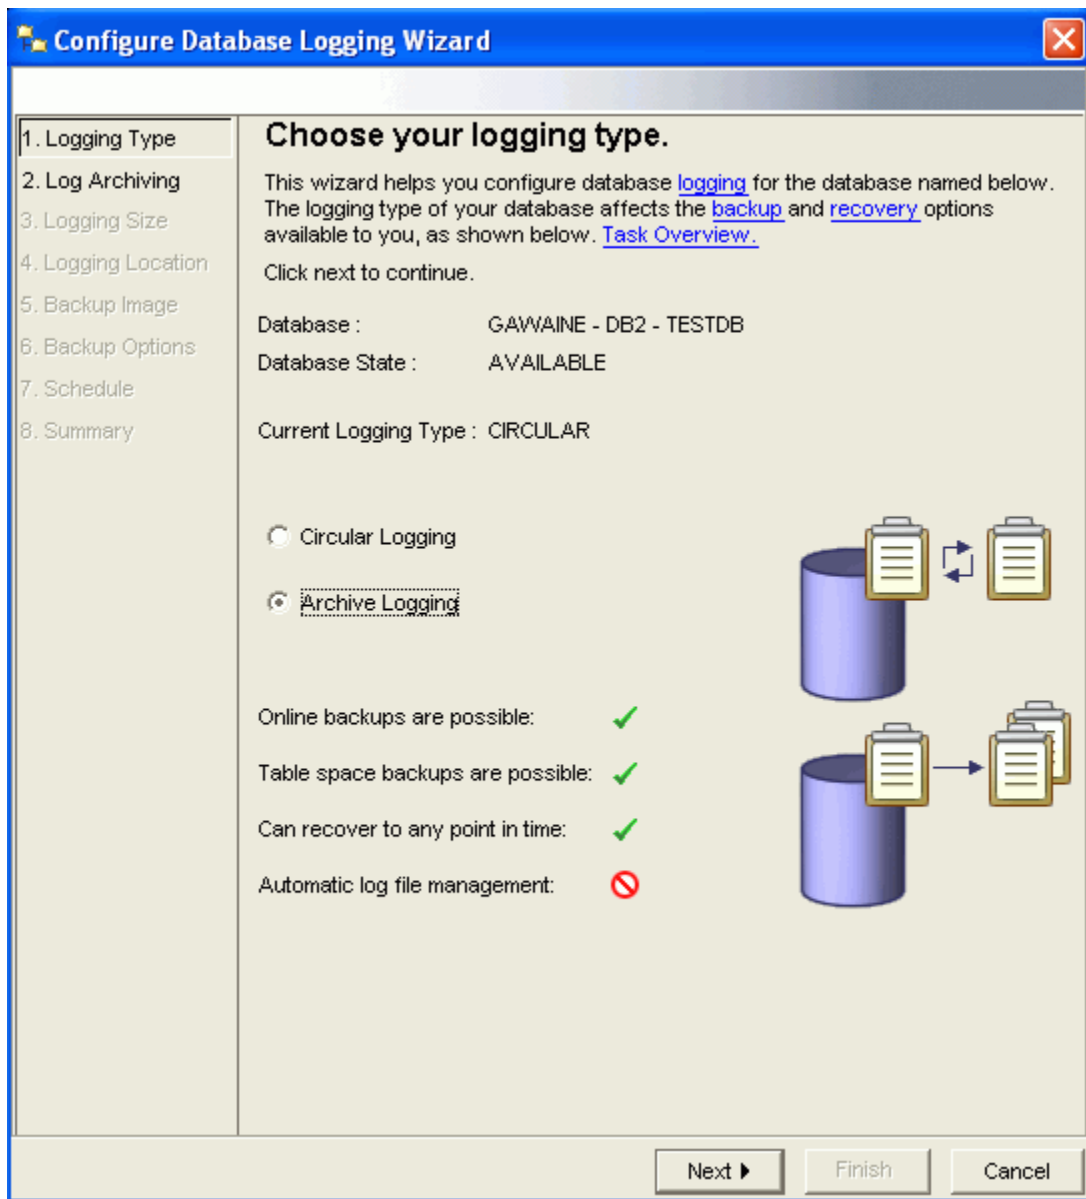


Figure 27. Archive logging selection

4.10 I am processing a large amount of data in my database and want to take a backup in case anything goes wrong. How can I take a manual backup before the automatic maintenance runs?

To back up the objects in your database manually:

1. From the Control Center, expand the object tree until you find the database object that you want to back up.
2. Right-click on the object and select **Backup** from the pop-up menu. The Backup wizard opens.
3. Complete each of the applicable wizard pages. Each page is discussed in the following text. The **Finish** push button is available when you complete enough information for the wizard to back up the objects in your database.
4. Click **Finish** to back up the objects in your database.

The wizard pages have the following uses:

Introduction

This page lists the object that you have selected to back up. Verify that it is the correct one. If you are using archive logging, you can choose to back up the entire database or to back up only selected table spaces.

Note: Backing up all table spaces is not equivalent to backing up the entire database.

Table Spaces (Advanced)

Your database storage can be divided into individual table spaces. Backing up by table space gives you additional flexibility in backing up or restoring only part of the database at one time. This page is available only if you are in archive logging mode and you have chosen to back up selected table spaces. Use this page to specify the table spaces you want to back up. All non-temporary table spaces in your database are eligible to be backed up, and are listed in the table to the left of the page. If you started the backup wizard from one or more table space objects, they will be pre-selected in the table to the right of the page.

Image

Use this page to specify where you want to store your backup image. Indicate the media type you want to use and information about the particular media type.

Note: Do not use the same directory that you use for automatic backups. Backup images are automatically pruned from the automatic backup directory.

Options

Use this page to specify whether you want to back up all data, back up data that has changed since the last full backup, or back up data that has changed since the last backup of any type. You can also specify whether you want the backup to occur while the database is online or offline. If the backup will run offline, you can indicate that you want to run a quiesce database command before the backup operation runs.

Note: If you are using circular logging, you can perform only an offline backup.

Performance (Advanced)

On this page, the wizard suggests recommended values for parallelism, the number of buffers, and the size of each buffer. You can modify these values. If you change the values on this page, you should click **Recommend** to have the wizard recalculate the recommended values.

Schedule

Use this page to indicate whether you want to run the backup operation immediately, or if you want to schedule the operation to run later.

Summary

This page allows you to review the decisions that you made for your backup operation and to view the command that will back up the objects in your database.

Backup Type

- Full Backup - Backup all data
- Incremental - Back up data that has changed since the last full backup
- Delta - Back up data that has changed since the most recent backup

Availability

- Online - Users can access the database during the backup
 - Include log files in backup image
- Offline - Users cannot access the database during the backup

Quiesce Database

- Quiesce the database before initiating the offline operation

Throttle

- Throttle this utility to regulate the performance impact on the database workload.

Set Priority...

Figure 28. Backup options

4.11 How can I take advantage of online backups and point-in-time recovery even after a disk crash when my log files are taking up a lot of room on my primary disk?

DB2 UDB can automatically move your archived log files to a different location after each one is closed. This gives you the extra security of moving your log files to a separate location and ensures that you won't run out of space where you keep your currently active log file. (If you run out of space for your active log file DB2 UDB will stop.)

If you select log archiving in the Configure Database Logging wizard, an additional page called Log Archiving is added to the wizard. You can select to **Use DB2 to automatically archive the log files**. This moves your log files out of the main log file location into the archive location as soon as they are closed. You can also specify a secondary archive location in case there is a problem with your primary location.

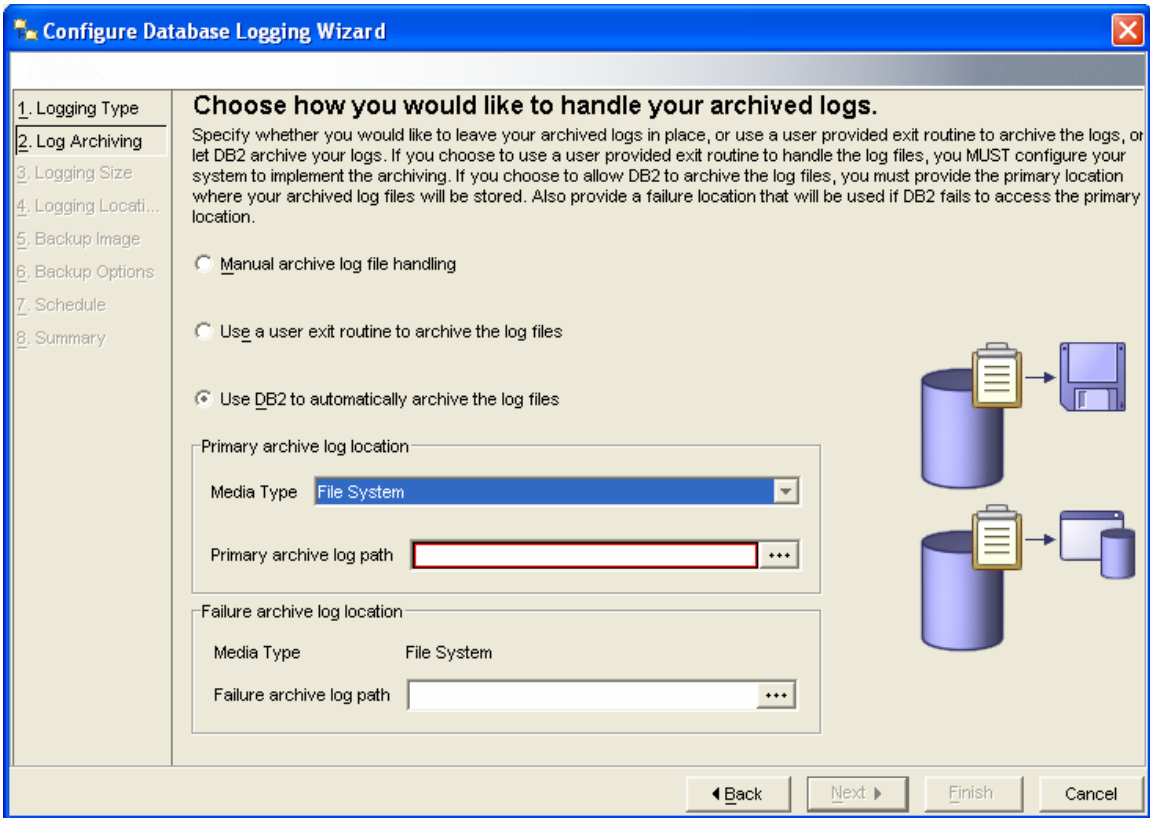


Figure 29. Archived logs management

4.12 How do I make sure that there are no users accessing the database before I back it up or perform manual maintenance?

Offline maintenance requires that no users are accessing the database. Quiesce forces all users off the specified database and puts it into a quiesced mode. In quiesced mode, users cannot connect from outside of the database engine. While the database instance or database is in quiesced mode, you can perform administrative tasks on it.

Quiesce is automatically added to the scripts created by the Backup and Logging wizards.

After administrative tasks are complete, use the UNQUIESCE command to activate the instance and database and to allow other users to connect to the database without having to shut down and perform another database start.

To manually quiesce a database:

1. From the Control Center window, expand the object tree until you find the database that you want to quiesce.
2. Right-click on the desired database and select **Quiesce** from the pop-up menu. The database is quiesced immediately.

4.13 What is best way to ensure that I never lose my log files?

Log mirroring is the best way to make sure you never lose your log files. Mirroring log files helps protect a database from:

- Accidental deletion of an active log file.
- Data corruption caused by hardware failure.

If you are concerned that your active log files may be damaged (as a result of a disk crash), you should specify a secondary path for the database to manage copies of the active log.

Recommendation: Place the secondary log file path on a physically separate disk, preferably one that is also on a different disk controller. That way, the disk controller cannot be a single point of failure.

When you first turn on log mirroring, it will not actually be used until the next database startup.

If there is an error writing to either the active log path or the mirror log path, the database marks the failing path as bad, writes a message to the administration notification log file, and writes subsequent log file records to the remaining good log path only. DB2 does not attempt to use the bad path again until the current log file is completed. When DB2 needs to open the next log file, it verifies that this path is valid, and if so, begins to use it. If not, DB2 will not attempt to use the path again until the next log file is accessed for the first time. There is no attempt to synchronize the log paths, but DB2 keeps information about access errors that occur, so that the correct paths are used when log files are archived. If a failure occurs while writing to the remaining good path, the database shuts down.




Figure 30. Log file mirroring for extra security

4.14 How do I recover my database on a different computer?

If you want to restore from a backup, all you need is the backup file. You can restore on a different computer as long as it is running the same or higher level of DB2 UDB on the same operating system. If you want to recover to the last transaction before your system failed, you must have been using archive logging and have the complete set of log files.

You will not have your database history file on the new computer. This file has the records of each backup that you made of your database and where they are stored. With the History file in place, you can go to the Journal and select the backup you want, right-click and select **Restore**. Figure 31 shows the Journal page that you use. You can also select **Restore** from the database object in the Control Center. The second page of the Restore wizard lets you pick the backup you want to use to restore your database.

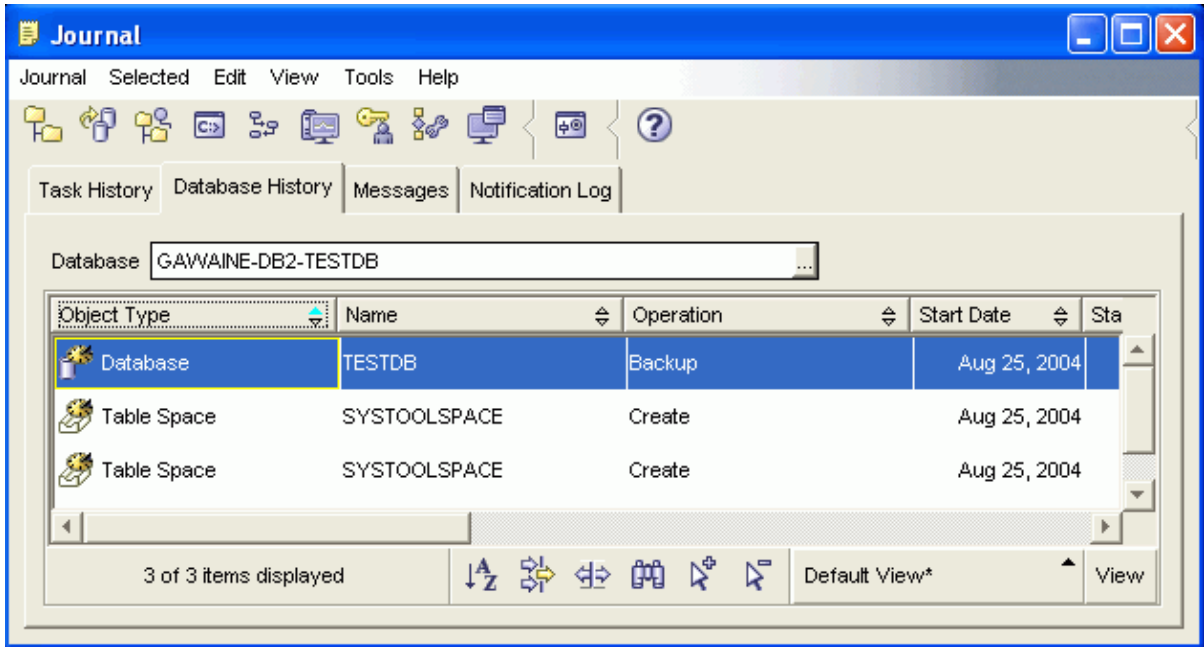


Figure 31. Database restoration using backup

If you want to create a new database from a backup image, select **Create Database** from the **All Databases** folder. Then select **From Backup...** The Backup wizard lets you create a new database with a new name from an existing backup image.

The complication occurs when you don't have a database history file. The Control Center only gives you the option of manually entering the backup image information. On Windows operating systems this means that you have to point to the backup file and specify the time and date when the backup was taken.

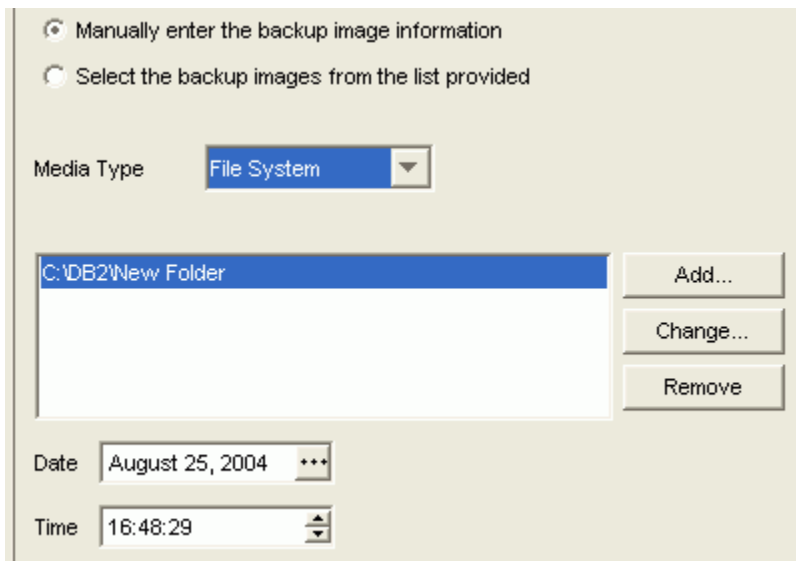


Figure 32. Entering backup image information manually

These actions are necessary because DB2 stores the backup image in a directory structure that includes the name of the database, as well as the date the backup image was created. The file itself is named after the time of the backup. For example, `D:\DB2\Backup\TEST.0\DB2\NODE0000\CATN0000\20040618\161432.01` is the first image created on June 18th, 2004 at 4:17:32 PM. To correctly point at this image, you need to fill out the Restore wizard using the `D:\DB2\Backup` directory and specify the correct date. The best way to handle this is to copy the entire backup directory onto the target computer, not just the file, and specify the right date and time in the wizard.

If you no longer have the directory structure and don't know the date and time the backup occurred, you can still restore the database from the file. However, you need to do it from the command line. First, place the backup file in a new directory. For example, `e:\db2\backup\newmachine`. Then issue the following command from the DB2 command line:

```
DB2 RESTORE DB TEST FROM e:\db2\backup\newmachine TO "D:" INTO NEWDB
```

If your backup file is the only file in this directory, DB2 UDB assumes it is the correct file and restores the database on your D: drive. The new database will be called "NEWDB".

5. Security

5.1. What is the difference between authentication and authorization?

To protect data and resources associated with a database server, DB2^(R) Universal Database uses a combination of external security services and internal access control information. To access a database server, you must pass some security checks before you are given access to database data or resources. The first step in database security is called authentication, where you must prove that you are who you say you are. The second step is called authorization, where the database manager decides if the validated user is allowed to perform the requested action, or access the requested data.

Authentication of a user is completed using a security facility outside of DB2^(R) Universal Database (DB2 UDB). The security facility can be part of the operating system, a separate product or, in certain cases, may not exist at all.

On the other hand, authorization is performed using DB2 facilities. DB2 tables and configuration files are used to record the permissions associated with each authorization name. The authorization name of an authenticated user, and those of groups to which the user belongs, are compared with the recorded permissions. Based on this comparison, DB2 decides whether to allow the requested access.

5.2 How do I add a user?

You can add users through the Control Center.

To add users:

1. Open the Control Center
2. Users are created for specific databases. Find the database for which you want to create a user, and expand the tree to see its constituent folders.
3. Expand the **User** and **Group Objects** folder; select the **DB Users** object.

If the **Users** and **Groups** folder is not visible, you are probably viewing the Control Center in Basic mode. You need to switch to Advanced mode in order to access the Users and Groups folder. To access this folder, click **View > Customize Control Center** from the menu bar, or click the **Customize Control Center** action in the Object Details pane at the root level of the Control Center's navigation tree. Figure 33 illustrates the Control Center.

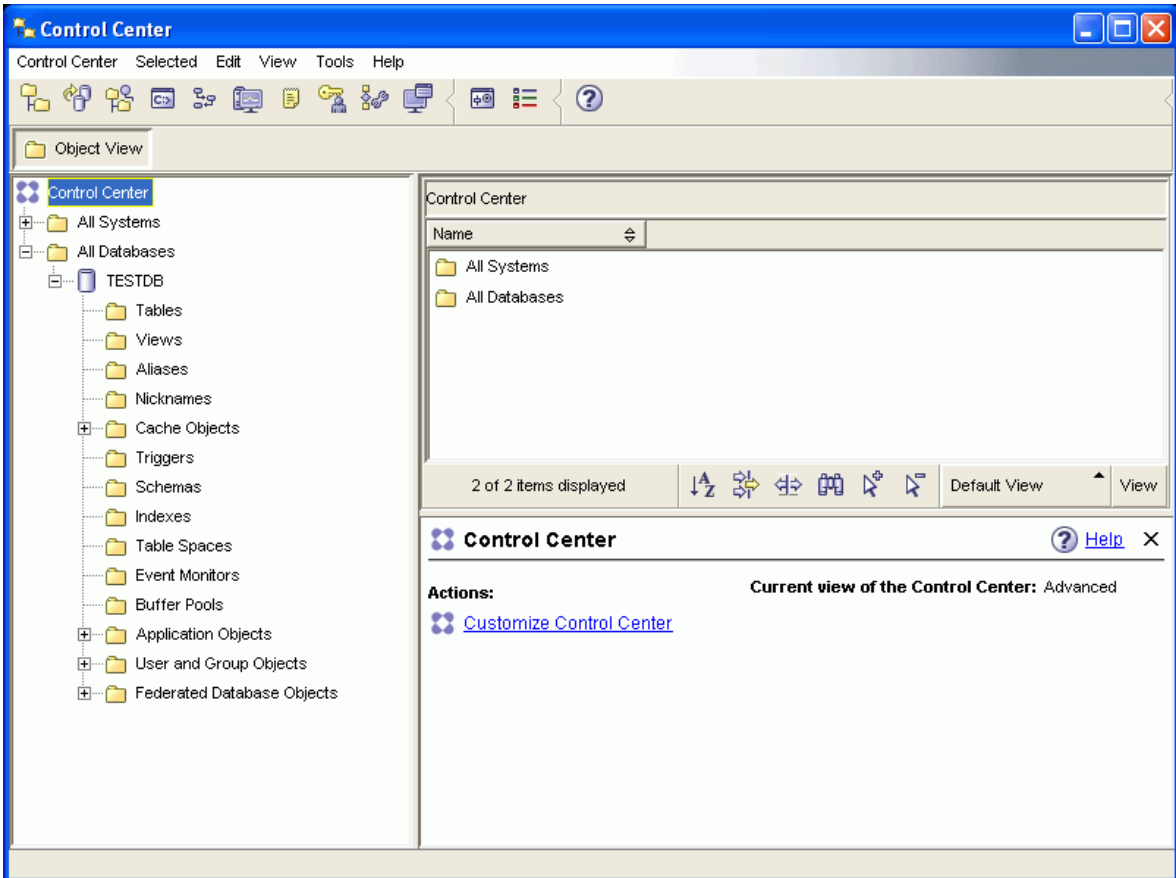


Figure 33. Control Center in Advanced mode

In the window that opens, switch to Advanced mode and click **OK**. The User and Group Objects folder should now be available.

4. From the Object Details pane, click the **Add New User** action. Alternatively, right-click on the **DB Users** object in the tree and select **Add...** Figure 34 illustrates the **Add New User** action.

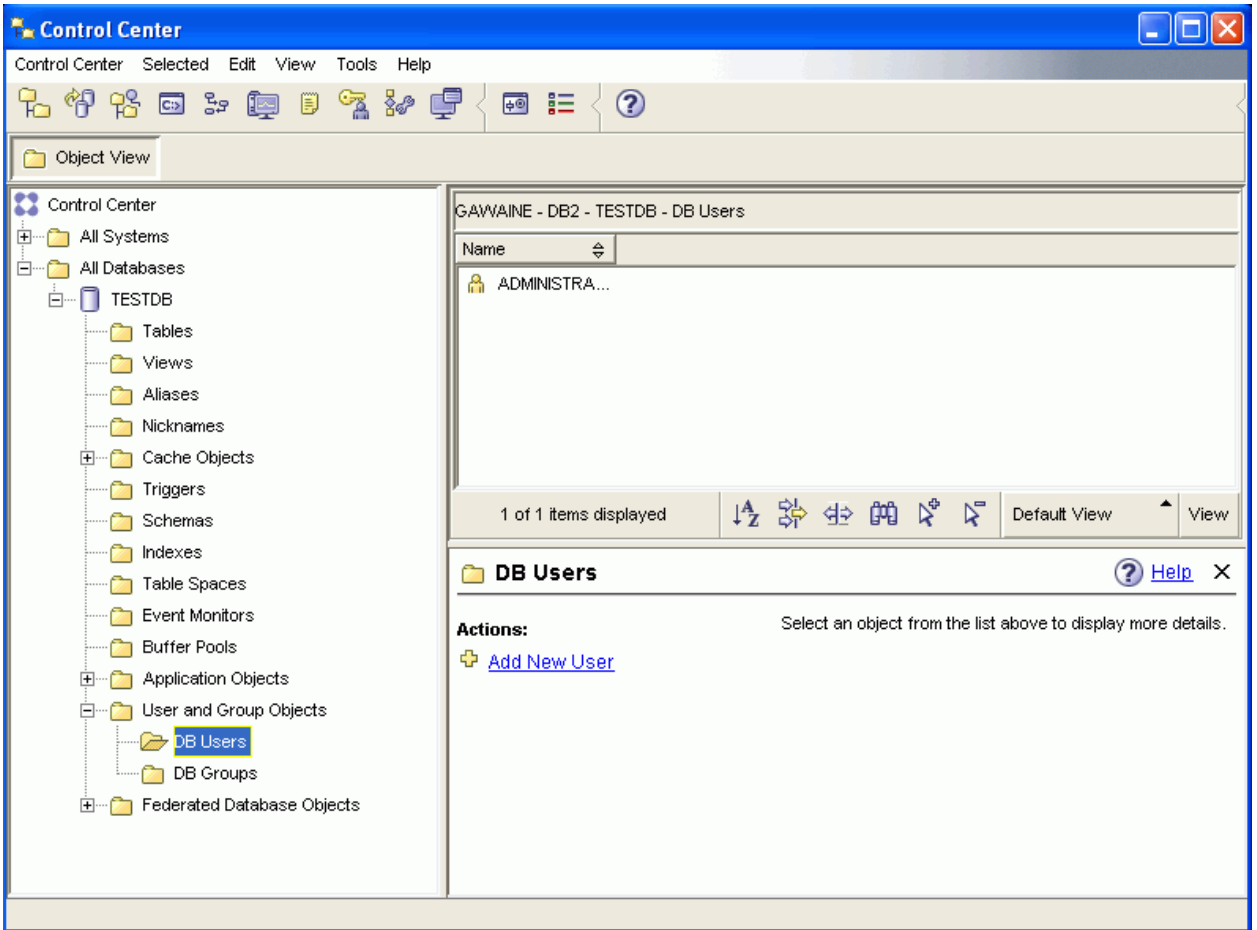


Figure 34. The Add New User action

5. The Add User window opens. On the first tab, marked **Database**, select the user you want to add from the **User** drop-down list. This is a list of all the operating system users on your local system (ie., Windows users or Linux users). When you have selected a user, you can give privileges to this user by selecting the appropriate check boxes under **Authorities**. Figure 35 provides an illustration of the Add User window.

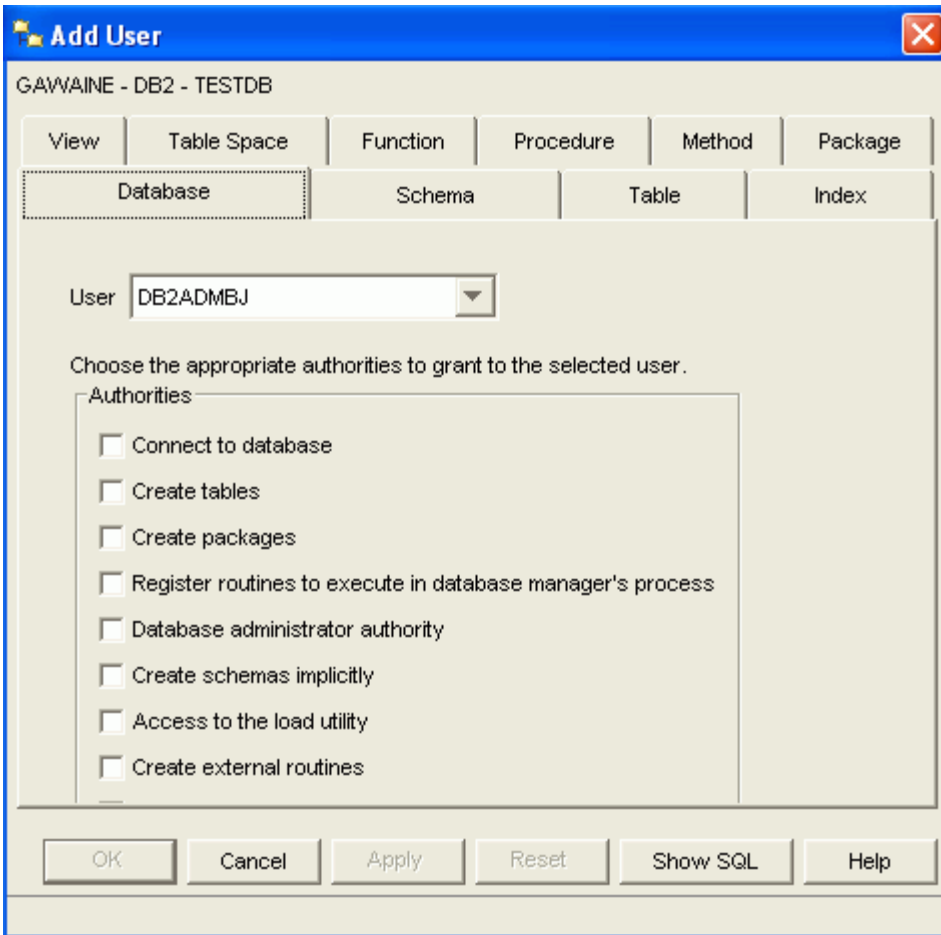


Figure 35. Add User

6. Each of the remaining tabs corresponds to a certain kind of DB2 UDB object (schemas, tables, indexes, views, and so on). Using each of these tabs, you can grant or deny privileges over individual objects to the new user.

7. Click **Apply** to commit your changes. When you are finished, click **OK** to close the window.

5.3 How is my user ID authenticated in Windows and in Linux?

User IDs are authenticated by a security facility outside of DB2 UDB. On Windows and Linux, these facilities will likely be the operating systems themselves, although a third-party system such as Kerberos may be used.

A security facility requires two items to authenticate a user: a user ID and a password. The user ID identifies the user to the security facility. By supplying the correct password (information known only to the user and the security facility) the user's identity (corresponding to the user ID) is verified.

Windows

In Windows^(R) NT, all the client and server computers in a domain share a single user accounts database called the Security Access Manager (SAM). One of the computers in the domain is the domain controller. The domain controller manages all aspects of user-domain interactions. The domain controller uses the information in the domain user accounts database to authenticate users logging onto domain accounts.

Linux

On Linux systems, the security facility is in the operating system itself.

Kerberos

DB2(R) Universal Database (DB2 UDB) provides support for the Kerberos authentication protocol on Red Hat Enterprise Linux Server Advanced Server 2.1 (32-bit Intel), and Windows(R) 2000 and later operating systems. The Kerberos support is provided as a GSS-API security plugin named "IBMkrb5" which is used as both a server and as a client authentication plugin. The library is placed in the following directories:

```
Linux:          sqllib/security{32|64}/plugin/IBM/{client|server}
Windows:       sqllib/security/plugin/IBM{client|server}
```

For more detailed information about user authentication, see the "Security" topic in the DB2 Information Center.

5.4 How can I change the authorization of my DB2 user?

There are two types of permissions recorded by DB2 Universal Database^(TM) (DB2 UDB): privileges and authorities (authorizations). An authority allows a user or group to perform a general task, such as connecting to a database, creating tables, or administering a system. A privilege gives a user or group the right to access one specific database object in a specified way.

You can change the authorities and privileges of your DB2 user through the Control Center.

1. Open the Control Center.
2. Users are created, and are assigned privileges, for specific databases. Find the database for which you would like to change a user's authorization, and expand the tree to see its constituent folders.
3. Expand the User and Group Objects folder; select the "DB Users" object.
4. From the list of users, select the user whose authorization you would like to change.
5. Click on the "Change User Privileges" action in the Object Details Pane or right-click on the user's name and select "Change."
6. In the "Change User" window, select authorities to grant to the user on the "Database" tab (see Figure 36).

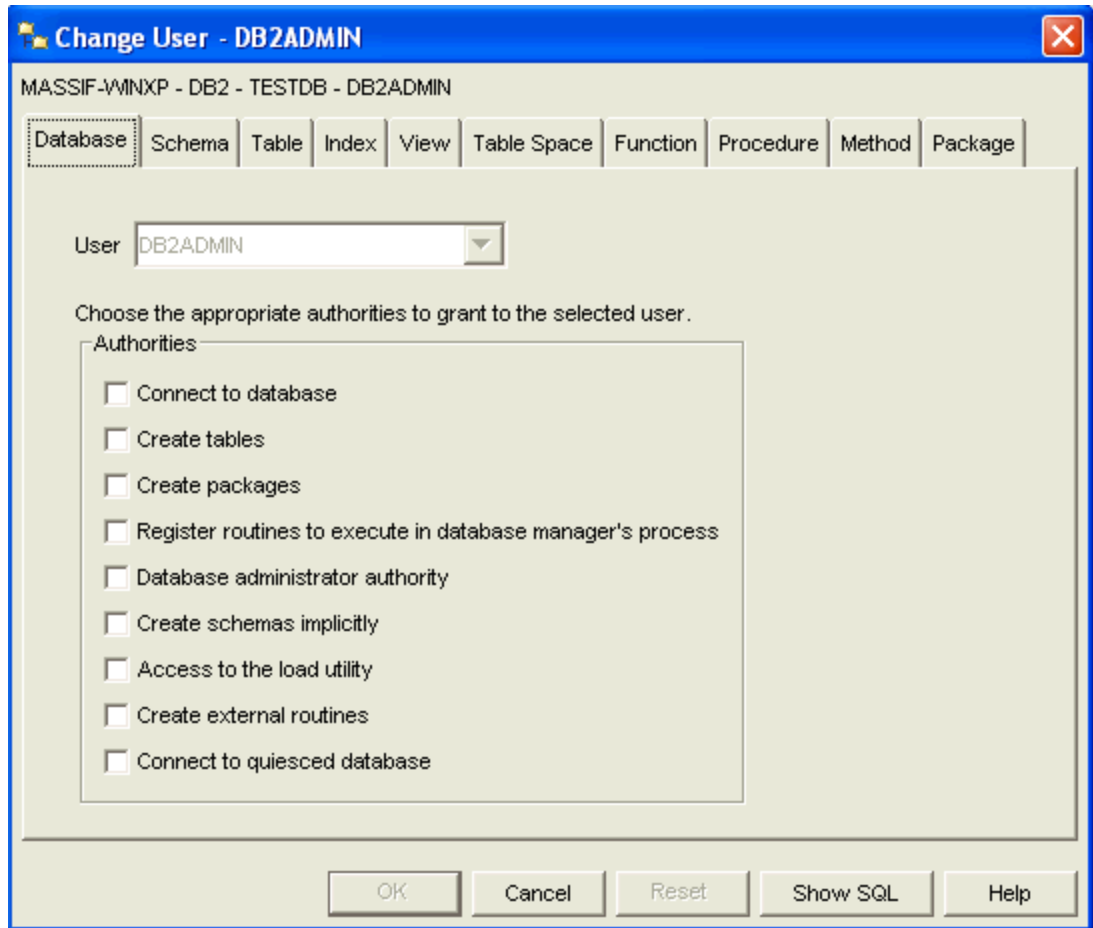


Figure 36. Change User window

7. To grant privileges, select the tab for the appropriate database object. You can add objects (such as tables) in order to specify privileges or revoke privileges for objects already listed (see Figure 37).

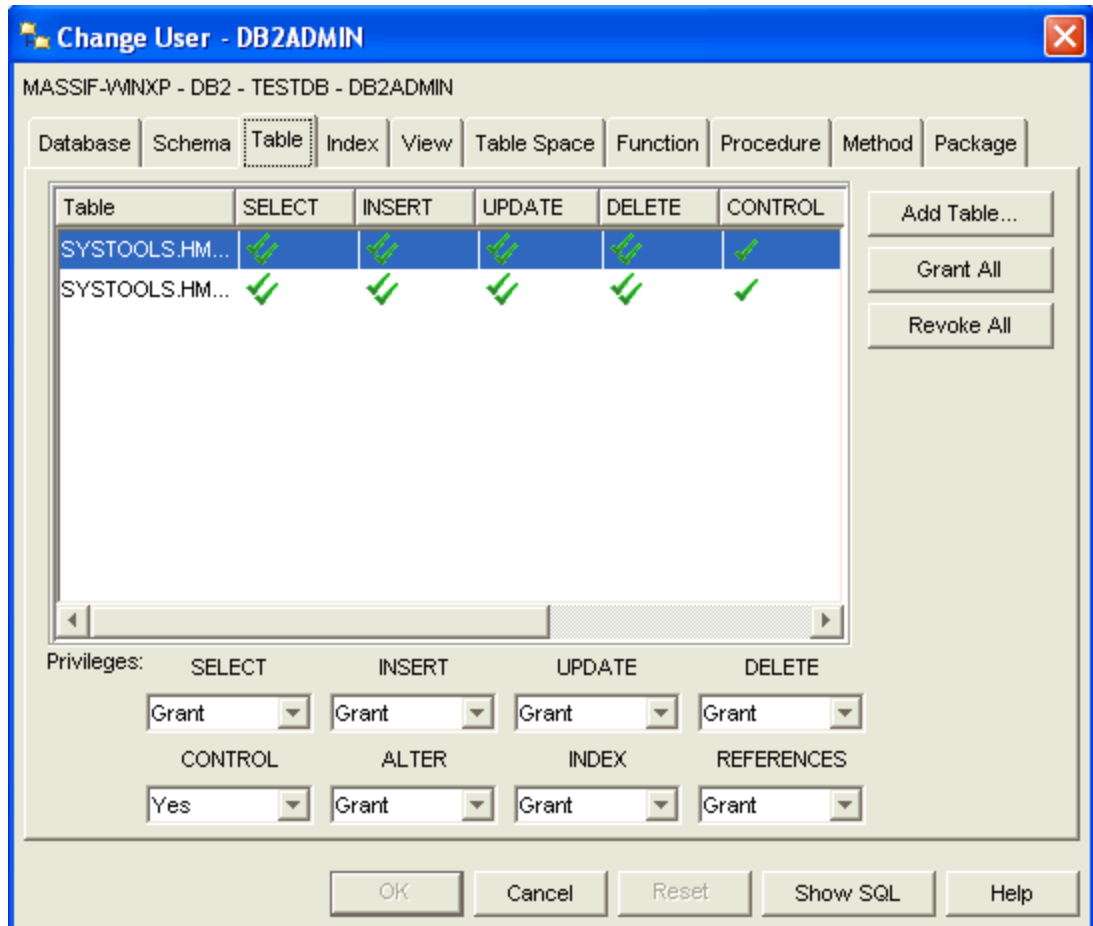


Figure 37. Table tab, Change User window

Appendix A: DB2 administration tools

This section describes the DB2 administration tools. It explains what each tool is used for and why you might want to use it. It tells you how to invoke a tool and presents basic usage information.

Navigation

The administration and development tools enable you to perform your job easily; however, you need to understand how the tools work and become familiar with their features. This section explains basic navigation features.

Start menu

The Start menu is the starting point for using the Administration tools in Windows(R) operating environments. From the Start menu, select **IBM DB2→General Administration Tools→Control Center**. You can also start other centers or tools such as the Journal, the Replication Center, and the Task Center from the Start menu as well as select the command line tools, the development tools, the monitoring tools, and the Information Center.

To start the Administration tools from the Linux desktop environment, open the IBM(R) DB2(R) folder on your Gnome or KDE desktop and then select Control Center.

Control Center menu bars

The Menu bar is a special panel that is displayed at the top of the window. It contains menu options that allow you to access drop-down menus. From the drop-down menus, you can select menu items. Items in the menu bar include actions that affect the objects in the center you are using. For example, menu items under Control Center include actions that affect the entire Control Center. Menu items under Selected dynamically change to include only those actions that apply to currently selected objects in the contents pane. Menu items under Edit include actions that let you work with objects in the contents pane. Menu items under View include actions that let you customize the display of objects in the contents pane. Menu items under Help include actions that display help information and allow you to work with the Information Center and Tutorials.

Toolbars

Toolbars are panels that contain icons representing functions you can perform. Toolbars are located below the menu bar. To see a brief description of a tool, place your cursor over its icon and hover help will tell you what function each icon represents. Toolbars provide quick access to the functions you can perform. The functions can also be selected in the View menu.

A Contents pane toolbar is located below the contents pane. It allows you to tailor the information in the contents pane.

Object trees

Object trees display the system and database objects graphically in the left navigation pane, enabling you to see the hierarchical relationship between different objects and to work with those objects. You can expand the object tree to expose the objects that are contained within the object. The exposed objects are displayed as folders beneath the object. Each folder represents an object type. If the object tree is collapsed, the folders contained within the object no longer appear in the view of the object tree.

Objects on your local workstation and any remote systems that can be connected to and from your local system are displayed in the object tree. Some objects, such as tables and views, do not contain other objects and appear at the bottom of the object tree.

When you select an object in the object tree, the objects that reside in the object are displayed in the contents pane. To invoke an action on an object, right-click the object to open a pop-up menu of available actions. You can also invoke an action on the object by selecting the object then clicking the **Selected** menu option from the menu bar.

Contents Pane

The Contents pane allows you to display large amounts of complex tabular data in a flexible manner.

You can filter, sort, and customize Contents pane columns. You can also save your changes as a view, enabling you to quickly switch between different views of the same data. To create and save customized views, select View from the toolbar at the bottom of the Contents pane.

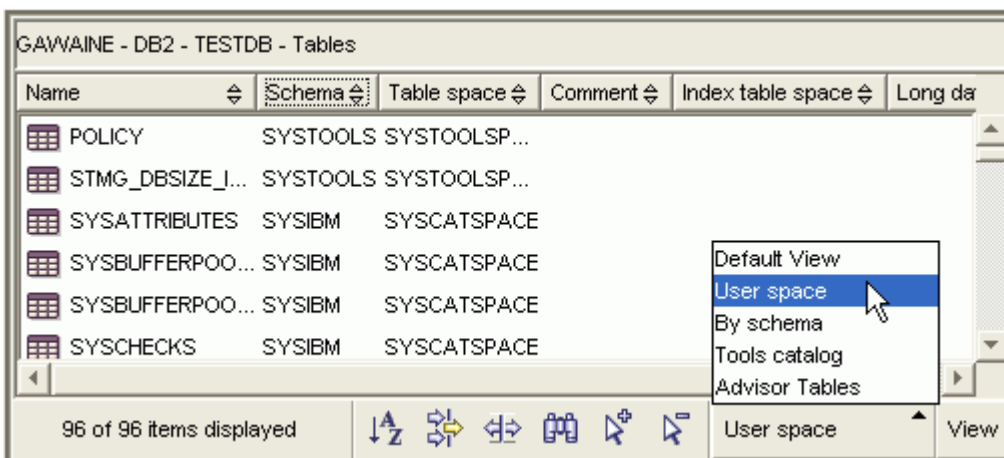


Figure 38. Working with views in the Contents pane

Filtering in the Contents pane simply hides data that has already been retrieved from the database. If your database has so many objects that retrieval time is an issue, you may prefer to retrieve only a subset of the objects (for example, tables). This filtering is referred to as Object Filtering and is explained in the next section.

Object filtering

Object filtering enables you to work with a subset of all the objects in the database in the Control Center. For example, if you only need to work with your own tables and not those created by someone else, you can set up an object filter so that only those tables will be retrieved from the database. You will not have to wait for data about the other tables to be retrieved.

Click on an object folder in the object tree and select **Filter→Create**. The Filter notebook opens allowing you to select which objects you would like to retrieve from the database.

Once you have created a filter, the filter icon will appear beneath the filtered object in the object tree to indicate that a filter has been set. To remove the filter and display all the objects in the Contents pane, right click the object folder and select **Filter→Remove**.

You can further customize your view by selecting the filter icon from the toolbar at the bottom of Contents pane described above. This filtering will only hide data that has already been retrieved from the database.

Infopops

An infopop is a pop-up window that is displayed when a control in a window or notebook has focus and you press F1. Holding the mouse cursor over a control in a window or notebook also causes the infopop to display. Infopops contain a description of the field or control. They may also list restrictions and requirements or provide instructions.

Infopops are disabled or re-enabled from the General page of the Tools Settings notebook. To disable or enable infopops for the Configuration Assistant, select the Display Infopops checkbox menu in the Help menu from the Configuration Assistant tool.

Mandatory fields

Mandatory fields are fields for which you must supply information. Mandatory fields are displayed with thick red borders around the field. When information is entered in the mandatory field, the red border is no longer displayed.

Help

Help is available from icons on the toolbar or by using the Help menu.



Help icon opens the Information Center so you can search for help on tasks, commands, and information in the DB2 library.

The Help menu displays menu items for displaying the online help index, general information about the Control Center, and keyboard help. The menu also displays the list of tutorials available with DB2.

Control Center

You can start the Control Center in the following ways:

- Select Control Center from the Tools menu of another tool.
- Click the **Control Center** icon from the toolbar of another tool.
- Enter the **db2cc** command.
- On Windows(R) systems, click the Start button and select **Programs→IBM DB2→General Administration Tools→Control Center**.
- On Linux systems, open the IBM DB2 folder on the desktop and select Control Center.

In the Control Center, you can administer all of your systems, instances, databases, and database objects such as tables and views. You can also open other centers and tools to help you optimize queries, jobs, and scripts; perform data warehousing tasks; create stored procedures; and work with DB2 commands.

The following are some of the key tasks that you can perform with the Control Center:

- Add DB2 systems, instances, databases, and database objects to the object tree.
- View the details of database objects. You can look at a table's constituent columns and see the current state of a database.
- Manage database objects. You can create, alter, and drop databases, table spaces, tables, views, indexes, triggers, and schemas. You can also manage users.
- Manage data. You can load, import, export, and reorganize data. You can also gather statistics.
- Perform preventive maintenance by backing up and restoring databases or table spaces.

- Manage database connections, such as DB2 Connect(TM) servers and subsystems.
- Manage applications.
- Launch other tools such as the Command Editor and the Health Center.

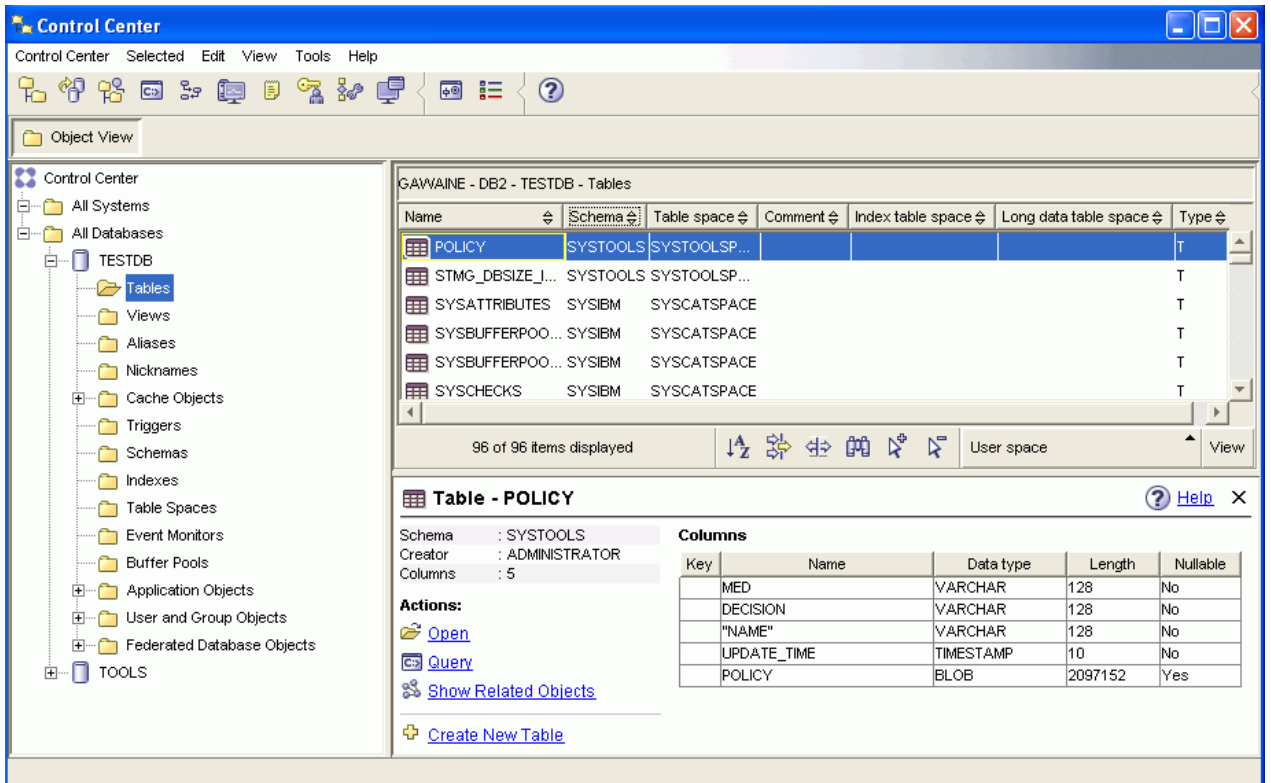


Figure 39. Control Center with a table selected

The main Control Center window is split into three sections, or panes. On the left side, there is an Object Tree. This tree displays all of your databases and their constituent objects (such as tables, views, and indexes).

When you select any one of the available database objects or object folders in the Object Tree, the Contents Pane, occupying the top half of the right side of the Control Center window, fills with a list of the relevant objects. For example, by selecting a database's Table folder in the object tree, you will fill the Contents Pane with all of that database's tables, as well as some information about each.

There is, however, more detail about any given object than will fit in the Contents Pane; the bottom pane on the right side, the Object Details Pane provide the additional information. Clicking on an object in the Contents Pane will bring up detailed information about that object in the Object Details Pane. When you click on a table object in the object list, you will see a listing of that table's columns, as well as other table attributes and some actions that are relevant to tables.

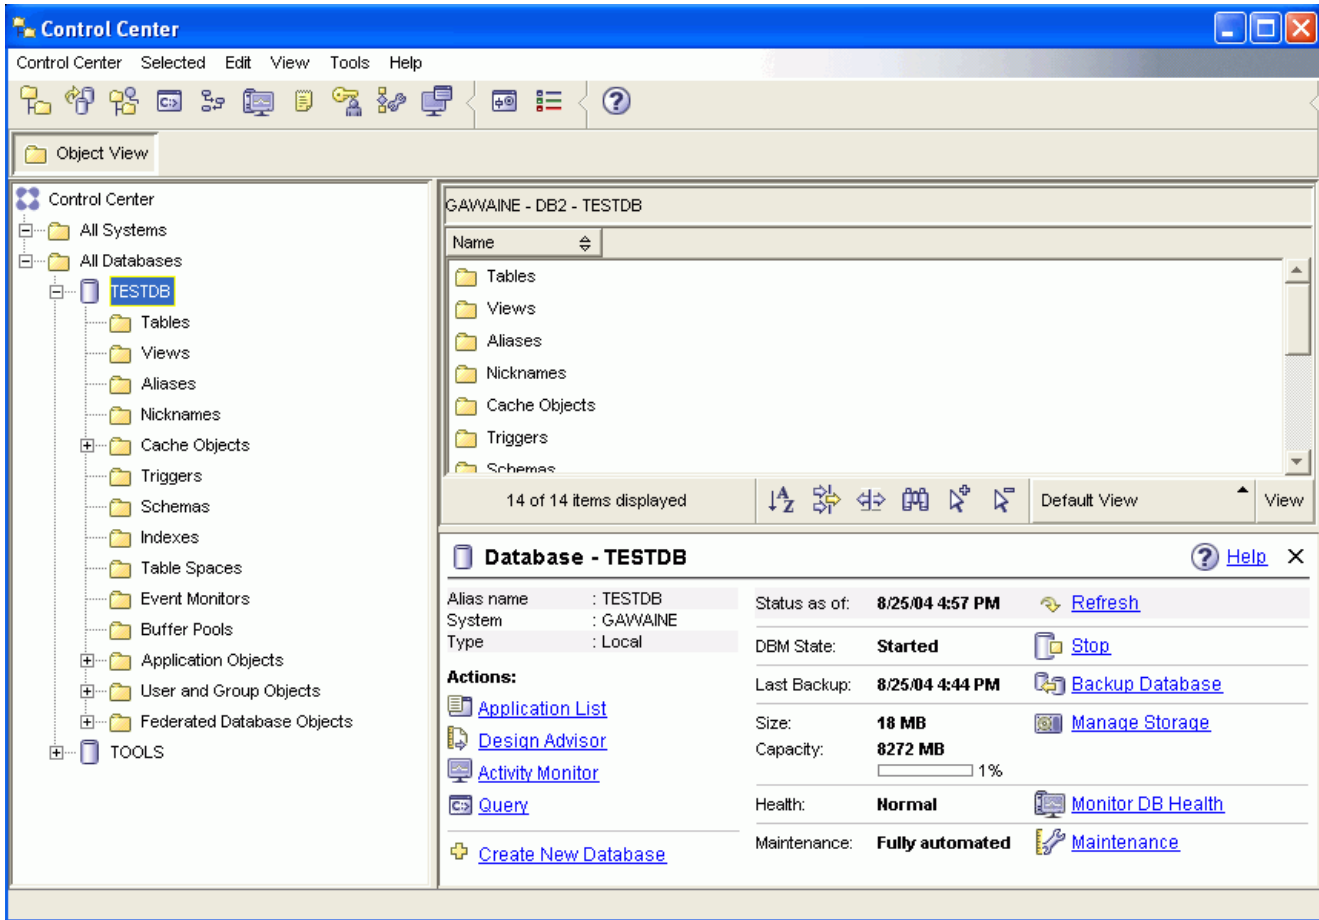


Figure 40. Control Center with a database selected

The screenshot above shows what you might see if you were to select the SAMPLE database. The object details pane displays operational information about the current state of the database, and offers possible actions related to aspects of its status.

The Maintenance status category will tell you whether or not automatic maintenance is enabled for the database in question. Clicking on Maintenance will allow you to enable, configure, and disable automatic maintenance for the database. If you created the database with the Create Database Using Automatic Maintenance wizard, from First Steps or within the Control Center, then Maintenance should be in the "enabled" state.

Note: To schedule tasks through the Control Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2, you can select **Tools**→**Tools Settings**→**Scheduler Settings** and create one now.

To customize objects and actions

Use the Control Center View window to select the view of the Control Center with which you want to work.

DB2 UDB is made up of various components and objects. These objects are used to store and work with data. The database itself is only one of a number of objects that make up a database product's hierarchy of objects. Other components include systems, tables, views, triggers, and indexes. DB2 UDB presents its hierarchy as an object tree in the left pane of the Control Center. Three different Control Center views are available.

To select a view

Open the Control Center View window by clicking **Tools→Customize the Control Center**. Select the Control Center view that you want displayed. This view provides you with the core DB2 functions. Its first folder is the All Databases folder.

Basic

From this folder you can work with all the databases to which you have been granted access, including their related objects such as tables and stored procedures. This view does not include the database objects that are at a higher level than the database in the DB2 hierarchy (systems and instances). It provides you with the essentials for working with data.

Advanced

This view provides you with all the folders, folder objects, and actions available in the Control Center. You should select this view if you are working in an enterprise environment and want to connect to DB2 for z/OS or IMS.

Custom

This view allows you to tailor the Control Center to your needs. You select the folders, folder objects, and object actions that you want to appear in your Control Center view. Click **Custom→Modify**. The Customize Control Center View window opens.

You can choose to have the Control Center View window open whenever you open the Control Center by selecting Show this window at startup time. Your Control Center now displays your view selection.

Command Editor

You can start the Command Editor in the following ways:

- Select Command Editor from the Tools menu of another tool.
- Click the Command Editor from the toolbar of another tool.
- On Windows systems, click the Start button and select **Programs→IBM DB2→Command Line Tools→Command Editor**.

Use the Command Editor to:

- Run DB2 commands and SQL statements
- Work with query results
- View a graphical representation of the access plan for explained SQL statements.
- Run commands in sequence.
- Create and save a script. You can optionally store a saved script in the Task Center, where you can schedule the script to run at a specific time.

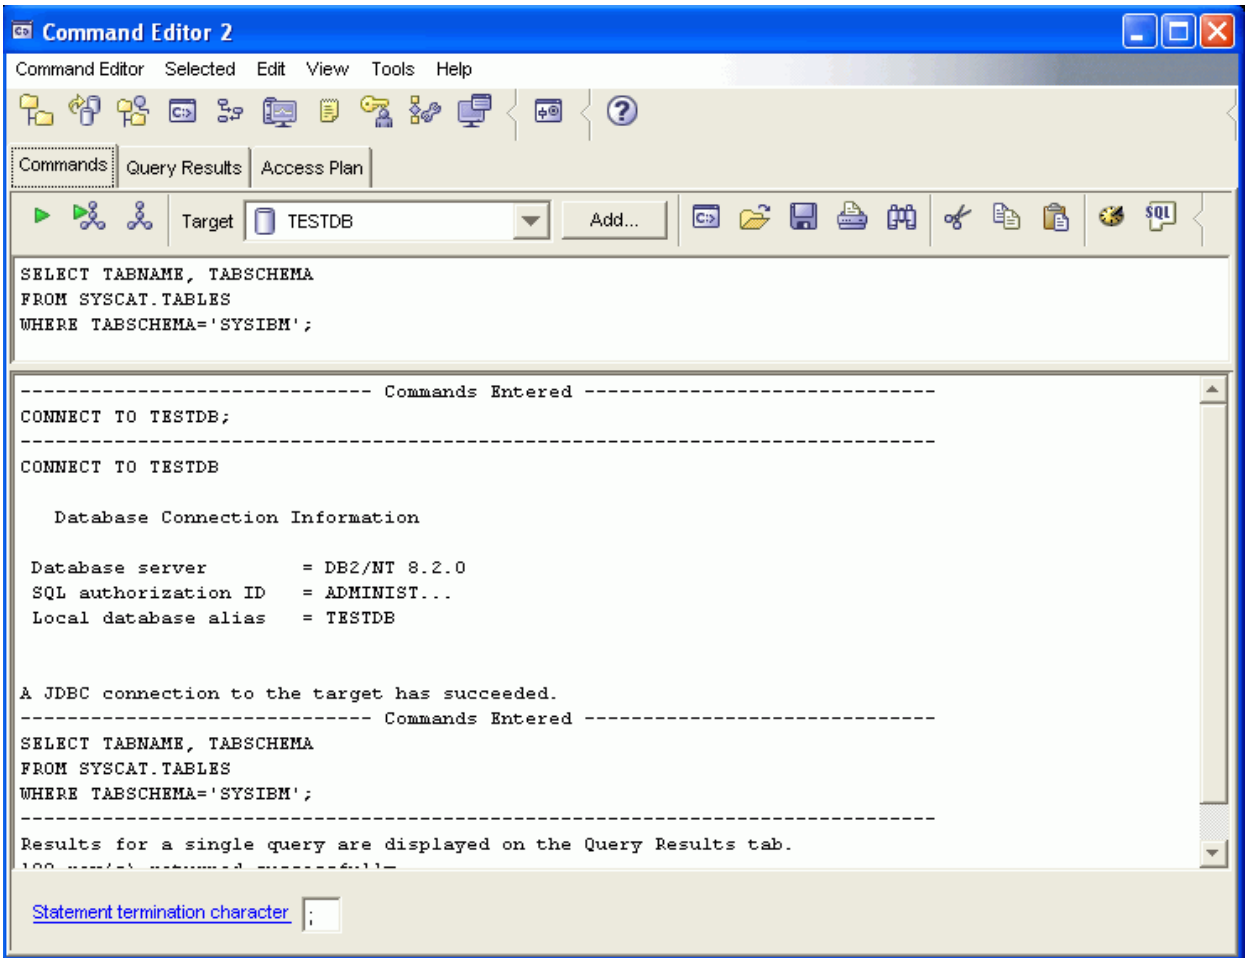


Figure 41. Command Editor window

On the **Commands** page, you can execute an SQL statement or DB2 CLP command. (You do not need to precede the command by **DB2**.)

Run a command or statement, by clicking the **Execute** icon (at the left of the toolbar).









On the **Query Results** page, you can see the results of the queries. You can also save the results or edit the contents of the table.

On the **Access Plan** page, you can see the access plan for any explainable statement that you specified on the Interactive page or the Script page. DB2 generates the access plan when it compiles the SQL statement. You can use this information to tune your queries for better performance.

If you specify more than one statement on the Script page, an access plan is not created unless you highlight a single statement.

You can get help writing SQL with the **SQL Assist** tool. Select **Selected** from the Command Editor menu and choose SQL Assist.

To execute commands and SQL statements:

1. To open an embedded Command Editor, expand the Control Center object tree until you find a DB2 Universal Database database. Right-click the object and select **Query** from the pop-up menu. A Command Editor tile opens in the Control Center.
2. To open a stand-alone Command Editor from the Start menu, select **Programs→IBM DB2→Command Line Tools→Command Editor**.
3. Select a target database to connect to. Targets that currently have connections or have had connections during the current session are listed in the Target drop-down list. If your target is not listed, click Add and select a target from those available on the Specify Target window.
4. In the top pane, type the commands or SQL statements that you want to execute.
5. You can select commands and statements previously executed in this session from the History window. Click  to open the History window.
6. You can execute commands and statements stored in a file. Click  to locate the file using the file browser Open window.
7. You can specify the character to be used as the termination character for your statements and commands in the Statement termination character field. This termination character will override the character that you may have specified in the Tool Settings notebook.
8. For scripts containing CREATE PROCEDURE or CREATE TRIGGER statements, it is suggested that a termination character other than a semicolon be used, as the body of these statements might themselves contain one or more semicolons.
9. When stored procedures are using the same character as a delimiter, as this would cause execution problems.
10. Click  to execute the commands or statements.
11. Click  to generate a graphical representation of the access plan for explainable SQL statements (for example, SELECT, INSERT, and UPDATE). An access plan indicates the order of operations for accessing requested data.
12. Click  to execute explainable SQL statements and generate access plans.
Note: The availability of push buttons changes depending on the target for your commands.
13. View the execution results. Command results are displayed in the output area of the Commands page, while SQL results are returned on the Query Results page. The access plan appears in graphical form on the Access Plan page.
14. Optional: The results from the most recently executed SQL statement are displayed on the Query Results page. To change the value in any cell, highlight the value and modify it. If the Automatically commit updates check box is selected, changes made to a cell are committed immediately when focus shifts off its row. If the check box is not selected, then click Commit to commit your changes, or click Rollback to rollback any uncommitted changes.
15. You can also add and delete rows in the table.
16. Optional: You can save your commands and statements in a file. Click  to specify a file using the file browser Save As window.

Health Center

You can start the Health Center in the following ways:

Select Health Center from the Tools menu of another tool.

1. Click the **Health Center** icon from the toolbar of another tool.
2. Enter the **db2hc** command.
3. On Windows systems, click the Start button and select **Programs→IBM DB2→Monitoring Tools→Health Center**.
4. Click the health center status beacon when it appears on a DB2(R) window. On the Tools Settings window, you can configure this beacon to appear on every DB2 window whenever an alert occurs.

Use the Health Center to resolve health alerts in the database environment and make any necessary changes.

When you use DB2, a health monitor continuously monitors a set of health indicators. If the current value of a health indicator is outside the acceptable operating range defined by its warning and alarm thresholds, the health monitor generates a health alert. DB2 comes with a set of predefined threshold values, which you can later customize. For example, you can customize the alarm and warning thresholds for the amount of space used in a table space.

Depending on the configuration of the DB2 instance, the following actions can occur when the health monitor generates an alert:

- An entry is written in the administration notification log, which you can read from the Journal. Also readable in the Windows Event Log
- The health center status beacon appears in the lower right corner of the DB2 Tools window.
- A script or task is executed.
- An e-mail or pager message is sent to the contacts that you specify for this instance.

The following are some of the key tasks that you can perform with the Health Center:

- View the status of the database environment. Beside each object in the navigation tree, an icon indicates the most severe alert for the object (or for any objects contained by that object). For example, a green diamond icon beside an instance means that the instance and the databases contained in the instance do not have any alerts.
- View the alerts for an instance or a database. When you select an object in the navigation tree, the alerts for that object are shown in the pane to the right.
- View details on the alert or use the Recommendation advisor to address the alert situation. Right click on an alert and select the Recommendation Advisor. The first page shows the details for the alert. The second page shows any recommended actions.
- Configure health indicator settings
- Select which contacts will be notified of alerts with an e-mail or pager message.
- Review the history of alerts for an instance.

You can use the toggle buttons at the top of the navigation bar



to filter the alerts according to their severity:

- Alarm toggle Displays objects in alarm state (the most severe type of alert).
- Alarm and warning toggle Displays objects in alarm or warning state.
- Alarm, Warning, and Attention toggle Displays objects in any alert state.
- Toggle for all states Displays all objects.

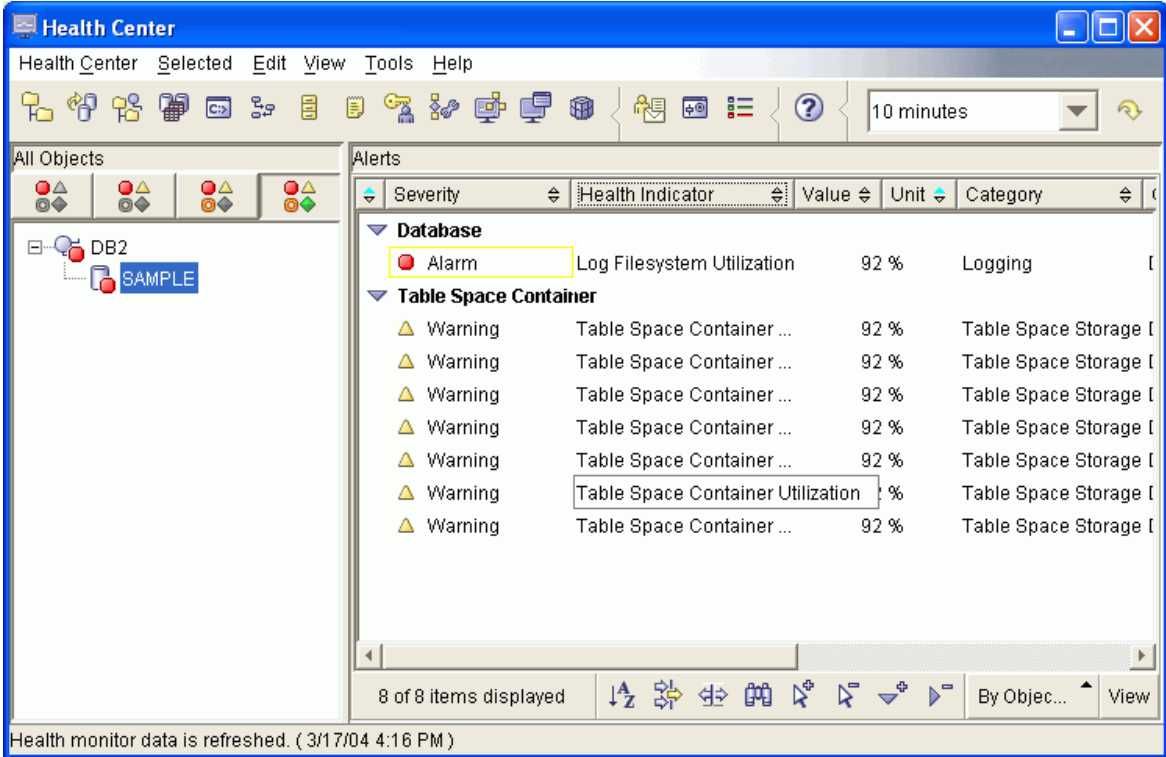


Figure 42. Health Center window

Setting up notification and contact lists

For DB2 to contact you or others regarding the health of your databases, you will have to set up notification and contact lists. If you did not do this during the DB2 installation, or you did not create your database with Automatic Maintenance, you can manually set up these lists.

Procedure

If the SMTP server was not setup during the installation process, you can configure it manually using the following command:

db2 update admin cfg using stmp_server host1

where **host1** represents the TCP/IP host name for the SMTP server used for e-mail notification. You can also update this information using the Create Database with Automatic Maintenance wizard, or the Troubleshoot Alert Notification wizard in the Health Center.

If you want the contact information to be located on a remote DB2 administration server (DAS), you can specify the `contact_host` parameter using the following command:

db2 update admin cfg using contact_host host2

where **host2** represents the TCP/IP host name for the DAS. If the `contact_host` parameter is not specified, the DAS assumes the contact information is local.

Turn on the scheduler using the following command:

```
db2 update admin cfg using sched_enable on
```

For these changes to take effect, you need to restart the DAS use the following commands:

```
db2admin stop  
db2admin start
```

In the Control Center, click on the **Contacts** icon in the Taskbar. Select System name, then click Add Contact. Type in contact information and then press OK. You can also add the new contact using the Task Center, the Health Center by configuring alert notification, or the Create Database with Automatic Maintenance.

Task Center

You can start the Task Center in the following ways:

- Select Task Center from the Tools menu of another tool.
- Click the **Task Center** icon from the toolbar of another tool.
- Enter the **db2tc** command.
- On Windows(R) systems, click the Start button and select **Programs→IBM DB2→General Administration Tools→Task Center**.

Note: To use the Task Center, you must first create a tools catalog. If you did not create a tools catalog when you installed DB2, you can select **Tools→Tools Settings→Scheduler Settings** and create one now.

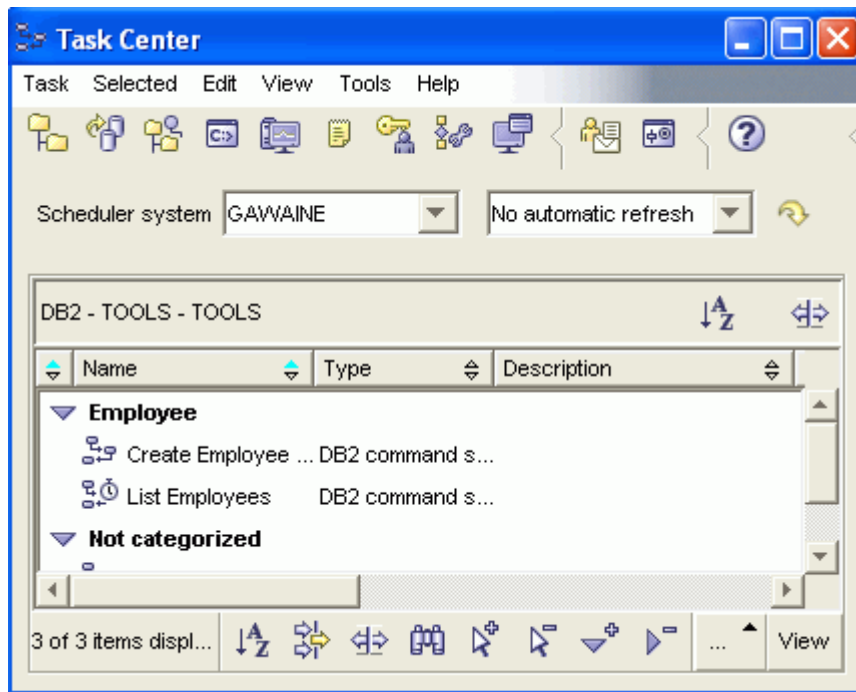


Figure 43. Task Center window

Use the Task Center to run tasks, either immediately or according to a schedule, and to notify people about the status of completed tasks. The Task Center includes functionality from the Script Center in previous versions of DB2, plus additional functionality.

A task is a script, together with the associated success conditions, schedules, and notifications. You can create a task within the Task Center, create a script within another tool and save it to the Task Center, import an existing script, or save the options from a DB2 dialog or wizard such as the Load wizard. A script can contain DB2, SQL, or operating system commands.

For each task, you can do the following:

- Schedule the task
- Specify success and failure conditions
- Specify actions that should be performed when this task completes successfully or when it fails
- Specify e-mail addresses (including pagers) that should be notified when this task completes successfully or when it fails

You can specify conditional coding by creating task actions. Each task action consists of a task and the action that should be performed on the task. For example, task 1 could have the following task actions:

- If task 1 is successful, task action A enables the schedule for task 2.
- If task 1 fails, task action B runs task 3.

You can also create a grouping task, which combines several tasks into a single logical unit of work. When the grouping task meets the success or failure conditions that you define, any follow-on tasks are run. For example, you could combine three backup scripts into a grouping task and then specify a reorganization as a follow-on task that will be executed if all of the backup scripts execute successfully.

To see the most recent information, click the **Refresh** icon button. You can view execution results as follows:

- To see the results from the last time a task was executed, right-click the task and select Show Results.
- To see the results from every time a task was executed, open the Journal. In the Journal, you can right-click the task execution that you are interested in and select Show Results.
- To see help text for a message that appears in the execution results, open the Command Center and enter the error message preceded by a question mark. For example, enter ? SQL0100 to get more information about message SQL0100W.

Several customizable saved views are available via the pull-up menu at the bottom of the Task Center.

Journal

You can start the Journal in the following ways:

- Select **Journal** from the Tools menu of another tool.
- Click the **Journal** icon from the toolbar of another tool.
- On Windows systems, click the Start button and select **Programs→IBM DB2→General Administration Tools→Journal**.

The Journal displays historical information about tasks, database actions and operations, Control Center actions, messages, and alerts.

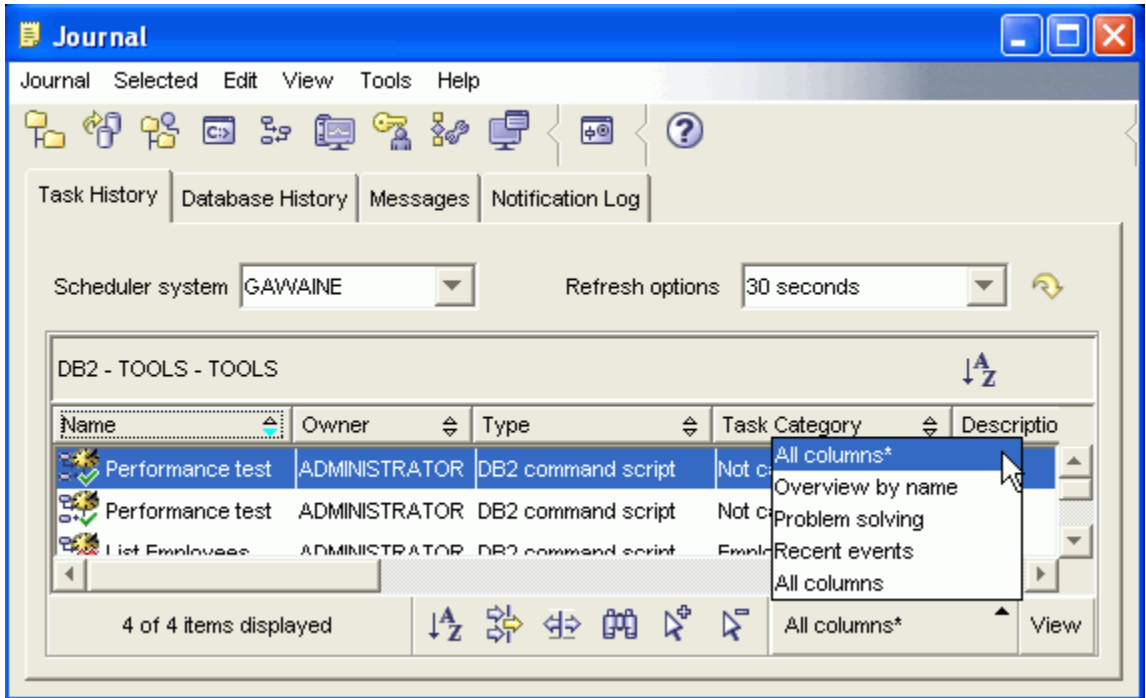


Figure 44. Journal window

To see the most recent information, click the **Refresh** icon button.

The Task History page shows the results of tasks that were previously executed. You can use this information to estimate how long future tasks will run.

The Task History page contains one row for each execution of a task. The Task Center contains only one row for each task no matter how many times the task is executed. Each row in the Task Center could be directly related to multiple rows in the Task History page of the Journal.

For each completed execution of a task, you can perform the following actions:

- View the execution results
- View the task that was executed
- Edit the task that was executed
- View the task execution statistics
- Remove the task execution object from the Journal

To perform one of these actions, right-click a completed task execution and select the corresponding action from the pop-up menu that appears.

Database History

You can check the status of key events in the life of your database using the database history. The Database History page in the Journal lists a record of recovery and administrative events

- A database or table spaces are backed up, restored or rolled forward
- A table space is created, altered, quiesced, renamed or dropped
- A table is loaded, dropped, reorganized.

- A database is recovered

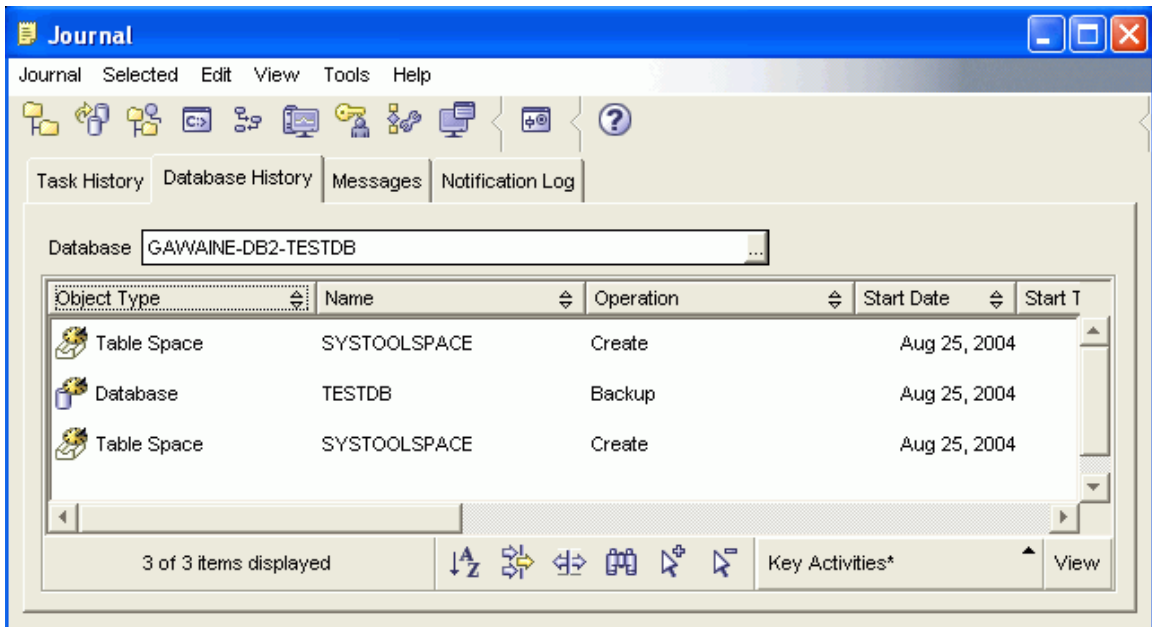


Figure 45. Database History

The Messages page shows messages that were previously issued from the Control Center and other GUI tools.

The Notification Log page shows information from the administration notification log. You can filter the list. For example, you can show only alerts from the health monitor.

Several customizable saved views are available via the pull-up menu at the bottom of the Journal.

Appendix B: Automatic maintenance

DB2 Universal Database(TM) (UDB) provides automatic maintenance capabilities for performing database backups, keeping statistics current, and reorganizing tables and indexes as necessary.

Automatic database backup provides users with a solution to help ensure their database is being backed up both properly and regularly, without either having to worry about when to back up, or having any knowledge of the backup command.

Automatic statistics collection attempts to improve the performance of the database by maintaining up-to-date table statistics. The goal is to allow the optimizer to choose an access plan based on accurate statistics.

Automatic statistics profiling advises when and how to collect table statistics by detecting outdated, missing, and incorrectly specified statistics and by generating statistical profiles based on query feedback.

Automatic reorganization manages offline table and index reorganization without users having to worry about when and how to reorganize their data.

Enablement of the automatic maintenance features is controlled using the automatic maintenance database configuration parameters. These are a hierarchical set of switches to allow for simplicity and flexibility in managing the enablement of these features.

Automatic database backup

A database may become unusable due to a wide variety of hardware or software failures. Automatic database backup simplifies database backup management tasks for the DBA by ensuring that a recent full backup of the database is performed as needed. It determines the need to perform a backup operation based on one or more of the following measures:

- A full database backup has never been completed.
- The time elapsed since the last full backup is more than a specified number of hours.
- The transaction log space consumed since the last backup is more than a specified number of 4 KB pages (in archive logging mode only).

Protect your data by planning and implementing a disaster recovery strategy for your system. If suitable to your needs, you may incorporate the automatic database backup feature as part of your backup and recovery strategy.

If the database is enabled for roll-forward recovery (archive logging), then automatic database backup can be enabled for either online or offline backup. Otherwise, only offline backup is available. Automatic database backup supports disk, tape, Tivoli(TM) Storage Manager (TSM), and vendor DLL media types.

Through the Configure Automatic Maintenance wizard in the Control Center or Health Center, you can configure:

- The requested time or number of log pages between backups
- The backup media.
- Whether it will be an online or offline backup.

If backup to disk is selected, the automatic backup feature will regularly delete backup images from the directory specified in the Configure Automatic Maintenance wizard. Only the most recent backup image is guaranteed to be available at any given time. It is recommended that this directory be kept exclusively for the automatic backup feature and not be used to store other backup images.

The automatic database backup feature can be enabled or disabled by using the **auto_db_backup** and **auto_maint** database configuration parameters. In a multiple database partitioned environment, the automatic database backup runs on each partition if the database configuration parameters are enabled on that partition.

Automatic statistics collection

When the SQL compiler optimizes SQL query plans, its decisions are heavily influenced by statistical information about the size of the database tables and indexes. The optimizer also uses information about the distribution of data in specific columns of tables and indexes if these columns are used to select rows or join tables. The optimizer uses this information to estimate the costs of alternative access plans for each query. When significant numbers of table rows are added or removed, or if data in columns for which you collect statistics is updated, the RUNSTATS utility needs to be run to update the statistics.

Automatic statistics collection works by determining the minimum set of statistics that give optimal performance improvement. The decision to collect or update statistics is taken by observing and learning how often tables are modified and how much the table statistics have changed. The automatic statistics collection algorithm learns over time how fast the statistics change on a per table basis and internally schedules RUNSTATS execution accordingly.

Normal database maintenance activities such as when a user performs RUNSTATS, REORG or altering or dropping the table, are not affected by the enablement of this feature.

If you are unsure about how often to collect statistics for the tables in your database, you can incorporate the automatic statistics collection feature as part of your overall database maintenance plan.

The automatic statistics collection feature can be enabled or disabled by using the **auto_runstats**, **auto_tbl_maint**, and **auto_maintdatabase** configuration parameters.

Tables considered for automatic statistics collection are configurable by you using the Automatic Maintenance wizard from the Control Center or Health Center.

Automatic reorganization

After many changes to table data, logically sequential data may be on non-sequential physical pages so the database manager has to perform additional read operations to access data.

Among other information, the statistical information collected by RUNSTATS show the data distribution within a table. In particular, analysis of these statistics can indicate when and what kind of reorganization is necessary. Automatic reorganization determines the need for reorganization on tables by using the REORGCHK formulas. It periodically evaluates tables that have had their statistics updated to see if reorganization is required. If so, it internally schedules a classic reorganization for the table. This requires that your applications function without write access to the tables being reorganized.

The automatic reorganization feature can be enabled or disabled by using the **auto_reorg**, **auto_tbl_maint**, and **auto_maint** database configuration parameters.

If you are unsure about when and how to reorganize your tables and indexes, you can incorporate automatic reorganization as part of your overall database maintenance plan.

Tables considered for automatic reorganization are configurable by you using the Automatic Maintenance wizard from the Control Center or Health Center.

Maintenance windows for automation

The automatic maintenance features described previously consume resources on your system and might affect the performance of your database when they are run. Automatic reorganization and offline database backup also restrict access to the tables and database when these utilities are run. It is therefore necessary to provide appropriate periods of time when these maintenance activities can be internally scheduled for execution by DB2 UDB. These can be specified as offline and online maintenance time periods using the Automatic Maintenance wizard available through the Control Center or Health Center.

Offline database backups and table and index reorganization are run in the offline maintenance time period. These features run to completion even if they go beyond the time period specified. The internal scheduling mechanism learns over time and estimates job completion times. If the offline time period is too small for a particular database backup or reorganization activity, the scheduler will not start the job the next time and relies on the health monitor to provide notification of the need to increase the offline maintenance time period.

Automatic statistics collection and profiling, as well as online database backups, are run in the online maintenance time period. To minimize the impact on the system, they are throttled by the adaptive utility throttling mechanism. The internal scheduling mechanism uses the online maintenance time period to start the online jobs. These features run to completion even if they go beyond the time period specified.

Storage

The automatic statistics collection and reorganization features store working data in tables in your database. These tables are created in the SYSTOOLSPACE table space. The SYSTOOLSPACE table space is created automatically with default options when the database is activated. Storage requirements for these tables are proportional to the number of tables in the database and should be calculated as approximately 1 KB per table. If this is a significant size for your database, you may want to drop and re-create the table space yourself and allocate storage appropriately. The automatic maintenance and health monitor tables in the table space are automatically re-created. Any history captured in those tables is lost when the table space is dropped.

Monitoring and notification

The health monitor provides monitoring and notification functionality for the automatic database backup, statistics collection, and reorganization features.

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country/region or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact:

IBM Canada Limited
Office of the Lab Director
8200 Warden Avenue
Markham, Ontario
L6G 1C7
CANADA

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information may contain sample application programs, in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

(C) (*your company name*) (*year*). Portions of this code are derived from IBM Corp. Sample Programs.
(C) Copyright IBM Corp. *_enter the year or years_*. All rights reserved.

Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both, and have been used in at least one of the documents in the DB2 UDB documentation library.

ACF/VTAM
AISPO
AIX
AIXwindows
AnyNet
APPN
AS/400
BookManager
C Set++
C/370
CICS
Database 2
DataHub
DataJoiner
DataPropagator
DataRefresher
DB2
DB2 Connect

iSeries
LAN Distance
MVS
MVS/ESA
MVS/XA
Net.Data
NetView
OS/390
OS/400
PowerPC
pSeries
QBIC
QMF
RACF
RISC System/6000
RS/6000
S/370
SP

DB2 Extenders
DB2 OLAP Server
DB2 Information Integrator
DB2 Query Patroller
DB2 Universal Database
Distributed Relational
Database Architecture
DRDA
eServer
Extended Services
FFST
First Failure Support Technology
IBM
IMS
IMS/ESA

SQL/400
SQL/DS
System/370
System/390
SystemView
Tivoli
VisualAge
VM/ESA
VSE/ESA
VTAM
WebExplorer
WebSphere
WIN-OS/2
z/OS
zSeries

The following terms are trademarks or registered trademarks of other companies and have been used in at least one of the documents in the DB2 UDB documentation library:

Microsoft, Windows, Windows NT, and the Windows logo are trademarks 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.

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

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

Sleepycat and the names of Sleepycat Software products referenced herein are trademarks, registered trademarks or service marks of Sleepycat Software, Inc.

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