

DB2 UDB SmarTeam Best Practices

DB2. Information Management Software

DB2 UDB V7.2 EE Windows platforms

DB2 Basic Administration

Backup & Restore

Performance Configuration using the Wizard


Wednesday, September 05, 2007

V1.01

Prepared by :

Amr Roushdi

IDSICC



IBM Dassault Systèmes International Competency Center

idsicc@fr.ibm.com



Trade Marks & Disclaimers

The information in this document concerns new products that IBM may or may not announce. Any discussion of OEM products is based upon information which has been publicly available and is subject to change. The specification of some of the features described in this document may change before the General Availability date of these products

REFERENCES IN THIS DOCUMENT TO IBM PRODUCTS, PROGRAMS, OR SERVICES DO NOT IMPLY THAT IBM INTENDS TO MAKE THESE AVAILABLE IN ALL COUNTRIES IN WHICH IBM OPERATES.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customers ability to evaluate and integrate them into the customers operational environment.

While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.
IBM MAY HAVE PATENTS OR PENDING PATENT APPLICATIONS COVERING SUBJECT MATTER IN THIS DOCUMENT. THE FURNISHING OF THIS DOCUMENT DOES NOT IMPLY GIVING LICENSE TO THESE PATENTS.

THE FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE IBM CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: IBM, AIX, APPROACH, AS/400, DB2, DB2 OLAP SERVER, NOTES, OPERATING SYSTEM/2, OS/400, ES/9000, OS/390, OS/2, RISC, RISC SYSTEM/6000, SQL, SQL/DS, VISUAL WAREHOUSE, VM/ESA. LOTUS AND 1-2-3 ARE REGISTERED TRADEMARKS OF LOTUS DEVELOPMENT CORPORATION

THE FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF THE MICROSOFT CORPORATION IN THE UNITED STATES AND/OR OTHER COUNTRIES: MICROSOFT, WINDOWS, WINDOWS NT, ODBC, WINDOWS 95

THE FOLLOWING TERMS ARE TRADEMARKS OR REGISTERED TRADEMARKS OF DASSAULT SYSTEMES IN FRANCE AND/OR OTHER COUNTRIES: CATIA, ENOVIA
SMARTEAM[®] is a registered trademark of SMARTEAM Corporation Ltd

UNIX IS A REGISTERED TRADEMARK IN THE UNITED STATES AND OTHER COUNTRIES EXCLUSIVELY THROUGH X/OPEN COMPANY LIMITED.

OTHER COMPANY, PRODUCT AND SERVICE NAMES MAY BE TRADEMARKS OR SERVICE MARKS OF OTHERS.

Any rights not expressly granted herein are reserved.

IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends upon the customer's ability to evaluate and integrate them into the customer's operational environment.

IBM grants SmarTeam the right to reproduce, distribute and display this Publications provided that all proprietary notices are preserved. However, SmarTeam may not make derivative works of these Publications without the express consent of IBM.

Table of contents

TRADE MARKS & DISCLAIMERS.....	2
TABLE OF CONTENTS.....	3
ACKNOWLEDGMENTS.....	5
<i>The Team that participated in this project.....</i>	<i>5</i>
<i>Other contributions</i>	<i>5</i>
INTRODUCTION AND OBJECTIVES	6
SECURITY REQUIREMENT FOR ADMINISTRATION A REMINDER :.....	7
BACKUP AND RECOVERY.....	9
<i>Introduction</i>	<i>9</i>
<i>Developing a Backup and Recovery Strategy.....</i>	<i>9</i>
<i>Performing the backup using the control center</i>	<i>12</i>
<i>Performing the restore</i>	<i>16</i>
BASIC DB2 TUNING USING THE PERFORMANCE WIZARD	18
APPENDIX A : RDBMS TERMINOLOGY MAPPINGS :	27
<i>Mapping of Oracle Terminology to DB2 UDB</i>	<i>27</i>
APPENDIX B. MOST USED DB2 COMMANDS	29
INSTANCE.....	30
DATABASE	31
CONNECTIVITY.....	31
PERFORMANCE.....	32
ADMINISTRATION.....	32
ADMIN SERVER.....	33
APPLICATION DEVELOPMENT	33
IDSICC.....	34
INTERNATIONAL COMPETENCY CENTER.....	34
OPTIMIZE COMPETITIVENESS THROUGH E-BUSINESS TECHNOLOGIES AND SERVICES.....	34
<i>Total Digital enterprise Solution Optimization.....</i>	<i>36</i>
IBM SOFTWARE.....	37
INTRODUCTION TO DB2 UDB GUI	38
DB2 GUI TOOLS OVERVIEW	38
<i>The Control Center</i>	<i>39</i>
<i>The Configuration Assistant</i>	<i>39</i>
<i>The Task Center</i>	<i>40</i>
<i>The Journal</i>	<i>41</i>
<i>The Command Center</i>	<i>41</i>
<i>Health Center</i>	<i>42</i>
PERFORMING BASIC DATABASE OPERATIONS	43
<i>Creating a database</i>	<i>43</i>
<i>Creating a table</i>	<i>46</i>
<i>Altering a table</i>	<i>47</i>
<i>Performing queries</i>	<i>49</i>
<i>Restricting access to your objects</i>	<i>51</i>
<i>Setting up connectivity</i>	<i>53</i>
<i>Summary</i>	<i>55</i>
<i>What's Next</i>	<i>55</i>

About the author	56
RESOURCES ON THE WORLD WIDE WEB.....	57



Acknowledgments

This white paper is the result of a joint effort performed at the IDSICC in Suresnes Paris France.

The Team that participated in this project.

Alex Bruskin : A former DB2 application developer .

Guillaum Testelin : PLM IT Architecture giving pre sales support to the IBM PLM community as well as business partners .

Christian Nicolas : PLM and in particular CATIA products specialist currently assigned to the IBM Dassault Systemes competence.

Amr Roushdi: Data Management consultant and providing support for IDSICC as well as SmarTeam and Dassault Systemes.

Other contributions

DB2 Technical Enablement Services IBM Toronto Lab for providing RDMS terms mapping to DB2

Paul YIP IBM Toronto lab for providing DB2 common command line processor commands

Introduction and objectives


This publication does not replace the existing documentation provided for DB2. nor by SmarTeam's documentation and recommendations .

It's primary objective is to let you start quickly using simple methods and rules for the different administration tasks evoked in this documentation .

For more information on DB2 administration please refer to the DB2 administration guide and for specific DB2 commands syntax please refer to the DB2 Command reference .

For more information on DB2 and or on SmarTeam please refer to the documentation CD's provided .

This document is hyperlinked and http links are active .



Security requirement for administration a reminder :

A user account that will be used to perform the administration .

The account you are logged in as must:

- Be defined locally

- Belong to the Local Administrator's group

- Have the following advanced user rights:

 - Act as part of the operating system*

 - Create token object*

 - Increase quotas*

 - Replace a process level token*

NB: You can perform the installation without these advanced user rights, but the setup program will be unable to validate accounts. We recommend that any user account used to install this product have these advanced user rights.

During installation, you will provide a user account that will be used by the DB2 Administration Server to log on to the system and to start itself as a service. The account you specify must be defined locally and belong to the Local Administrator's group.

By default, the setup program will create a user account using the username db2admin and the password that you specify. You can accept the default user account, create your own user account by modifying the default values, or provide your own. If you create or provide your own user account, ensure that it conforms to DB2's naming rules.

To verify that DB2 installed correctly, you will need to have a user account that belongs to the DB2 System Administrative (SYSADM) group. The account name must comply with DB2's naming rules, as described in Username, User ID, Group Name, and Instance Name Rules. By default, any user that belongs to the Local Administrators group, on the local machine where the account is defined, has SYSADM authority on the instance. For more information, see Working with the System Administrative Group. For more information on valid DB2 usernames, see Appendix E, Naming Rules in the quick beginning guide .

LDAP Considerations :

During the installation, you may also be prompted for user names and passwords for the products and services that you install.

If you want to use LDAP with Windows 2000, you must extend the directory schema to contain DB2 object classes and attribute definitions. You must do this once, before you install any DB2 product.

To extend the directory schema, execute the db2schex.exe program from the installation CD with Schema Admins authority. You can execute this program with Schema Admins authority, without logging off and logging on again, as follows:

```
runas /user:MyDomain\Administrator x:\db2\common\db2schex.exe
```

where x: represents the CD-ROM letter.

When db2schex.exe completes, you can continue with the installation.



Backup and recovery

Introduction

Developing a Backup and Recovery Strategy

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, and application failures, and different failure scenarios require different recovery actions. 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? Will table space level backups be sufficient, or will full database backups be necessary?

A database recovery strategy should ensure that all information is available when it is required for database recovery. It should include a regular schedule for taking database backups and, in the case of partitioned database systems, include backups when the system is scaled (when database partition servers or nodes are added or dropped). Your overall strategy should also include procedures for recovering command scripts, applications, user-defined functions (UDFs), stored procedure code in operating system libraries, and load copies.

Different recovery methods are discussed in the sections that follow, and you will discover which recovery method is best suited to your business environment.

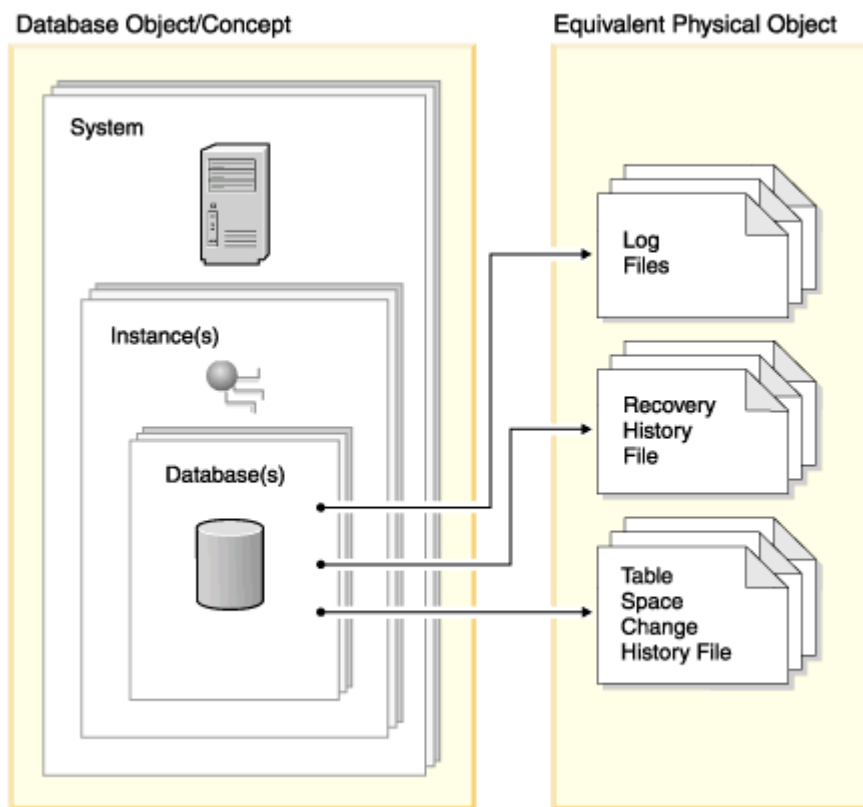
The concept of a database *backup* is the same as any other data backup: taking a copy of the data and then storing it on a different medium in case of failure or damage to the original. The simplest case of a backup involves shutting down the database to ensure that no further transactions occur, and then simply backing it up. You can then rebuild the database if it becomes damaged or corrupted in some way.

The rebuilding of the database is called *recovery*. *Version recovery* is the restoration of a previous version of the database, using an image that was created during a backup operation. *Rollforward recovery* is the reapplication of transactions recorded in the database log files after a database or a table space backup image has been restored.

Crash recovery is the automatic recovery of the database if a failure occurs before all of the changes that are part of one or more units of work (transactions) are completed and committed. This is done by rolling back incomplete transactions and completing committed transactions that were still in memory when the crash occurred.

Recovery log files and the recovery history file are created automatically when a database is created. These log files are important if you need to recover data that is lost or damaged. You cannot directly modify a recovery log file or the recovery history file; however, you can delete entries from the recovery history file using the PRUNE HISTORY command. You can also use the *rec_his_retentn* database configuration parameter to specify the number of days that the recovery history file will be retained.

Recovery Log Files and the Recovery History File



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.

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. It is used to track recovery-related events such as backup and restore operations, among others. This file is located in the database directory.

The *table space change history file*, which is also located in the database directory, contains information that can be used to determine which log files are required for the recovery of a particular table space.

Data that is easily recreated can be stored in a non-recoverable database. This includes data from an outside source that is used for read-only applications, and tables that are not often updated, for which the small amount of logging does not justify the added complexity of managing log files and rolling forward after a restore operation. *Non-recoverable databases* have both the *logretain* and the *userexit*

database configuration parameters disabled. This means that the only logs that are kept are those required for crash recovery. These logs are known as *active logs*, and they contain current transaction data. Version recovery using *offline* backups is the primary means of recovery for a non-recoverable database. (An offline backup means that no other application can use the database when the backup operation is in progress.) Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken and rollforward recovery is not supported.

Data that *cannot* be easily recreated should be stored in a recoverable database. This includes data whose source is destroyed after the data is loaded, data that is manually entered into tables, and data that is modified by application programs or users after it is loaded into the database. *Recoverable databases* have either the *logretain* database configuration parameter set to "RECOVERY", the *userexit* database configuration parameter enabled, or both. Active logs are still available for crash recovery, but you also have the *archived logs*, which contain committed transaction data. Such a database can only be restored offline. It is restored to the state it was in when the backup image was taken. However, with rollforward recovery, you can roll the database *forward* (that is, past the time when the backup image was taken) by using the active and archived logs to either a specific point in time, or to the end of the active logs.

Recoverable database backup operations can be performed either offline or *online* (online meaning that other applications can connect to the database during the backup operation). Database restore and rollforward operations must always be performed offline. During an *online* backup operation, rollforward recovery ensures that *all* table changes are captured and reapplied if that backup is restored.

If you have a recoverable database, you can back up, restore, and roll individual table spaces forward, rather than the entire database. When you back up a table space online, it is still available for use, and simultaneous updates are recorded in the logs. When you perform an online restore or rollforward operation on a table space, the table space itself is not available for use until the operation completes, but users are not prevented from accessing tables in other table spaces.

Having said that we have opted for a simple backup strategy to show you the mechanics of the backup and restore thru the GUI interface .We have taken of the following consideration :

Full backup and non recoverable which means that we cannot make a point in time recovery:

This activity should take place when nobody is connected it is advisable to do it for example at night , during the weekends

For different backup restore strategies please refer to the DB2 administration guide and to the DB2 command reference for the different command syntaxes .

Performing the backup using the control center

Log on as the instant owner from an administration client ..

It is recommended to update the statistics of the DataBase before doing the backup ..

Get a DB2 Command window (Start menu , db2cmd from a command window followed by db2) .

You can of course reorganize table by table . The method I use does in one shot and gives you a lot of information.

Now ,from the command window we issue the following command

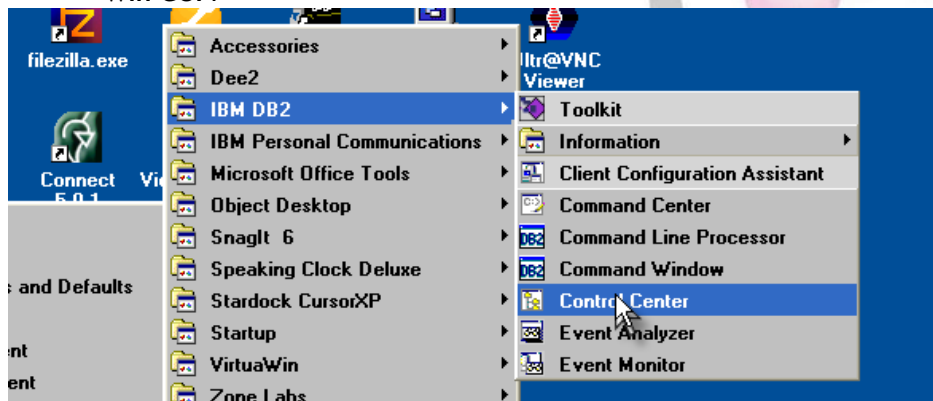
**REORGCHK UPDATE STATISTICS
ON TABLE ALL**

Now we start the control center

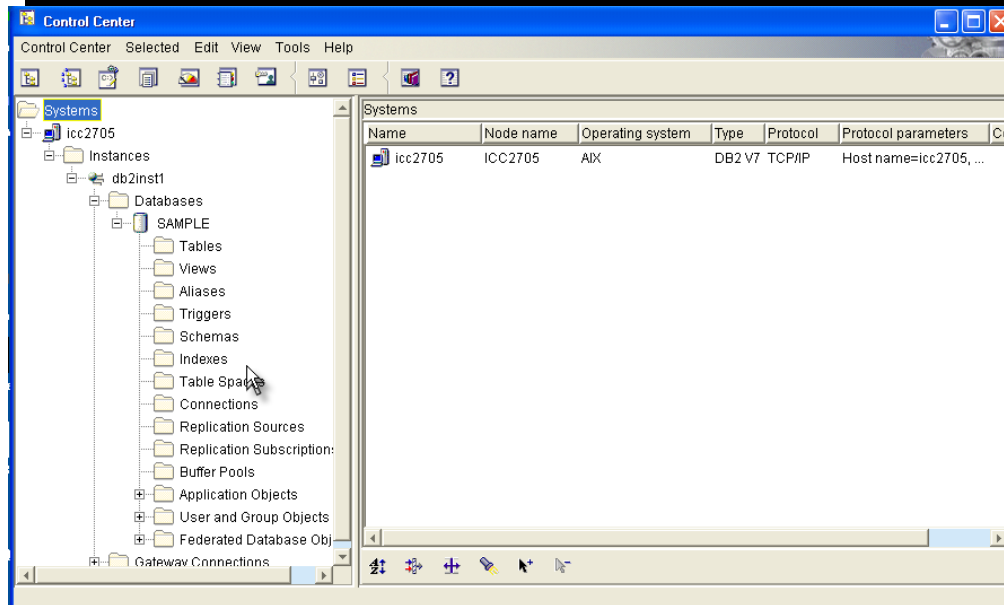
AIX user type db2cc from a command shell .

WIN from command line db2cc

Win GUI .



You will be presented with the Control Center that looks something like the following



Select the instance

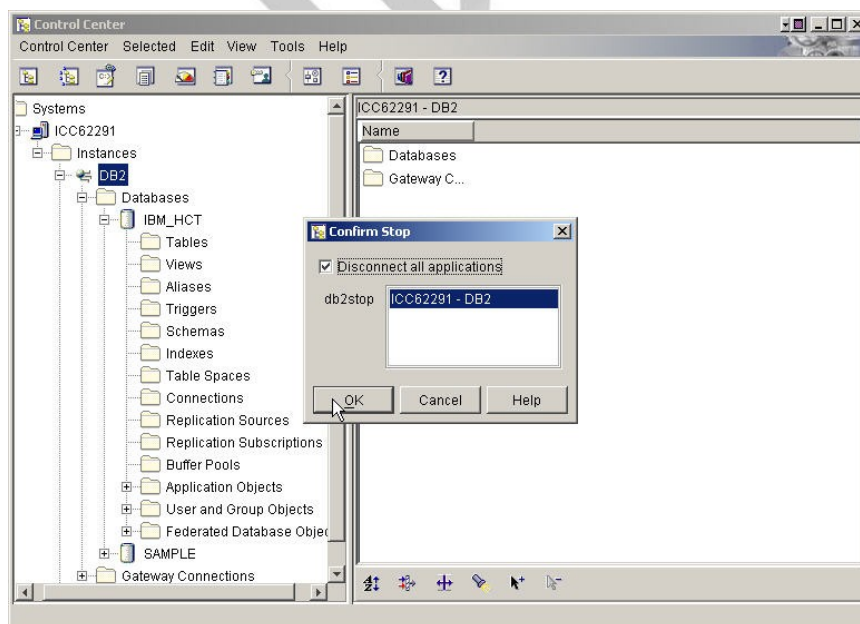
Right click → stop

You will be presented with the acknowledgement panel

Select the instance and check Disconnect all application .

OK

This procedure ensures that all the connections are closed

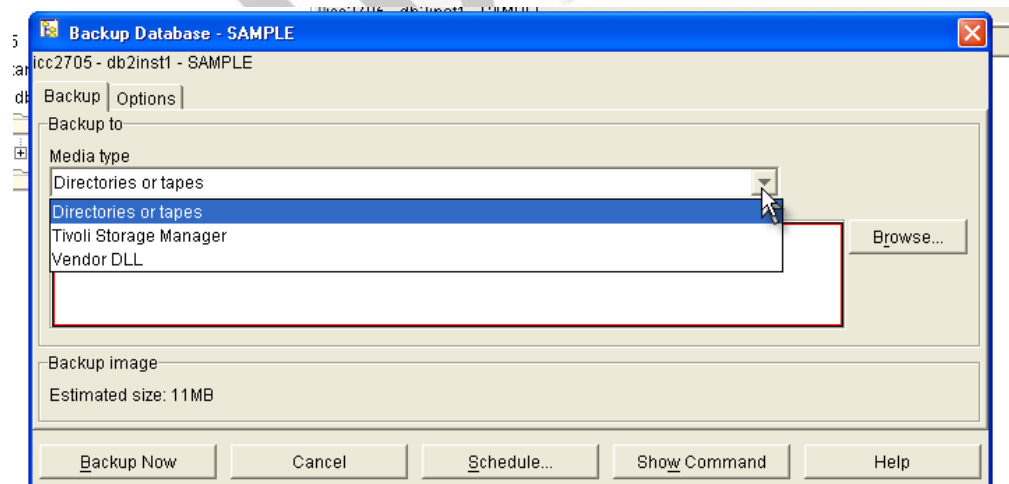
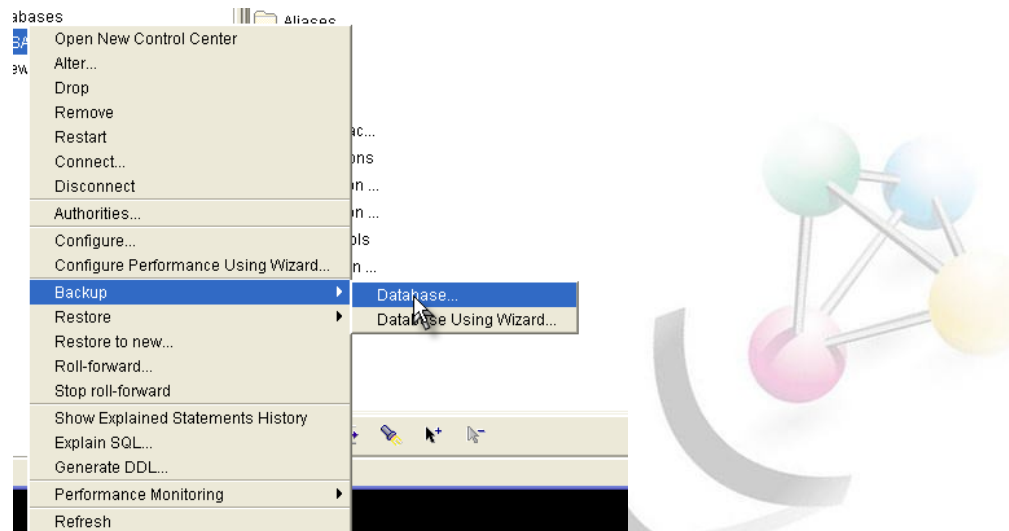


Select the instance

Right click → start

You will be acknowledged that the instance has started successfully .

Select the DB . → Backup → DataBase

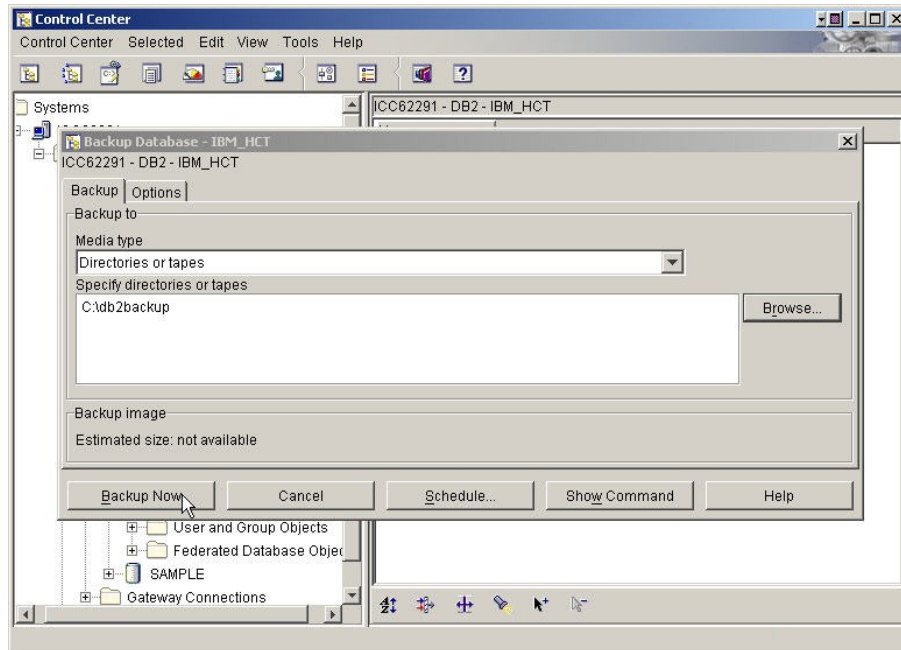


Pressing the arrow on the right for Media type you will be presented with different alternatives :

Directories or tapes

TSM if you are using Tivoli Storage Manager

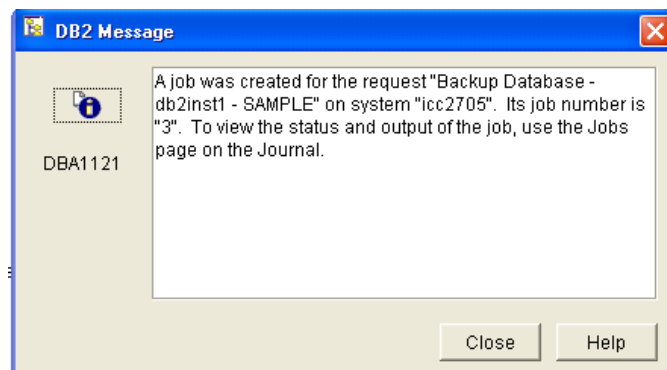
If you are using a different backup system other than TSM you should select Vendor DLL .



With the browse button we reach the directory where the backup will reside .

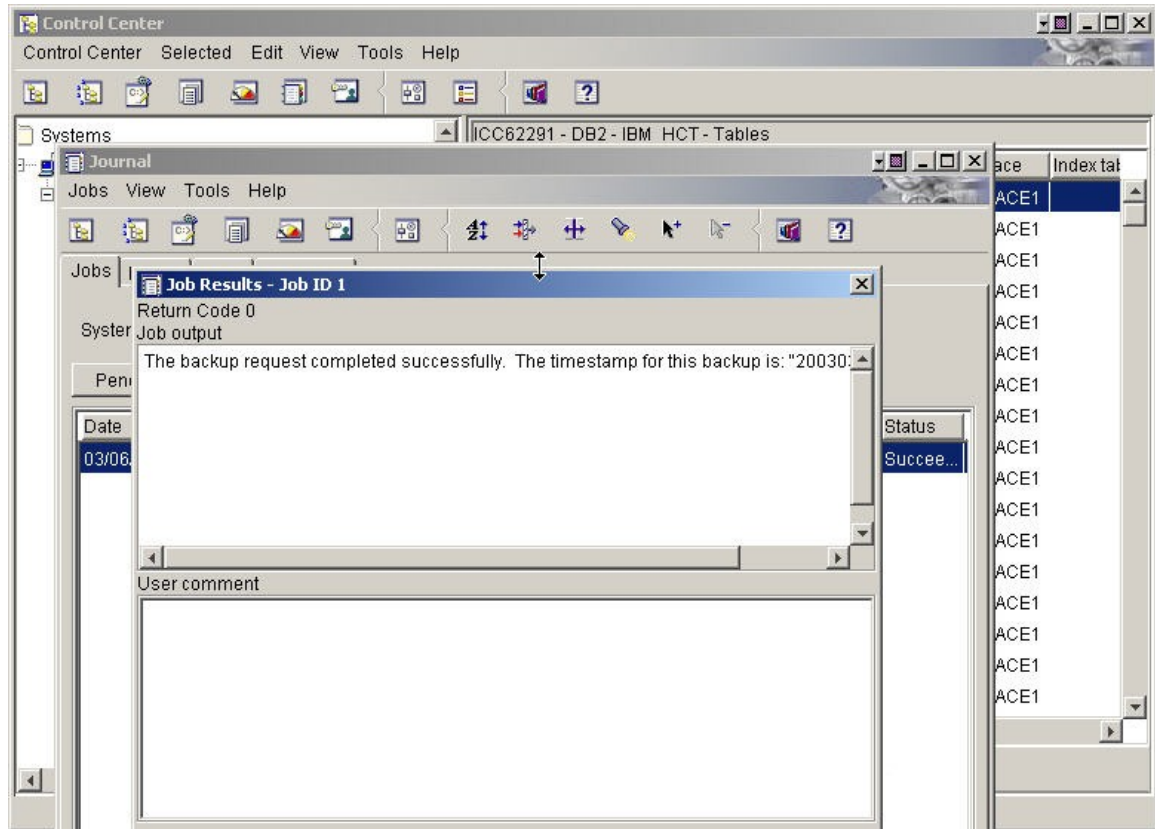
NB: This directory should exist or created before .

You will be presented with an acknowledgment that the backup job has been submitted
A message for the success should follow .



Status of the backup job can be retrieved from the journal .

To access the journal either press the journal button



Double click on the line containing the job number of the backup job .

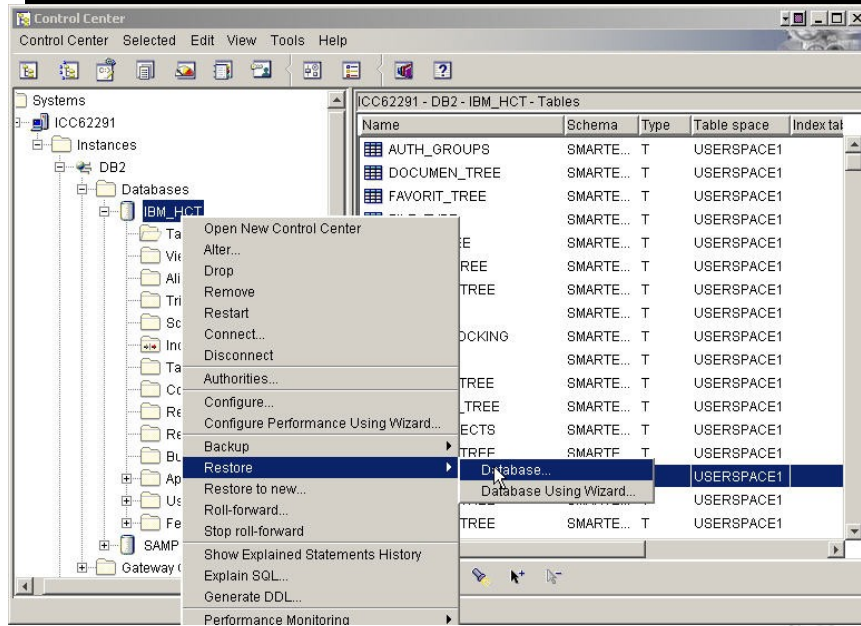
Performing the restore

The restore operation is the opposite of the backup command .. From the control center will select the restore image we want , then submit it for execution .

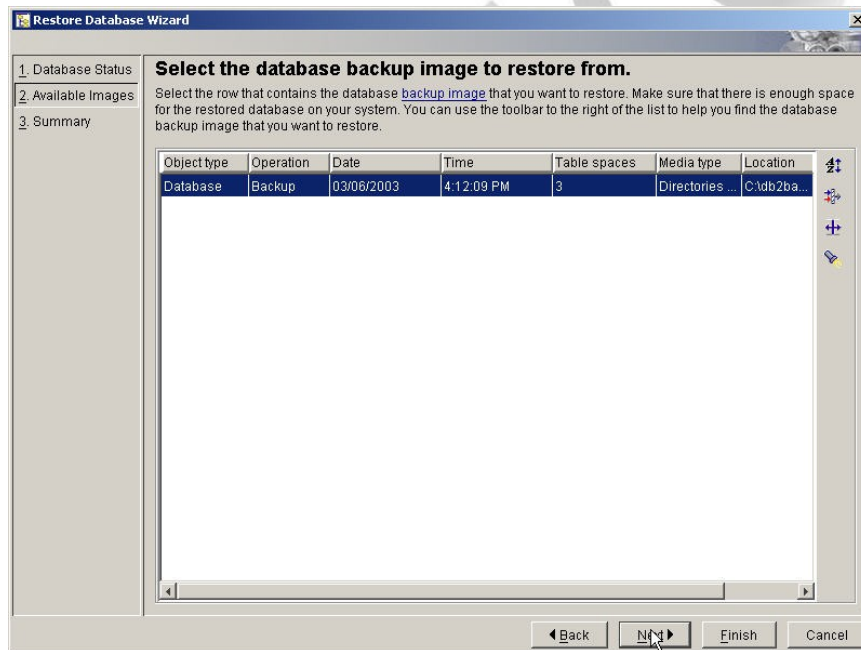
Open the control center .

Select your database

Right click → Restore → Database using wizard



Select the desired restore image and press next . Job will be submitted and you can see the result and messages in the Journal as previously in the backup step .



Basic DB2 tuning using the performance wizard

Tuning a db2 data base is a vast subject and depends upon different factors such as memory available , file systems , striping mirroring methods , No of CPU's and their speed , etc

It is beyond the scope of this document to enumerate the different parameters and their optimum settings . For further reading please consult the DB2 administration performance publication .

The objective of this document is to get you going quickly and efficiently and here are the steps to do so:.

➔ Log on as the instant owner .

It is mandatory to update the statistics of the DataBase before starting the performance expert due to the fact that the optimizer uses in his calculations the statistics available at that time .

➔ Get a DB2 Command window (Start menu , db2cmd from a command window followed by db2) .

You can of course reorganize table by table . The method I use does in one shot and gives me a lot of information.

Now from the command window we issue the following command

```
REORGCHK UPDATE STATISTICS  
ON TABLE ALL
```

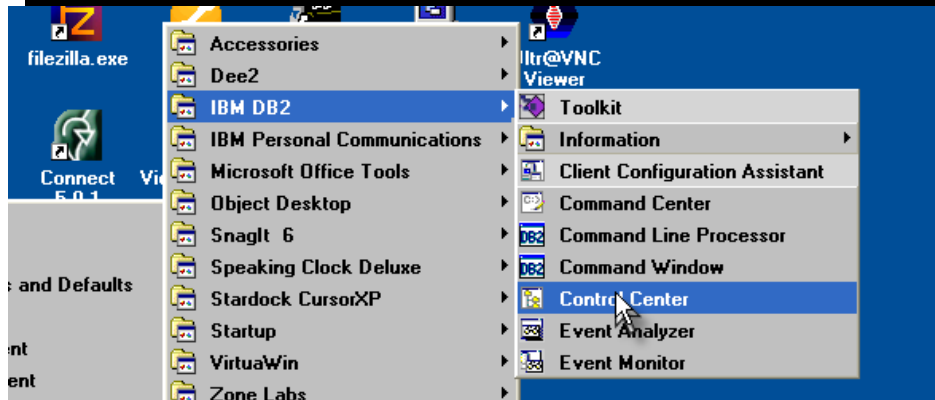
➔ MAKE A BACKUP OF THE DATABASE NOW !!!!!

Now we start the control center

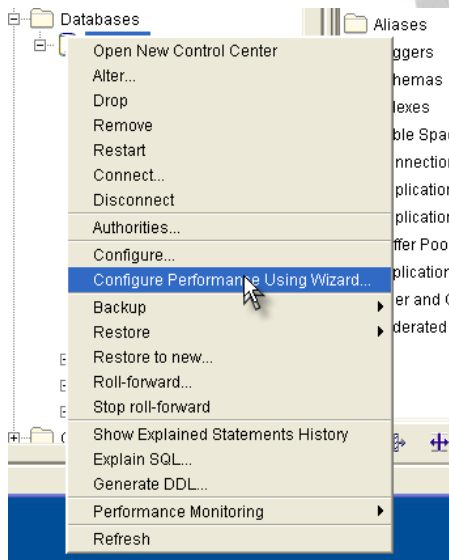
AIX user type db2cc from a command shell .

WIN from command line db2cc

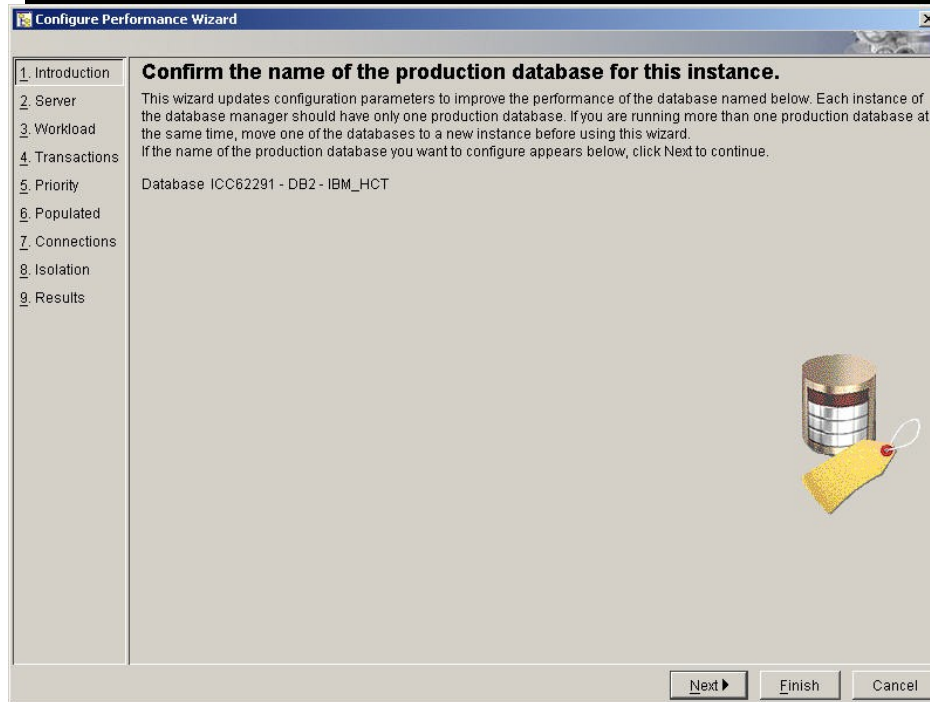
Win GUI .



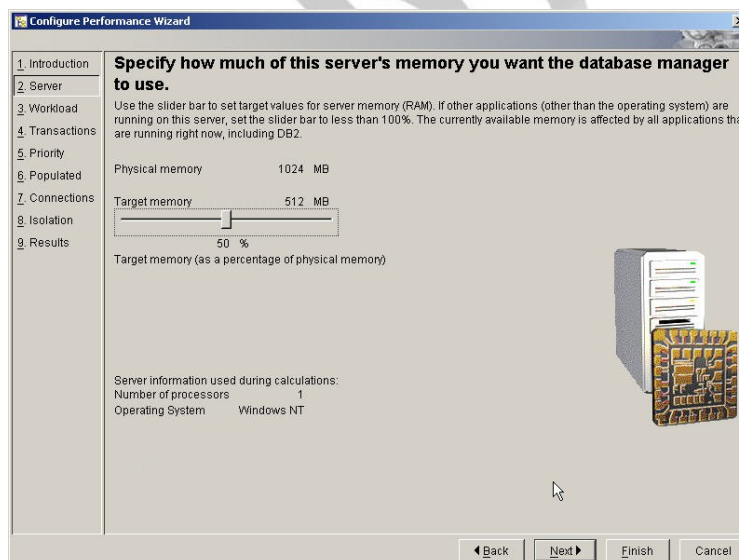
From the control center select your DataBase and click right and select Configure performance using wizard .



You will be presented with a series of dialogues for 9 steps .



Here you are asked to confirm the database name . Just to confirm .

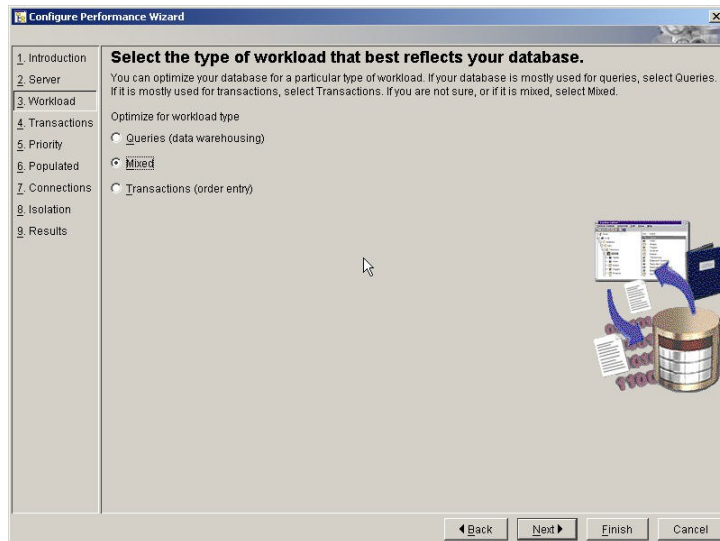


The wizard detects the total memory installed . At this stage a decision has to be made on how much memory will DB2 consume . This depends largely on the number of applications running on the server or is it a dedicated DB2 server .

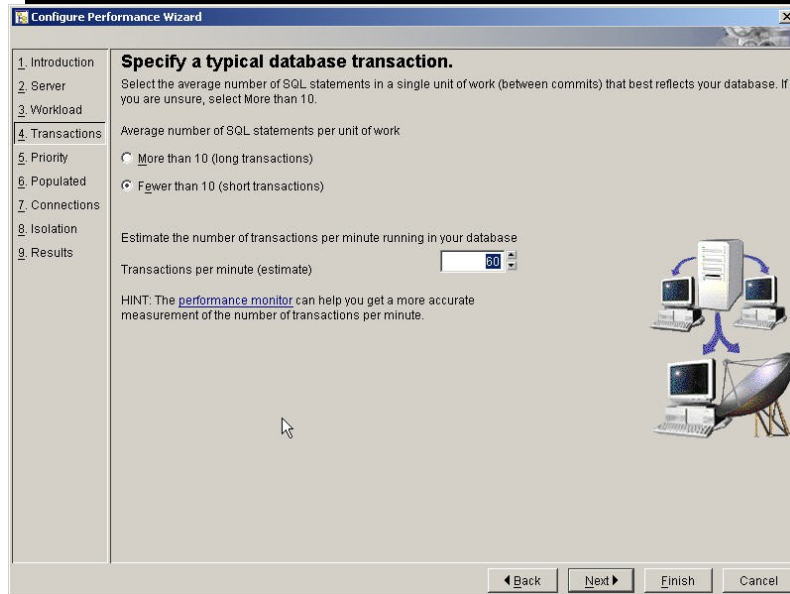
Using the slide bar adjust the percentage .

We have opted for 50 % assuming that SmarTeam is running as well on this server .

For a DB2 dedicated server you can go up to 75-80 % of the available memory .

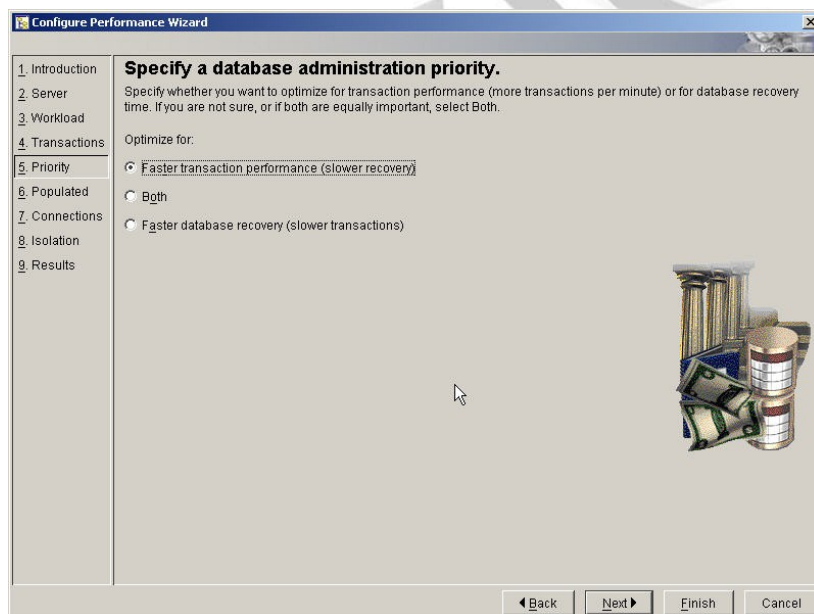


We are asked to select the transaction mixes either short transaction OLTP , Long queries , or a mix of both . SmarTeam application is a mixture of both

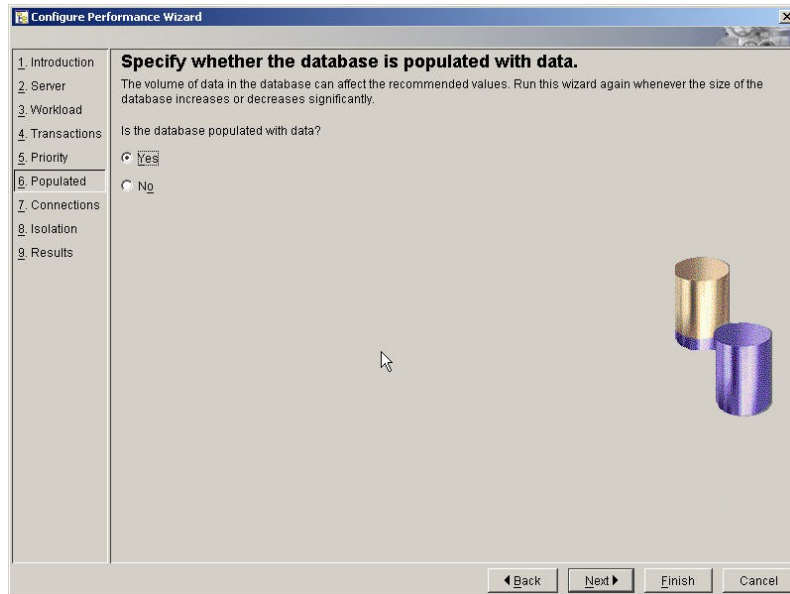


In this dialogue we are asked for a rough estimate of number of transactions per minute as well as an estimate of the Unit Of Work .

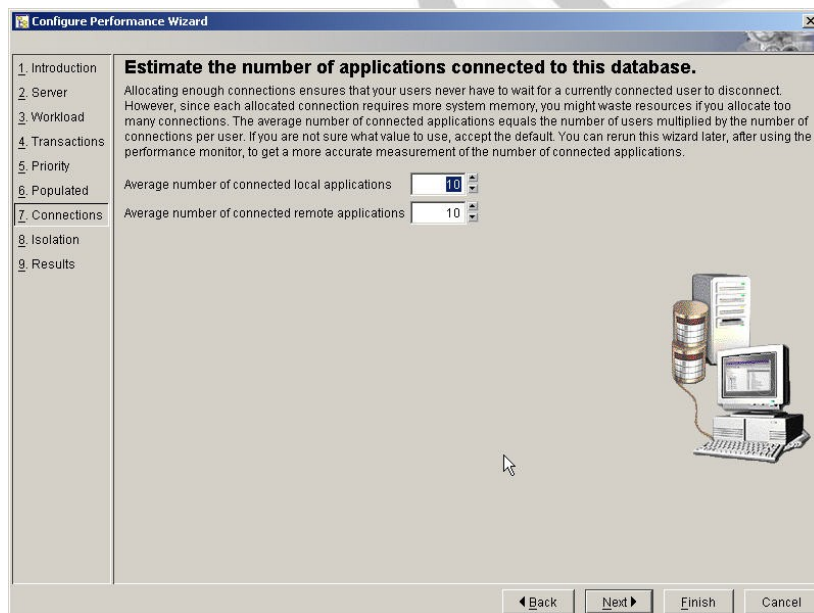
We have opted for a unit of work fewer than 10 and 60 transactions a minute.



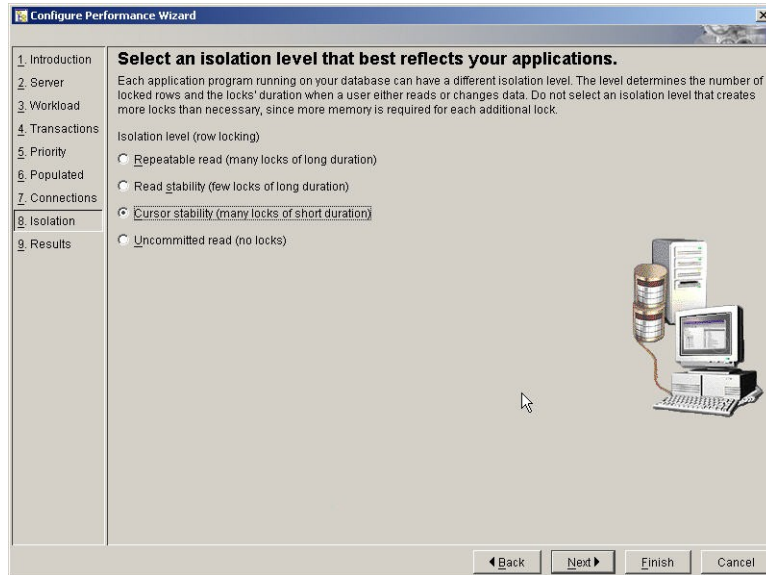
We specify here if we favor the fast recovery versus faster transaction or a bit of both . We have opted for faster transactions



Here we instruct the advisor to pick up the current statistics assuming that you have run **RUNSTAT** as prescribed earlier in this document .



We inform the performance wizard how many client are connected locally and how many are remote . These parameters are installation dependant . We have assumed for our exercise 10 local and 10 remote connections .



DB2 UDB offers many levels of concurrent access and uses different locks and different lock mechanisms . The recommended value here is cursor stability which means many locks for a short time . .

MAKE SURE THAT YOU MADE A BACKUP OF THE DATABASE IF NOT CANCEL AND DO IT NOW !!!!!

THE results shown here are examples and the results differ from machine to a machine , database to another , the size of the db , etc

Configure Performance Wizard

1. Introduction
2. Server
3. Workload
4. Transactions
5. Priority
6. Populated
7. Connections
8. Isolation
9. Results

Review the performance configuration recommendations.

Based on your selections in this wizard, as well as the volume of data in the database, and system information, this wizard recommends the following values. Below the list, specify if you want to save the new values to a script, or apply them to the database immediately.

Parameter	Current value	Suggested value	DB2 Parameter
Application ...	128	128	app_ctl_heap_sz
Buffer pool ...	250	49979	bufpage
Catalog cac...	32	902	catalogcache_sz
Changed p...	60	60	chngpgs_thresh
Database h...	1073742424	1762	dbheap
Default deg...	1	1	dft_degree
Default pref...	16	32	dft_prefetch_sz
Maximum s...	50	646	locklist
Log buffer s...	8	32	logbufsz
Log file size	250	250	logflsiz
Number of ...	3	6	logprimary
Number of ...	2	10	logsecond
Maximum n...	40	200	maxappls
Maximum l...	22	9	maxlocks

☒ Apply these recommendations immediately
☐ Save these recommendations to the Script Center

Script name:
 Script description: Configure Performance wizard recommendations

Back Finish Cancel

Configure Performance Wizard

1. Introduction
2. Server
3. Workload
4. Transactions
5. Priority
6. Populated
7. Connections
8. Isolation
9. Results

Review the performance configuration recommendations.

Based on your selections in this wizard, as well as the volume of data in the database, and system information, this wizard recommends the following values. Below the list, specify if you want to save the new values to a script, or apply them to the database immediately.

Parameter	Current value	Suggested value	DB2 Parameter
Sort heap s...	256	256	sortheap
Statement ...	2048	2048	stmtheap
Utility heap ...	5000	33402	util_heap_sz
Number of ...	1024	1024	fcm_num_buffers
Number of ...	512	512	fcm_num_rqb
Enable intr...	0	0	intra_parallel
Maximum q...	-1	1	max_querydegree
Maximum n...	200	200	maxagents
Agent pool ...	-1	-1	num_poolagents
Initial numb...	0	0	num_initagents
Sort heap t...	10000	10000	sheapthres
BP8 - Buffer...	250	25114	BP8
IBMDEFAU...	250	49979	IBMDEFAULTBP

☒ Apply these recommendations immediately
☐ Save these recommendations to the Script Center

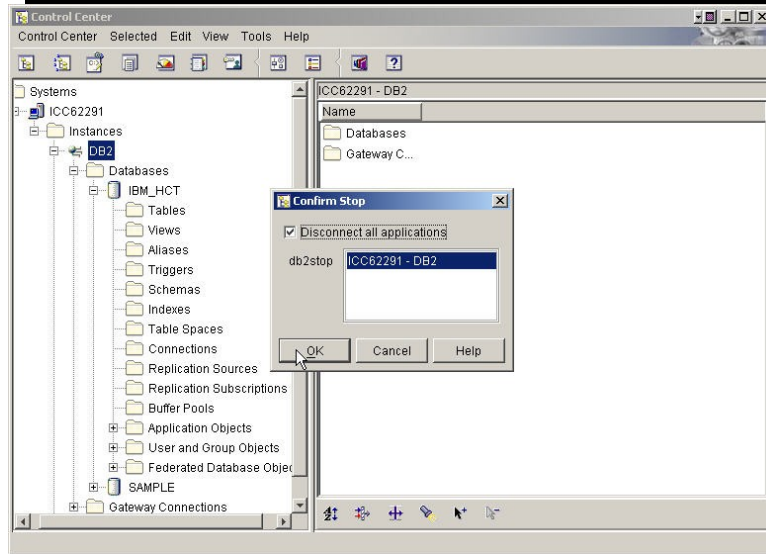
Script name:
 Script description: Configure Performance wizard recommendations

Back Finish Cancel

MAKE SURE THAT YOU MADE A BACKUP OF THE DATABASE IF NOT CANCEL AND DO IT NOW !!!!!

You are presented with a scrollable panel showing the original and in bold the modified values to be applied to the configurations .

You can cancel the job or click on finish where the modifications will be applied and comes into effect when you restart the database



Stop the database

Start the database ;

Now all modifications done and comes into effect .

appendix A : RDBMS Terminology mappings :

Mapping of Oracle Terminology to DB2 UDB

The following table gives readers who are familiar with Oracle a quick mapping to DB2 UDB terminology:

DB2 UDB	Oracle	Comments
DB2 UDB EE	Oracle EE	Enterprise product.
DB2 UDB EEE	Oracle Parallel	Support node partitioning.
DB2 Connect	Oracle Gateway	DRDA access to hosts.
SQL Control Statements	PL/SQL	Programming language extension to SQL. DB2 UDB stored procedures can be programmed in SQL Control Statements (subset of PSM standard), Java, C, C++, COBOL, Fortran, OLE, and REXX. DB2 functions can be programmed in Java, C, C++, OLE, or SQL control statements.
DB2 CLP	SQL*Plus	Command line interface to the server.
Instance	Instance	Processes and shared memory. In DB2 it also includes a permanent directory structure: an instance is usually created at install time (or can be later) and must exist before a database can be created. A DB2 instance is also known as the <i>database manager</i> (DBM).
Database	Database	Physical structure containing data. In Oracle, multiple instances can use the same database, and an instance can connect to one and only one database. In DB2, multiple databases can be created and used concurrently in the same instance.
DBM and database configuration files, etc.	Control files and .ora files	In Oracle, files that name the locations of files making up the database and provide configuration values. In DB2, each instance (DBM) and database has its own set of configuration parameters stored in a binary file; there are also other internal files and directories: none is manually edited.
Federated System	Database Link	In Oracle, an object that describes a path from one database to another. In DB2 a federated system is used. One database is chosen as the federated database and within it wrappers, servers, nicknames, and other optional objects are created to define how to access the other databases (including Oracle databases) and objects in them. Once an application is connected to the federated database it can access all authorized objects in the federated system.
Table spaces	Table spaces	Contains actual database data.
Containers	Data files	Entities inside the table spaces.
Objects	Segments	Entities inside the containers/data files.
Extents	Extents	Entities inside the objects/segments.
Pages	Data blocks	Smallest storage entity in the storage model.

N/A	Clusters	Data structure that allows related data to be stored together on disk; can be table or hash clusters. The closest facility to this in DB2 is a <i>clustering index</i> , which causes rows inserted into a table to be placed physically close to the rows for which the key values of this index are in the same range.
System catalog	Data dictionary	Metadata of the database.
SMS	N/A	System-managed table space.
DMS	Data files	Database-managed table space.
Buffer pools	Data cache	Buffers data in the table spaces to reduce disk I/O.
Package cache	Statement cache	Cache prepared dynamic SQL statements.
Log files	Redo logs	Recovery logs.
N/A	Rollback segments	Store the old version of data for a mutating table. In DB2 the old version of an updated row is stored in the log file along with the new version.
Database manager and database shared memory	SGA	Shared memory area(s) for the database server. In Oracle there is one, while in DB2 there is one at the database manager (instance) level and one for each active database.
Agent / application shared memory	UGA	Shared memory area to store user-specific data passed between application process and the database server.
Package	N/A	A precompiled access plan for an embedded static SQL application stored in the server.
N/A	Package	A logical grouping of PL/SQL blocks that can be invoked by other PL/SQL applications.

Appendix B. Most used DB2 Commands**Quick Reference: Common DB2® Command Line Processor****(CLP) Commands**

Maintained by Paul Yip IBM® Toronto Lab November 2002

The following is a quick reference for the most commonly used DB2 CLP commands. While not a complete reference, this document should prove valuable at the desk side of any new DB2 user. Feedback can be sent to ypaul@ca.ibm.com.

Arguments enclosed in square brackets are optional, and variables are denoted by enclosing angle brackets. For example, the syntax for `CONNECT` is:

```
connect to <dbname>
[ [user <userid>]
  using <pwd>]
```

This means that a `CONNECT` to database `sample` as `user1` using password `mypass` can have the following forms:

```
Connect to sample
Connect to sample user user1
Connect to sample user user1 using mypass
```

All commands are organized by their scope. For any CLP command, you may type in the keyword prefixed by a question mark (“?”) to see its full syntax.

Have fun!

. © 2002 International Business Machines Corporation. All rights reserved

Instance

CLP Command	Description
db2start	Starts the database manager instance.
db2stop	Stops the database manager instance.
get dbm cfg	Returns the database manager configuration settings.
get dbm cfg show detail	Displays current and delayed values of database manager parameters (as of V8).
update dbm cfg using <p> <v>	Updates the database manager configuration parameter <p> with value <v>.
get instance	Returns the value of DB2INSTANCE environment variable.
list active databases	Lists active databases and number of connections.
list application [show detail]	Returns information about the currently connected applications.
force application (h1 [,h2,...,hn])	Disconnects specific application(s) by handle number.
force application all	Disconnects all applications from the database.
attach to <node> user <userid> using <pwd>	Attaches to remote instance identified by <node> as user <userid> using password <pwd>.

¹ Some parameters take effect immediately. Others require instance stop and restart.

© 2002 International Business Machines Corporation. All rights reserved

Database

create database <dbname>	Creates a database called <dbname>.
activate database <dbname>	Explicitly activates the database.
Deactivate database <dbname>	Explicitly deactivates the database.
connect to <dbname> [[user <userid>] using <pwd>]	Connects to database <dbname> using explicit user <userid> and password <pwd> if needed.
connect reset	Disconnects from current database.
get db cfg show detail	Displays current and delayed values of database configuration parameters (v8 only).
zget db cfg for <dbname>	Returns the database configuration settings for database <dbname>.
update db cfg for <dbname> using <p> <v>	Updates the database configuration parameter <p> with value <v> for database <dbname>.
list tables [for {user all system schema <schemaname>}] [show detail]	Lists tables in the database. If no arguments are specified, the default is to list the current user's tables.
describe table <tablename>	Displays column information for a table or view.
list tablespaces [show detail]	Displays table space ID, name, type, contents and state.
list tablespace containers for <tablespace_id> [show detail]	Displays container information for the specified table space using <tablespace_id>.
quiesce tablespaces for table <tablename> reset	Resets the state of a table space to normal.

Connectivity

3catalog [admin] <protocol> node ...	Creates an entry in the node directory for .protocol <protocol>.
list [admin] node directory	Returns contents of the node directory.
4catalog database <dbname>...	Creates an entry in the database directory for database <dbname>.
list database directory [on <path>]	Returns contents of the database directory.

² Some parameters take effect immediately. Others require database deactivation and reactivation.

³ Use “? catalog <protocol>” to get full syntax. For example “? catalog tcip”. Be careful: It is possible to catalog a node that does not exist.

⁴ Use “? catalog database” to get full syntax. Be careful: It is possible to catalog a database that does not exist.

© 2002 International Business Machines Corporation. All rights reserved

Performance

get monitor switches	Returns the state of the session's monitor switches.
update monitor switches using <monitor> <on off>	Sets the state of the session monitor switch for <monitor>.
reset monitor all	Resets performance monitor values.
get snapshot for dbm	Returns performance information at the instance level.
get snapshot for all on <dbname>	Returns all performance information at the database level for database <dbname>.
get snapshot for dynamic sql on <dbname>	Returns the contents of the dynamic SQL cache.
runstats on table <tblschema>.<tblname>	Gathers statistics for table <tblname>. Table name must be fully qualified with <tblschema>.
reorgchk on table all	Determines if tables need to be reorganized. Useful for automatically performing <i>runstats</i> on all tables.
reorg table <tablename>	Reorganizes a table by reconstructing the rows to eliminate fragmented data and compacting information.

⁵ It is generally recommended to reset monitors before each performance test⁶ More snapshot variations are available. More information is available when monitors are on. ⁷Use db2rbind to rebind packages using latest statistics ⁸ Use db2rbind to rebind packages using latest statistics. Use REORG if tables or indexes need to be reorganized.

© 2002 International Business Machines Corporation. All rights reserved

Administration

export	Extracts database data into a flat file.
import	Imports data into the database using the IMPORT utility.
load	Imports data into the database using the LOAD utility.
load query table <tblname> [to local-message-file] [nosummary summaryonly] [showdelta]	Returns the progress of the LOAD utility.
backup database <dbname> [to <path>]	Performs a database backup.
restore database <dbname> [from <path>]	Performs a database restore.
get health snapshot for dbm	Returns health snapshot information for instance (v8 only).
get health snapshot for all on <dbname>	Returns all health snapshots for database <dbname> (v8 only).

Admin server

<code>get admin cfg</code>	Returns the admin server configuration settings.
<code>update admin cfg using <p> <v></code>	Updates the admin server configuration parameter <p> with value <v>.

Application development

<code>get routine into <filename> from [specific] procedure <routine-name> [hide body]</code>	Extracts SQL procedure to binary file.
<code>put routine from <filename> [owner <new-owner> [use registers]]</code>	Deploys SQL procedure from binary file.

° See “? export” for more information ¹⁰ See “? import” for more information ¹¹ See “? load” for more information ¹² If more than one backup image exists at the specified path, the TAKEN AT clause is required. ¹³ More variations of health snapshots are available. Use “? get health snapshot” for more information.

© 2002 International Business Machines Corporation. All rights reserved

IDSICC



International Competency Center

Optimize competitiveness through e-business technologies and services

The e-business challenge

The competitive challenge for management of manufacturing companies has moved on from ERP to Product Development Management. Delivering the right product to the right place, at the right time and at the right price is the best way to grow business.

Challenges are: faster customer service, greater product diversity, shorter product life cycles, globalization and just-in-time manufacturing.

Responsiveness is the key when business today is dependent on many different players. Most companies are both customers and suppliers for other companies at the same time, hence improving collaboration is essential in today's world.

IBM's e-business revolutionizes the supply chain by providing its customers with dynamic and instant access to the most current data and applications, allowing them to plan, schedule, analyze and make appropriate decisions linking together manufacturers, suppliers, distributors and customers, without regard to functional, geographic or enterprise boundaries. It is the **Digital Enterprise Vision**.

Irrespective of the size of your company, IBM's e-business solutions can give your current customers and millions of potential new customers access to your business at any time, from anywhere, all year-round.

Today, IBM is helping businesses around the world to maximize their supply chains by integrating product design, production, distribution and transportation functions into a seamless enterprise that links internal divisions, outside suppliers, distributors and especially customers into a single, strategically orchestrated entity.

Taking full advantage of opportunities creates business and requires the right planning and processes. This is where IBM Product Lifecycle Management can help.

IBM Product Lifecycle Management solutions based on CATIA, ENOVIA and SMARTEAM portfolio allow you to use data as a weapon to beat competition through innovation.

IBM solutions enable knowledge capture and reuse. They drive sharing of knowledge across the extended enterprise helping you to create faster, smarter organizations.

IBM business driven solutions leverage data within your company. IBM knowledge management solutions help you capitalize on human experience to create a more responsive culture. These solutions change your relationships with customers, employees and suppliers. They let you move from mass marketing to mass personalization, strengthen partnerships throughout your supply chain and therefore agility making you more responsive to change. You achieve maximum value as you make the transition from traditional business to e-business. Product Development Management projects may impact all the activities and services of the extended enterprise.

Considerations

- Availability
- Security
- Development of proposed solution
- Integration with existing systems
- Good systems exploitation conditions
- IT infrastructure

Digital enterprise requirements

- Global teams
- Distributed databases
- Function integration
- Support of
 - Digital mockup
 - Virtual product
 - Process
 - Resources
 - Model management



Deployment of CATIA, ENOVIA or SMARTEAM solutions provides you with a consistent functional and technical architecture to support the Digital Enterprise requirements.

IDSICC helps you to build the optimized total solution

In April 1999, IBM and Dassault Systèmes announced the establishment of the **IBM / Dassault Systèmes International Competency Center (IDSICC)** located at Dassault Systèmes headquarters in Suresnes, France.

A team of highly skilled specialists from both IBM and Dassault Systèmes staff the IDSICC. All of whom have extensive hands-on experience with implementing IBM and Dassault Systèmes Digital Enterprise solutions for customers. The center's mission is to provide worldwide technical experience and comprehensive Digital Enterprise systems recommendations to customers as well as to IBMers and IBM Business Partners to address every phase of development from research to implementation and production.

IBM is the only technology provider for CATIA, ENOVIA and SMARTEAM solutions to develop a competency center with this mission of total solution optimization.

IBM provides you with the most powerful combination available in the industry through the integration of key e-business technologies and services to implement the Digital Enterprise vision, with the added security of the support from the IDSICC.



Best of all, IBM offers the hardware, services and middleware needed to transform standard Computer-Aided Design (CAD) and Product Data Management (PDM) implementations into a true, Product Lifecycle Management engine for innovation.

IBM Hardware:

- e-servers : pSeries, xSeries, zSeries
- RS/6000, IntelliStation, Thinkpad,
- IBM TotalStorage products.

IBM e-business middleware :

DB2 Data Management Software

DB2 Universal Database offers a level of functionality and technology in Relational-Object support unmatched by any other database product in these engineering environments.

Lotus. software

Lotus Notes, Domino and Sametime foster efficient communication and collaboration throughout the entire process, sharing the knowledge inside and beyond the organization.

Tivoli software

Tivoli Systems Management delivers a unique life-cycle approach to managing clients by deploying software, monitoring resources, managing events and automating routine tasks in a centralized management way.

WebSphere software

WebSphere speeds up transition from publishing to web-based transactions, delivering a consistent open standard web application platform for CATIA and ENOVIA .

Skilled engineers from IBM labs as well as Business Consulting Services are available on request to become part of the core staff within the center working to provide integration support and total enterprise solution implementation for customers

Rational. software

Rational software delivers Software and systems development tools to assist you in getting better quality products . It covers the following areas :

Requirements & analysis, Visual modeling & development, Automated testing, Project management, Software configuration management, Configurable process

Total Digital enterprise Solution Optimization

This total solution environment proven by the IDSICC can be demonstrated at the centre running in real time on the Centre's own Installations.

The IDSICC provides skills transfers through webinars, system sizing, configurations, implementation and methodology guidelines, technical white papers and reports, assistance and participation in international customer events as COE, ECF and MICAD. This enables companies to deploy solutions more effectively and build higher quality products more quickly and with improved business efficiency.

IDSICC builds on IBM / Dassault Systèmes partnership and demonstrates IBM's commitment to deliver unmatched world-class technologies and services solutions for the competitive Digital Enterprise.

For more information:

IBM / Dassault Systèmes International Competency Center

fax :+33 (0)141449510 or visit our website:

<http://www.ibm.com/solutions/plm/idsicc>

email : idsicc@fr.ibm.com

IBM is a registered trademark of International Business Machines Corporation

CATIA, ENOVIA, DELMIA and SMARTEAM are registered trademarks of Dassault Systèmes.

All other products are trademarks or registered trademarks of their respective companies.

Information in this flyer is for guidance only.

Reference to IBM products, programs or services does not imply that IBM intends to make these available in all countries in which IBM operates.

IBM Software



DaimlerChrysler

→ steers toward deeper insights with IBM data warehouse

IBM software makes e-business work, from small businesses to enterprises. Our middleware software allows for interoperability between components of your infrastructure.

→ [Learn more about IBM softw](#)

IBM software brands

DB2. Information Management Software

Information, integration, storage
[DB2](#), [DB2 Content Manager](#), [IMS](#),
[Informix](#), [DB2 Information Integrator](#)

Lotus. software

Collaboration, communication
[Messaging](#), [Instant collaboration](#),
[e-Learning](#), [Collaborative portals](#),
[Knowledge & expertise](#)

Rational. software

Software and systems development
[Requirements & analysis](#), [Visual modeling & development](#), [Automated testing](#), [Project management](#), [Software configuration management](#),
[Configurable process](#)

Tivoli. software

Integration, automation, security
[Security](#), [Privacy](#), [Performance & availability](#), [Storage](#), [Configuration](#)

WebSphere. software

Dynamic e-business, integration
[Application servers](#), [Commerce](#),
[Portals](#), [Integration](#), [Application development](#)

IBM software advantage

Learn the big picture about IBM software

[Business Activity Management](#)

[Tivoli integrated identity management](#)

[zSeries Software and Tools](#)

IBM software in action

Read news and customer success stories

[IBM introduces first privacy language to automate policy compliance among enterprise systems](#)

[IBM expands software portfolio designed for medium-size business](#)
[More news](#)

[Terra Lycos brings IBM Lotus Instant Messaging to the masses](#)

[IBM, Linux rescue GE Interlogix-CASI security solution from platform paralysis](#)

[More customer success stories](#)

Software synergy

Extend the capabilities of your IT investment

Optimize

WebSphere and DB2 performance

Webcast 18 July

Speed-start Web services

Tools, training, and technical support

Introduction to DB2 UDB GUI

An Introduction to DB2 UDB Express GUI tools (Part 1)

[Raul F. Chong](#)

IBM Toronto Lab July 2003

© 2003 International Business Machines Corporation. All rights reserved.

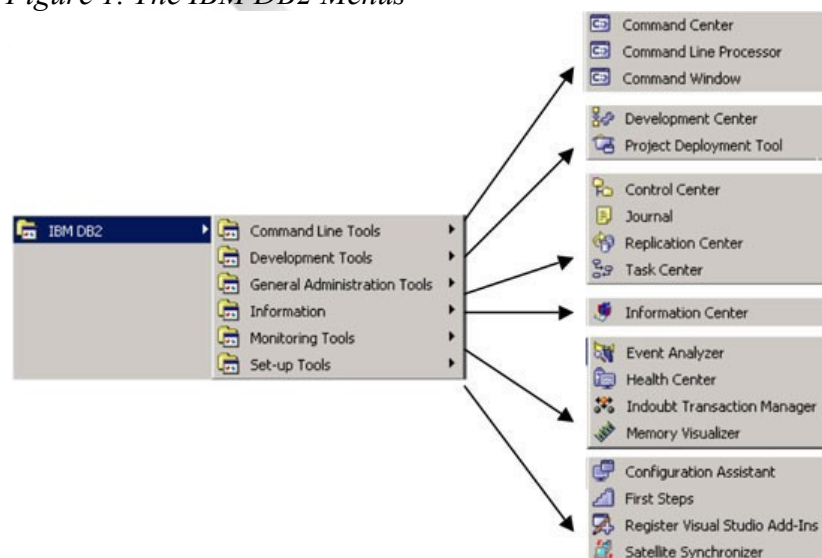
Introduction

If you are new to relational Database Management Systems (RDBMS), or have used other database products like Microsoft Access®, Fox Pro® or Microsoft SQL® Server, but have never worked with IBM® DB2® Universal Database™ (UDB), this document is for you. Using DB2 UDB Express -- the DB2 product specifically targeted to small and medium sized businesses -- we will show you, using a step-by-step approach, how to accomplish some of the most important database tasks. As with many other database products, database tasks can be performed in several ways. This document focuses on how to perform those tasks using only GUI Tools.

DB2 GUI Tools overview

Let's first start with a quick overview of the DB2 UDB Express GUI Tools. Figure 1 shows the different tools that can be invoked from the IBM DB2 Menu. We won't cover all of these tools in this article, just the ones you most likely will need for your day-to-day database operations.

Figure 1. The IBM DB2 Menus

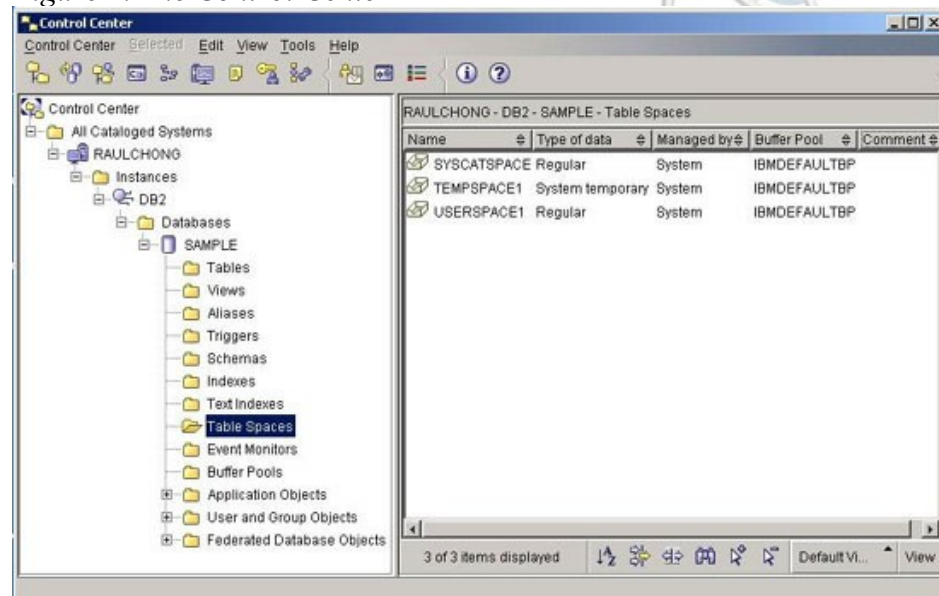


From the above menus, I'll use the following GUI Tools to describe how to perform daily database operations:

The Control Center

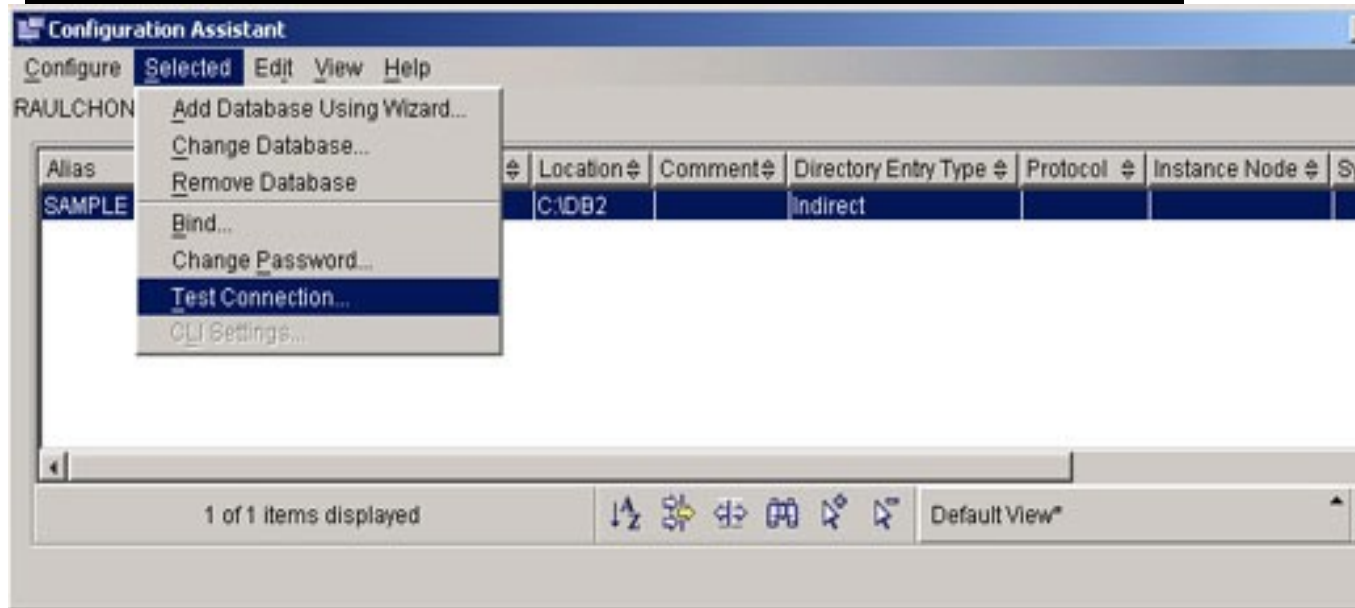
This is without a doubt the most important DB2 GUI tool. It provides you with a whole picture of your instances and databases, and allows you to perform most database operations in DB2. As you can see in Figure 2, the left panel (object pane) shows you the tree structure in your local and remote systems, and the right panel (contents pane) provides more detail about the specific item selected. We will describe the Control Center in more detail in the next sections.

Figure 2. The Control Center



The Configuration Assistant

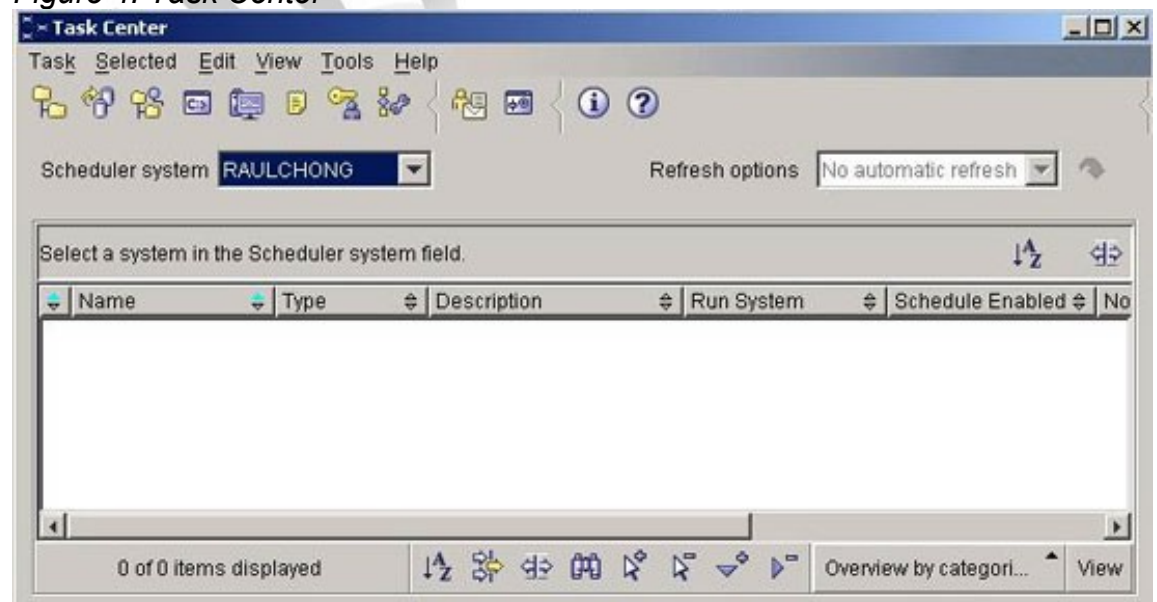
You have DB2 UDB Express installed in your desktop in Toronto, but would like to connect to a DB2 Server in Sao Paulo, Brazil. How can you connect to this remote server? The Configuration Assistant GUI tool can help you with the set up required, and you can also test that the connection is working using this tool. I'll describe how to use the Configuration Assistant in detail in the section [Setting up connectivity](#).



The Task Center

Say you would like to back up your database every day; however, you would prefer to perform this operation at 3:00 a.m. when there is no activity in your system. You can either get up at 3:00 a.m. every morning to perform this task, or you can save some sleep by automating the backup process. The task center can help you with this. Moreover, you can set up this task to send you an email if the operation was successful, and to page you if it was not.

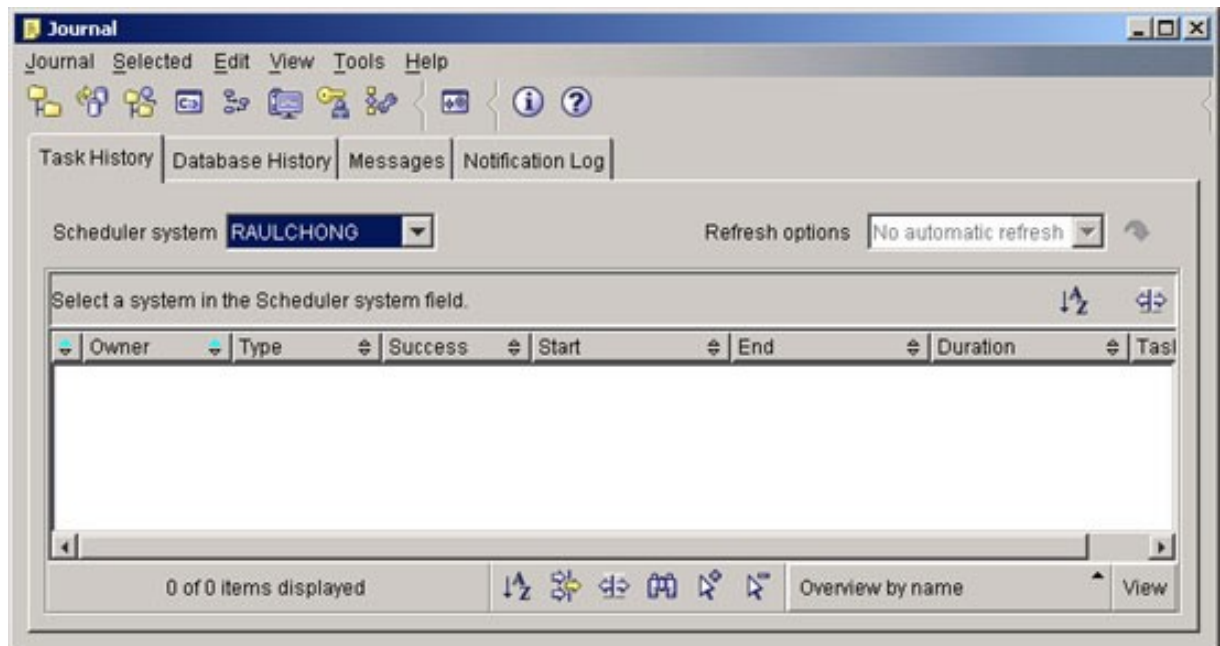
Figure 4. Task Center



The Journal

This GUI tool keeps track of all script invocations, DB2 messages, and DB2 recovery information. Should you need to investigate a problem, you can use this tool to find out what happened.

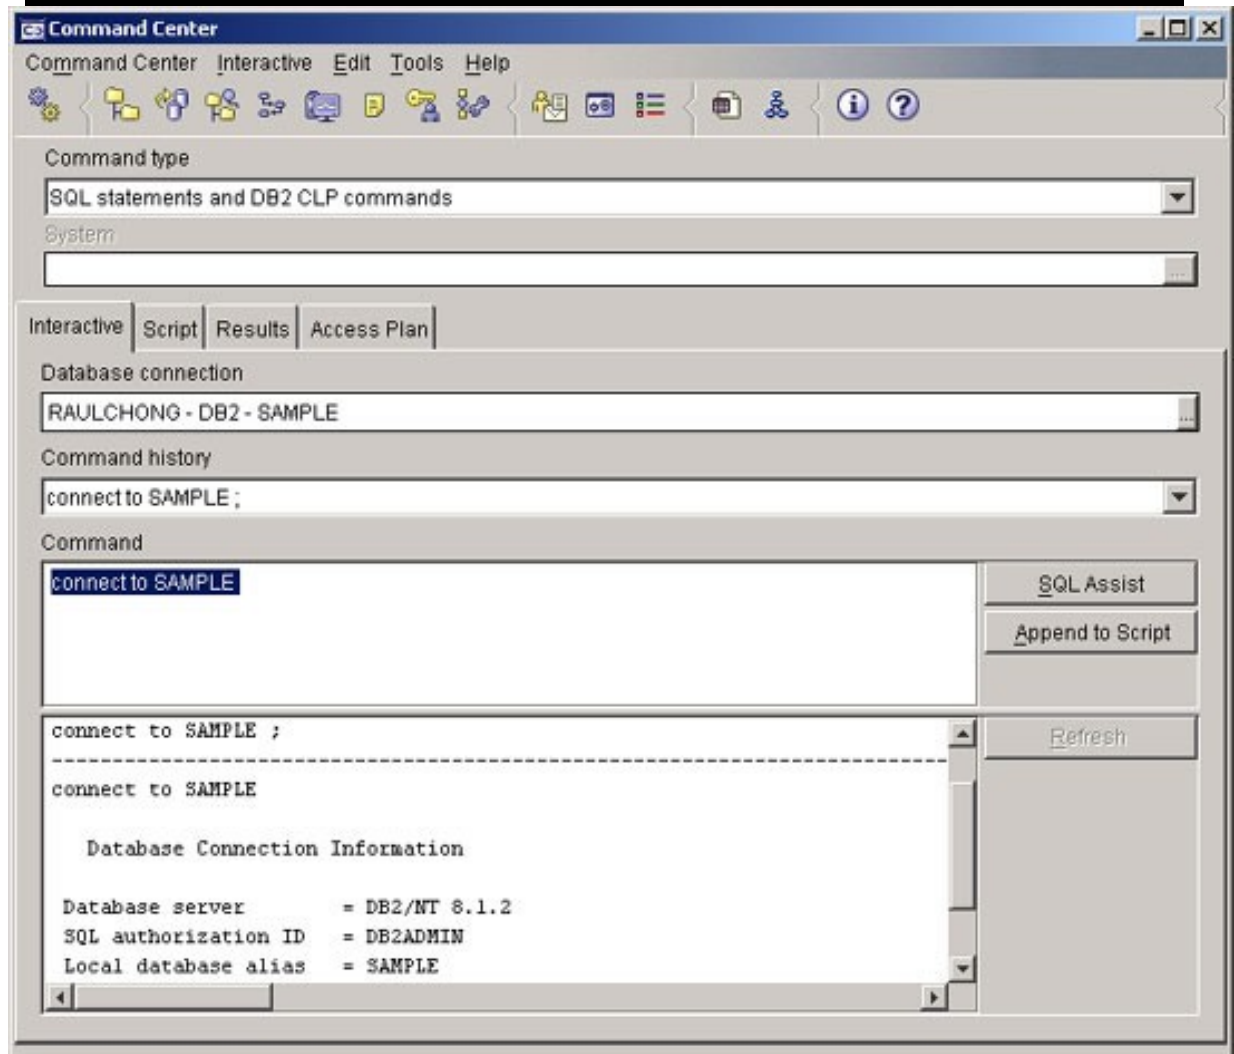
Figure 5. Journal



The Command Center

This tool allows you to input SQL statements or DB2 commands in an interactive window and see the results.

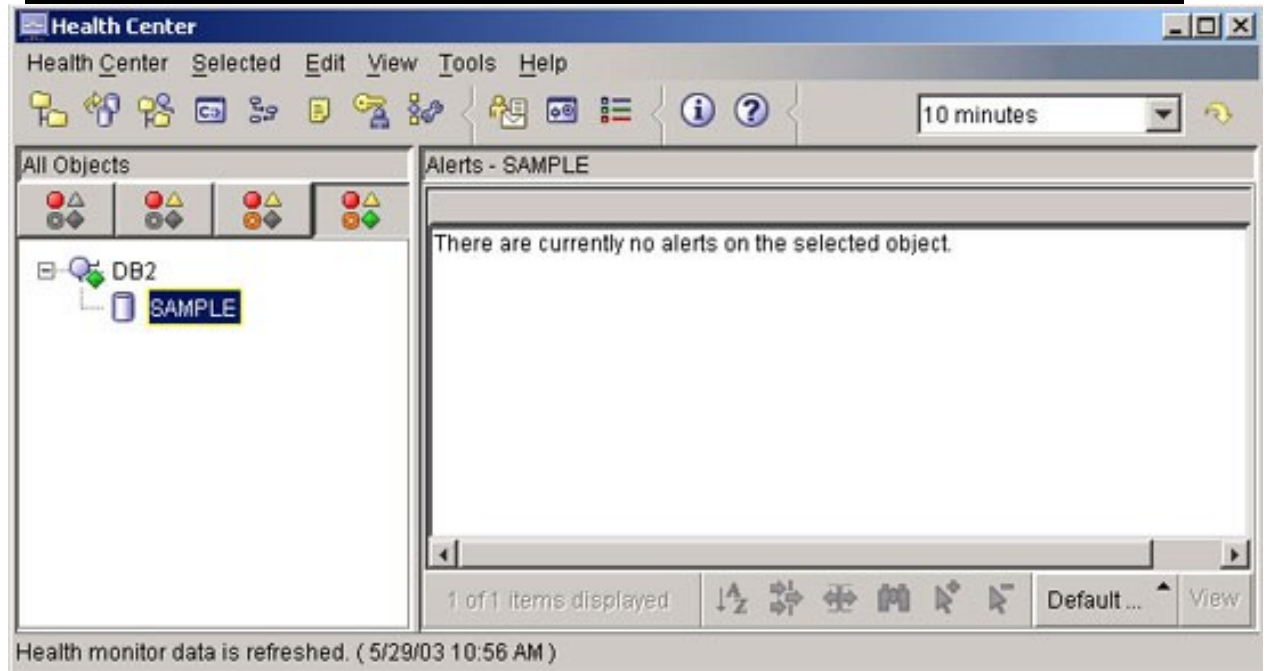
Figure 6. Command Center



Health Center

Use the Health Center GUI tool to set up thresholds that, when exceeded, will prompt alert notifications, or even actions to relieve the situation. In other words, you can have the database manage itself!

Figure 7. Health Monitor



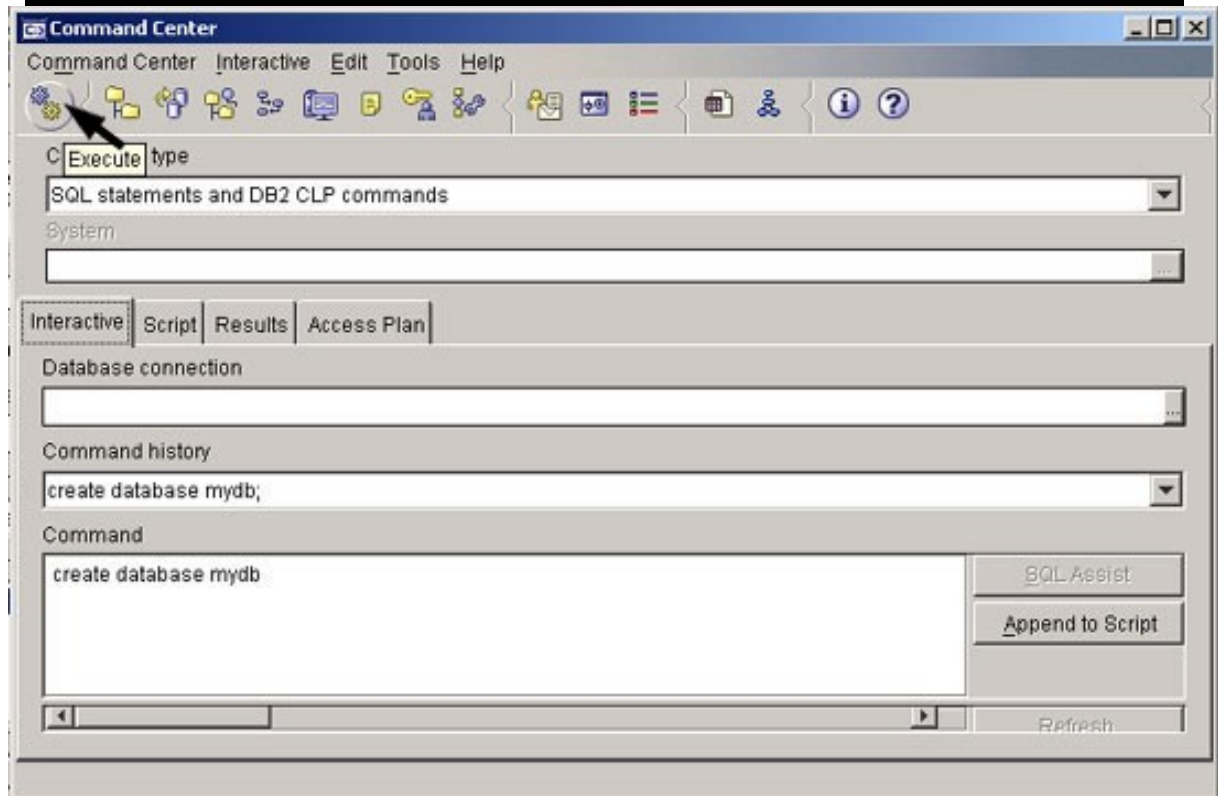
Performing basic database operations

In this section we will describe how to perform some basic operations using GUI tools:

Creating a database

The easiest way to create a database using all the defaults is to type the following from the Command Center and clicking on the **Execute** button as shown in Figure 8:

Figure 8. Create Database using the Command Center



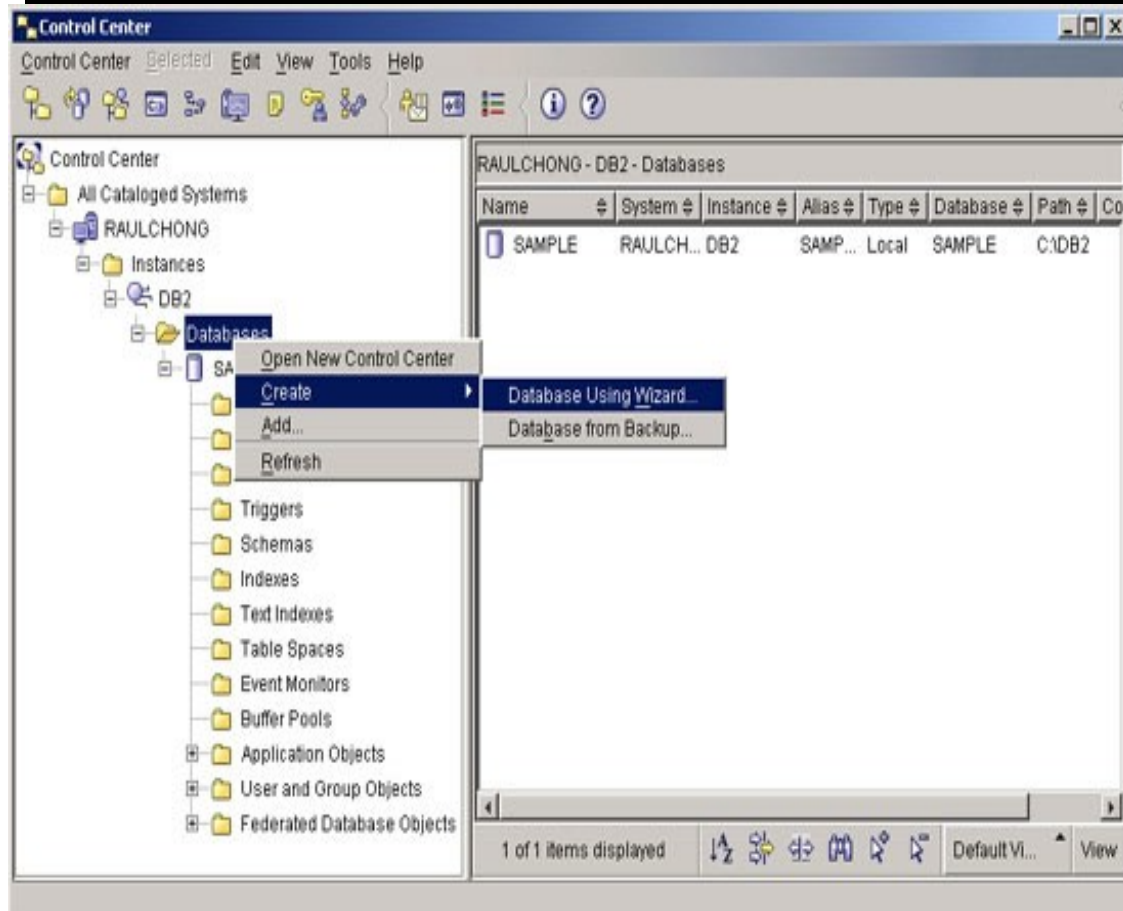
For the above example, the database 'mydb' will be created using all the defaults. Note that you must perform the `create database` command while you are *not* connected to any other database. If you were connected, execute the command `connect reset` prior to creating the database.

Another easy way to create a database which allows you to perform some customization is by using the 'create database wizard' from the Control Center. The three steps to follow are:

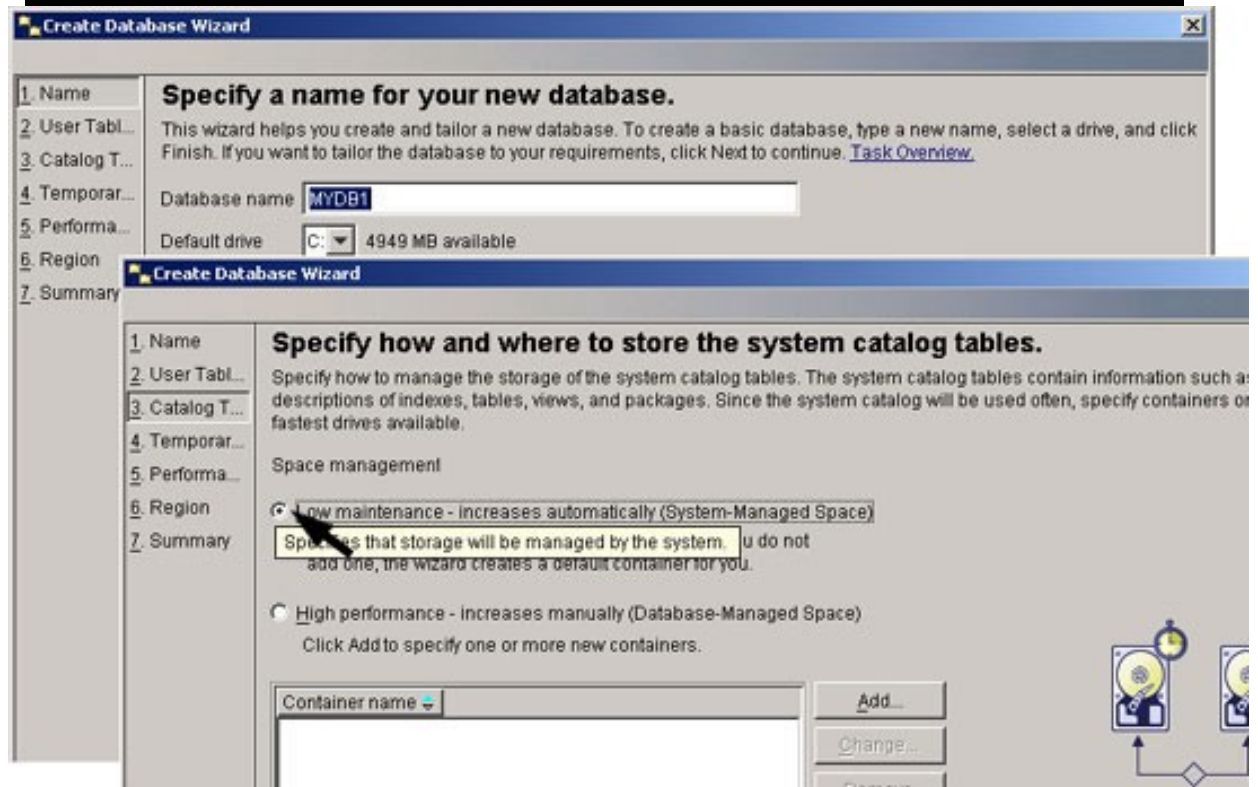
- 1.1. Click on **Databases** on the left panel of the Control Center
- 2.2. Right click on **Databases** and choose **create** -> **Database Using Wizard**.
- 3.3. Follow the panels from the create database wizard.

As you can see, like using any other GUI tool, performing operations with the Control Center is very intuitive.

Figure 9. Create database Wizard from the Control Center



The 'create database wizard' will walk you through several self-explanatory steps where you can specify the name and alias for the database, the type of table spaces to use for user, system catalog and system temporary tables, the locale and so on. Figure 10 shows some of the 'create database wizard' options.



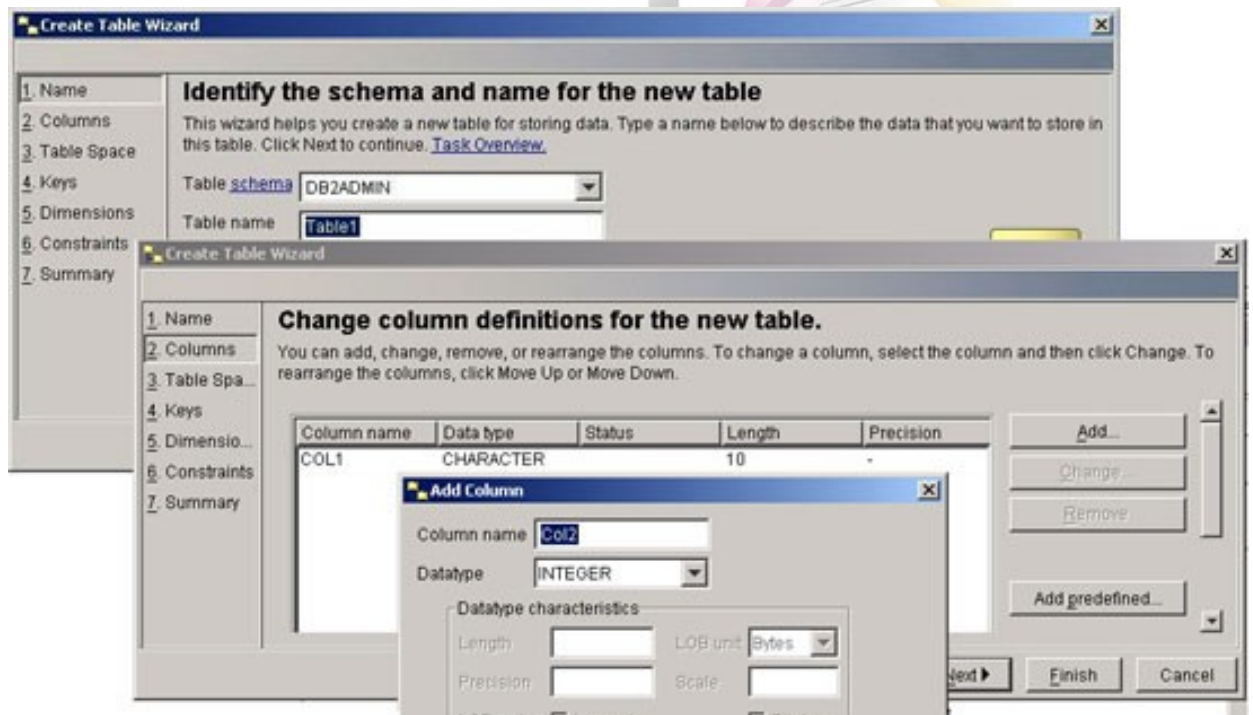
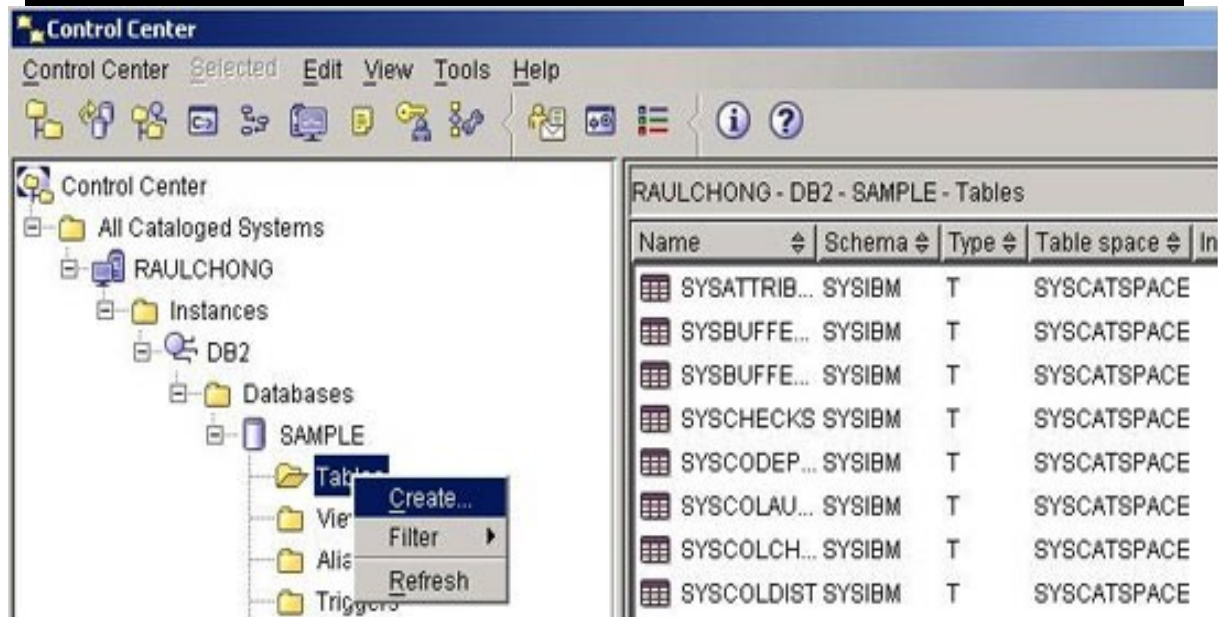
If you need more information about an option, you can either point your mouse to the given option, and hover help will come up. Alternatively, you can select the option in question and press F1. This action also displays the hover help. This is illustrated by the box in Figure 10 above containing the text 'Specifies that storage will be managed by the system.'

Creating a table

From the Control Center again we would need three steps:

- 1.1. Click on **Tables** on the left panel of the Control Center.
- 2.2. Right click on **Tables** and choose **Create**.
- 3.3. Follow the panels from the create table wizard.

Figure 11. Create Table from the Control Center



By now it should be obvious to you how an object can be created using the Control Center. Simply look for the object in the left panel, right click on it and choose create. Following this same procedure you can create aliases, views, indexes, triggers, schemas, and so on.

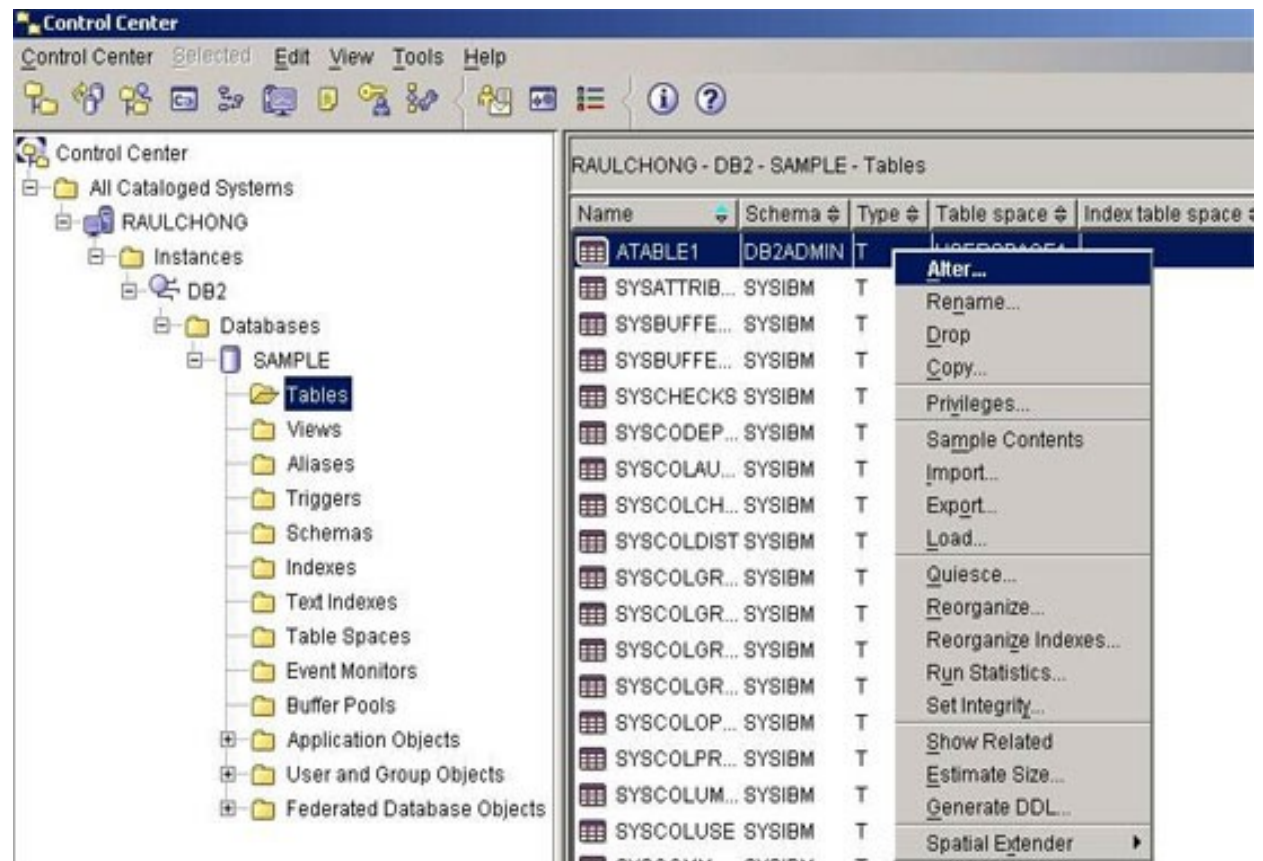
Altering a table

If you want to alter some of the characteristics of a table, you first need to

specify which table you want to alter. Thus, using the Control center's right pane (contents pane), follow these four steps:

<http://www7b.software.ibm.com/cgi-bin/wsdd/prnterfriendly.cgi> (11 of 20)23-Jul-03 12:51:45

- 1.1. Click on **Tables** on the left panel of the Control Center.
- 1.2. Click on the table name you want to alter.
- 1.3. Right click on the table name and choose **Alter**.
- 1.4. Follow the 'Alter Table' panels.

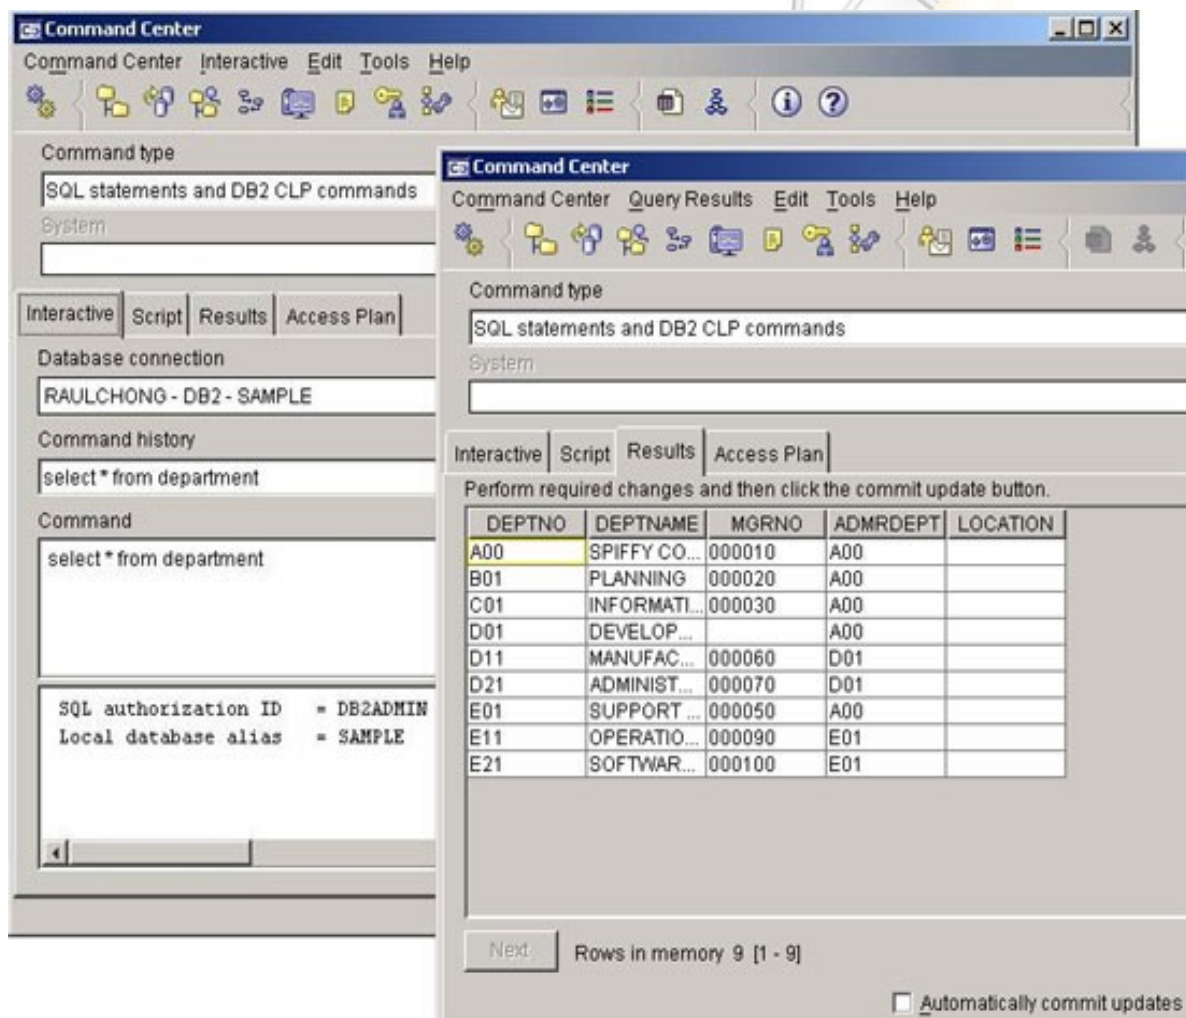


The example above used a table as the object to alter; however, a similar approach can be used to alter other types of objects. Note that not all attributes of an object can be altered. The GUI tool will only let you make the changes that are allowed. For example, in the case of a view, it can only be altered to add a comment. If you need other types of changes, you would need to drop the object and recreate it. As you can see from the above figure, the **Drop** option is also included in the menu when you right click on the object to be affected. From the menu, you can also choose **Sample Contents**, which will retrieve rows from your table; you can then alter some of the values directly if you want. Other operations like importing data to a table or exporting data from a table are also shown in the menu. We encourage you to explore these options on your own.

Performing queries

Now that you know how to create a database and its objects, you probably want to insert, update, delete or select information from these objects. The Command Center is the best GUI Tool to perform such operations. The figure below shows the Command Center with the **Interactive** tab selected. Based on the

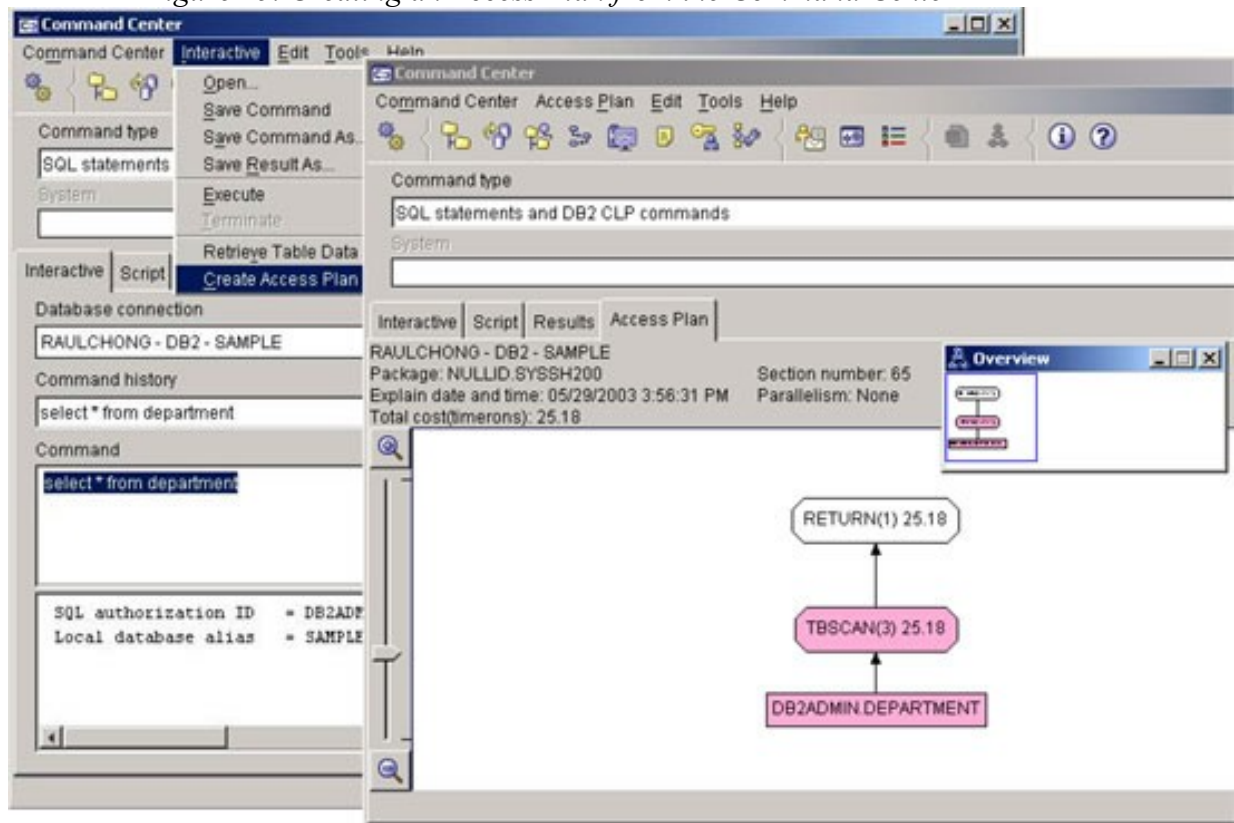
tab that is selected the menu bar will change its options. For Figure 14, a connection to the sample database has already been made. After the query 'select * from department' is executed, the output is immediately displayed in the results tab by default.



If you would like to see which access plan DB2 chose to retrieve a query, type the query in the Interactive tab, and then choose in the menu bar the option: **Interactive -> Create Access Plan**. If this is the first time you have

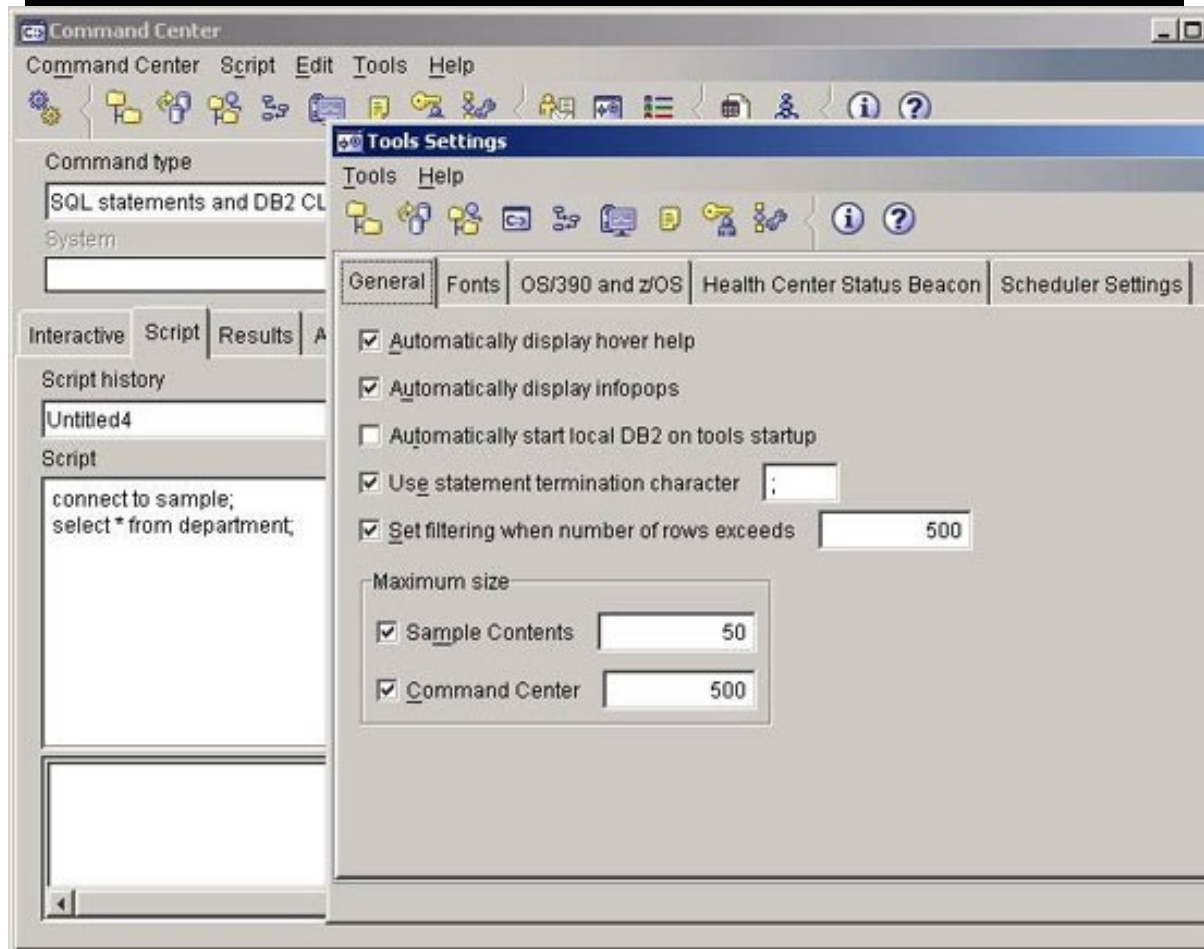
performed such an operation, you will get an informational message indicating some 'Explain' tables were created to store access plan information. Then you will note that the **Access Plan** tab is selected showing you the path chosen by the DB2 optimizer. Figure 15 below shows the access plan for the query 'select * from department.'

Figure 15. Creating an Access Plan from the Command Center



The Interactive tab of the Command Center is normally used when you want to perform one SQL statement at a time. If you would like to perform several statements at a time, you probably want to use the **Script** tab. Figure 16 shows the Command Center with the **Script** tab selected. For the example in the figure, two statements will be executed one after the other after clicking on the 'execute' button. Note that each statement ends with a semicolon. This is needed by default to indicate the end of a statement. If you choose in the menu bar the option **Tools -> Tools Settings**, you will see that the entry 'use statement termination character' is checked by default with a value of ';'.

Figure 16. Performing scripts with the Command Center



Restricting access to your objects

If you would like to keep confidential data inaccessible to regular users, you can always specify the privileges a user has against objects. Using the Control Center, it is easy to view which authorizations or privileges a user has. Follow these steps:

- 1.1. On the left panel of the Control Center and within the desired database tree, click on **User and Group Objects -> DB Users**.
- 2.2. Click on the desired user on the right panel of the Control Center.
- 3.3. Right-click on the desired user and choose change.

Figure 17. Managing Security using the Control Center

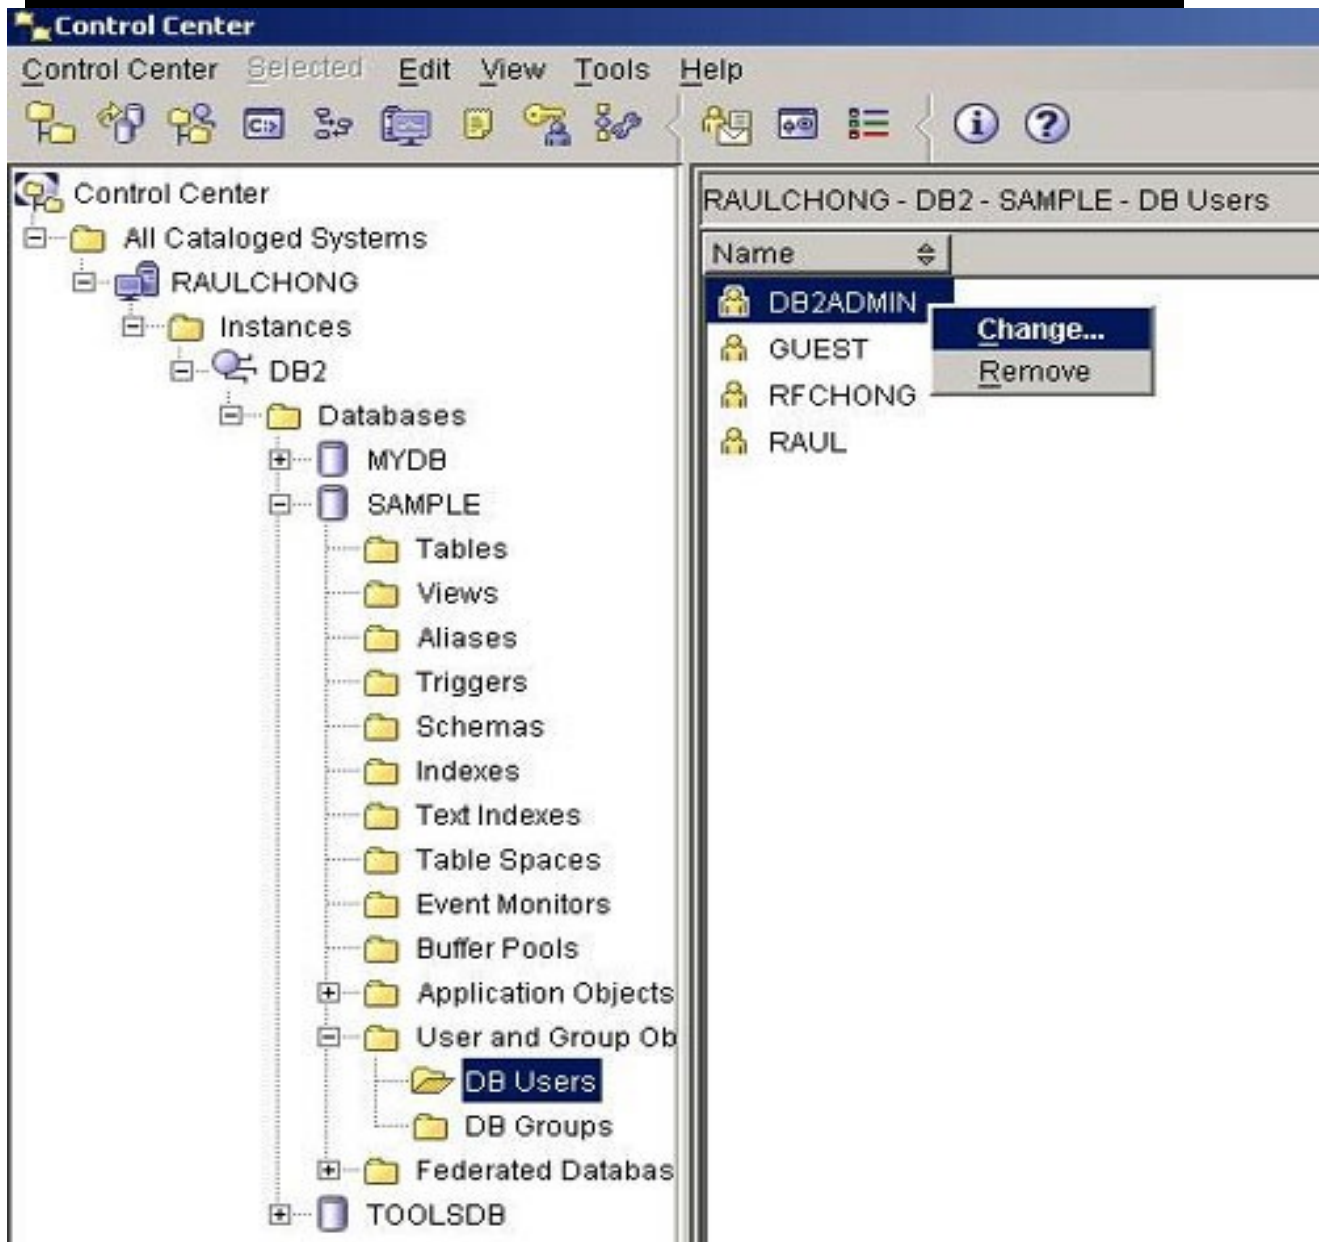
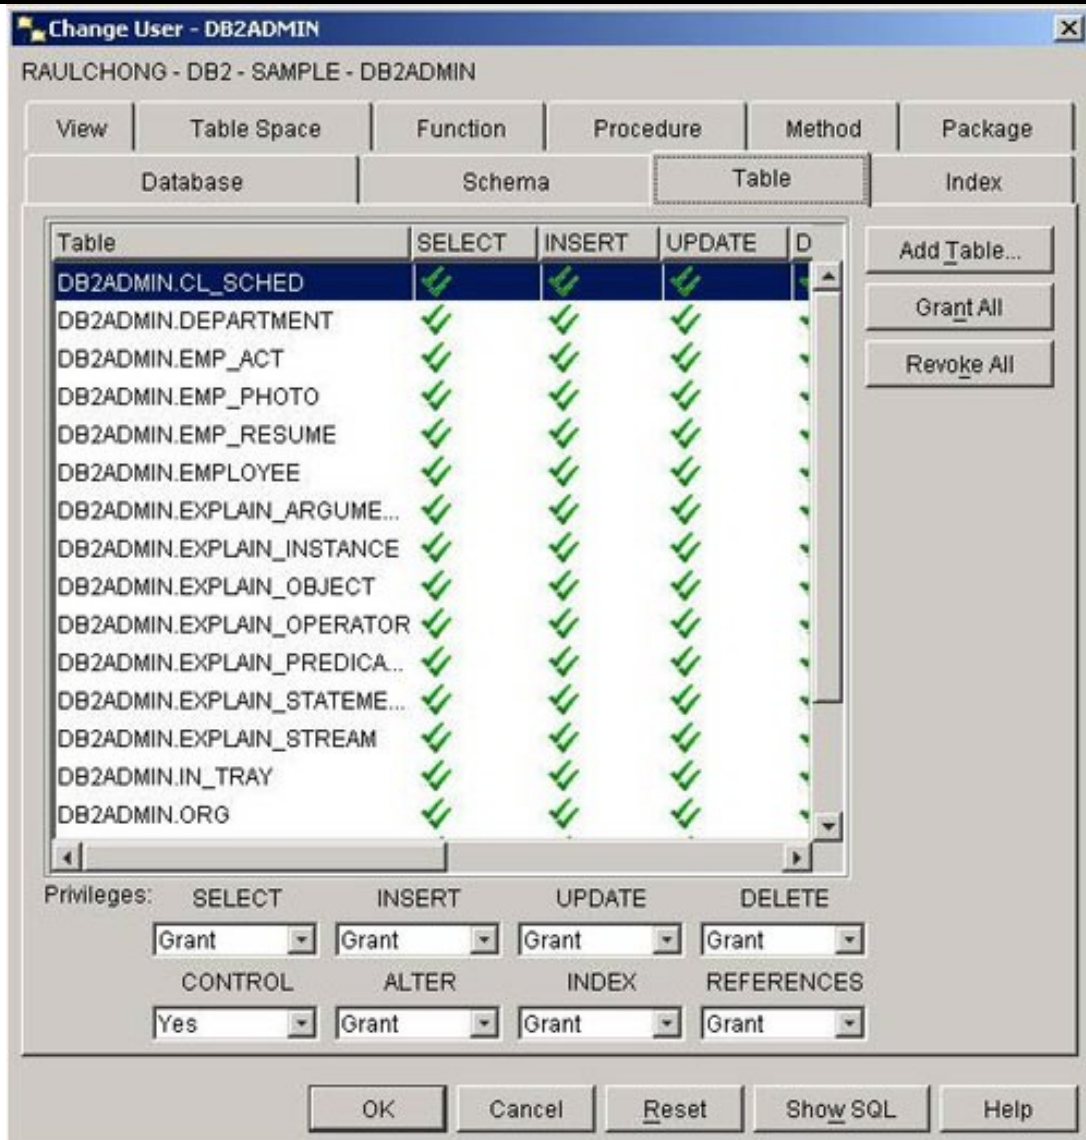


Figure 18. Managing Security using the Control Center (Continued)



Setting up connectivity

In this section we will describe in more detail the Configuration Assistant GUI Tool, which is used to set up the connectivity between your DB2 UDB Express server and other DB2 Servers.

Choose **Selected -> Add Database Using Wizard** from the menu bar as shown in the Figure 19. The 'Add Database Wizard' window comes up where you have three choices:

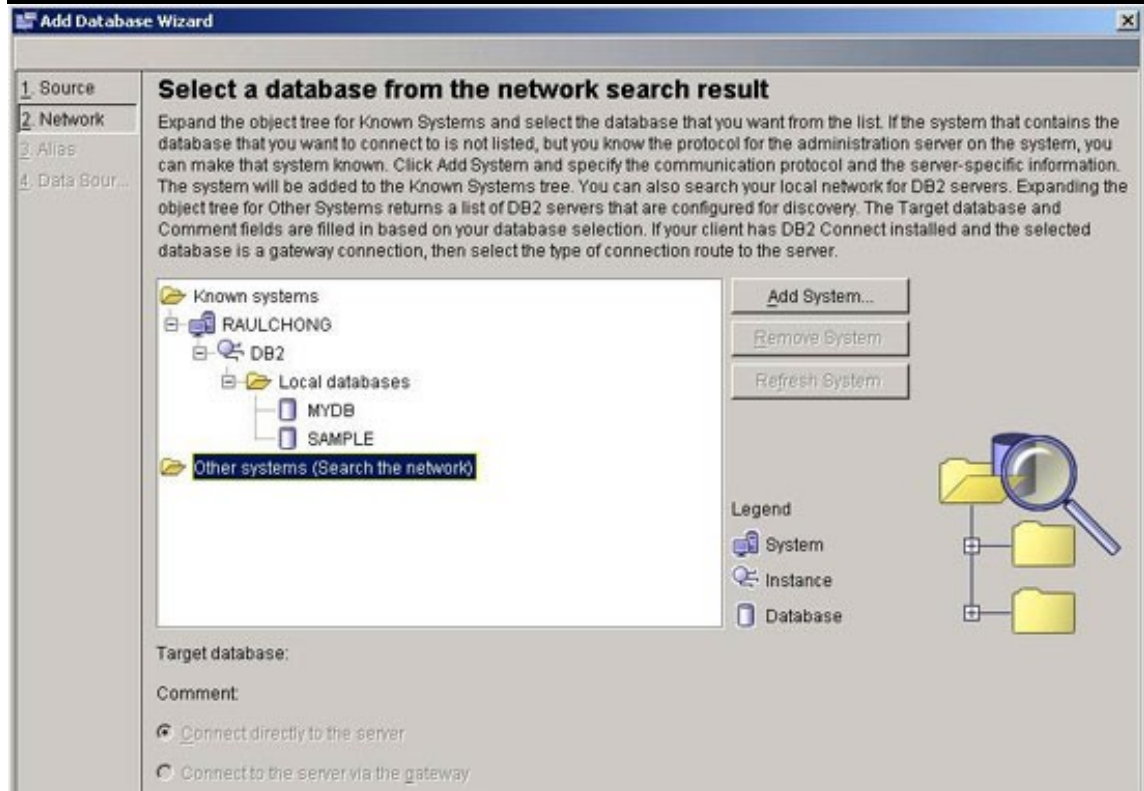
- 1.1. Use a Profile
- 2.2. Search the network
- 3.3. Manually configure a connection to a database

http://www7b.software.ibm.com/cgi-bin/wsdd/printerfriendly.cgi (17 of 20)23-Jul-03 12:51:45



The panel in Figure 19 contains a brief description of each option. Assuming your network is not large and does not contain many routers or hubs, the second choice 'Search the network' may be the easiest to use to set up connectivity to other DB2 Servers. Once you select this option and click **Next**, you will see the following panel:

Figure 20. Search the Network using the Configuration Assistant



You have the option to add the system in the 'Known Systems' folder by clicking on **Add System**, or alternatively, if you don't know the system you are trying to access, you can let DB2 search your entire network, and will display a tree with your server machine names, instance names and database names. All you need to do to establish connectivity to a particular database in this tree is to select it. The figure above shows the panel for 'Known Systems'. The tree structure would be similar in the case of 'Other Systems (Search the Network)'.

Summary

In this article we have introduced you to several DB2 UDB Express GUI Tools to perform basic database operations like creating a database, creating a table, altering a table, performing queries, restricting access to objects, and so on. The DB2 UDB Express GUI Tools are the same across the DB2 UDB family. With this article we have demonstrated how powerful and easy to use these tools are. We hope with this introduction you feel more comfortable working with DB2 UDB.

What's Next

Stay tuned for [Part 2](#) of this introduction article. In part 2 we'll cover the topics of database automation, basic performance tuning and GUI Tools troubleshooting. Again, the focus will be on DB2 GUI Tools.

We hope you have enjoyed your introduction to the DB2 UDB GUI tools and we welcome you to the world of DB2 UDB!

About the author

Raul F. Chong is a database consultant from the IBM Toronto Laboratory and works primarily with IBM Business Partners. Raul has worked for five years in IBM, three of them in DB2 Technical Support, and two of them as a consultant specializing in database application development and migrations from other RDBMS to DB2.

DB2, DB2 Universal Database, and IBM are trademarks or registered trademarks of IBM Corporation in the United States, other countries, or both.

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

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

Resources on the World Wide Web

IBM home page <http://www.ibm.com/>

IBM data management home page (English only) <http://www.software.ibm.com/data/>

DB2 home page (English only) <http://www.software.ibm.com/data/db2/>

DB2 Technical Library (English only) <http://www.software.ibm.com/data/db2/library/>

Developers Works <http://www.ibm.com/developersworks>

