

IBM DB2 Recovery
Expert for Multiplatforms



User's Guide

Version 1 Release 1

IBM DB2 Recovery
Expert for Multiplatforms



User's Guide

Version 1 Release 1

Note:

Before using this information and the product it supports, be sure to read the general information under "Notices."

First Edition (September 2002)

This edition applies to DB2[®] Recovery Expert for Multiplatforms, Version 1 Release 1 (program number 5724-B91) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2002. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

© Copyright Rocket Software 2002.

Contents

About this book	v
Who should read this book	v
How to send your comments	v
Online resources	v
Contacting customer support	v
Chapter 1. Introduction	1
Product description	1
Highlights	1
Required Privileges, Authorities, and Authorization	1
Hardware requirements	1
Software requirements	1
Required DB2 settings	2
Constraints	2
Chapter 2. Installation and configuration	5
To install DB2 Recovery Expert for Multiplatforms	5
Configuration	10
Configuration files	10
To configure DB2 Recovery Expert:	11
Configuring for a DB2 UDB EEE environment	12
Chapter 3. Getting started with DB2 Recovery Expert	13
Before using DB2 Recovery Expert for Multiplatforms	13
Launching DB2 Recovery Expert	13
Database status	14
Database status fields	14
Selecting a database	15
Exiting DB2 Recovery Expert	15
Other operations available in DB2 Recovery Expert	16
Accessing log analysis	16
Creating or updating the versioning repository	16
Chapter 4. Recovering a database	19
Database object selection	19
Database point-in-time selection	20
To select other points in time:	21
Database recovery path	23
Database recovery options	24
Redirected restore	25
Database summary	26
Chapter 5. Recovering a table space	29
Table space object selection	29
Table space selection display columns	30
Table space point-in-time selection	31
To select other points in time:	32
Table space recovery path	32
Table space recovery options	32
Table space summary	32
DMS considerations	33
Chapter 6. Recovering a table.	35

Table object selection	35
Table selection display columns	36
Table point-in-time selection	36
Table recovery path	36
Table recovery options	36
Table summary	37
Chapter 7. Log analysis	39
Getting started with log analysis	39
Generating the report	41
Log analysis in the report	43
Generating the details	44
Reviewing the details	45
Reviewing the SQL	46
Chapter 8. Recovery Expert Versioning Repository command line interface	49
Versioning Repository	49
Versioning Repository update command	50
Saving DDL	50
Chapter 9. DB2 minilog management tool command line interface	53
Minilog management tool command	53
Command line examples	55
Example 1	55
Example 2	55
Minilog management tool output	55
Screen output	55
File output	55
Scenarios for use	57
Scenario 1	57
Scenario 2	57
Scenario 3	57
Chapter 10. Log analysis command line interface	59
Log analysis command	59
Command line examples	61
Example 1	61
Example 2	62
Sample reports	62
Chapter 11. DB2 object translation tool command line interface	69
DB2 object translation tool command	69
Object translation constraints	70
Command line examples	70
Example 1	70
Example 2	70
Appendix. Notices	73
Trademarks	75
Index	77

About this book

This book contains information on installing and using DB2 Recovery Expert for Multiplatforms.

Who should read this book

This book is intended for all users of DB2 Recovery Expert for Multiplatforms.

How to send your comments

Your feedback is important in helping to provide the most accurate and high-quality information. If you have any comments about this book or any other IBM® DB2 Recovery Expert documentation:

- Use the online reader comment form located at www.ibm.com/software/data/db2imstools/rcf/.
- Send your comments by e-mail to dmtinfo@us.ibm.com. Be sure to include the name of the book, the part number of the book, the version of DB2 Recovery Expert, and, if applicable, the specific location of the text that you are commenting on (for example, a page number or table number).
- Print and fill out the reader comment form located at the back of this book. You can use the following methods to return the form:
 - Print and fill out the reader comment form located at the back of this book. You can use the following methods to return the form:
 - Give it to your local IBM branch office or IBM representative.
 - Send it to the address printed on the reader comment form
 - For customers in the U.S. and Canada, fax it to 1-800-426-7773.

Online resources

For a complete and up-to-date source of DB2 Recovery Expert information, including information on issues discovered after this user's guide was published, visit the following Web sites:

- DB2 Recovery Expert Product:
<http://www.ibm.com/software/data/db2imstools/index.html>
- DB2 Recovery Expert Publications:
<http://www.ibm.com/software/data/db2imstools/library.html>
- DB2 Recovery Expert Online Support:
<http://www.ibm.com/software/data/db2imstools/support.html>

If you are not able to find your answer searching the listed URLs, call 1-800-237-5511 to speak to an IBM representative for assistance.

Contacting customer support

There are a few things that it would be helpful to prepare if you need to contact customer support for any reason. They are as follows:

- Operating system name and version
- The name and version of your DB2 database
- The version of DB2 Recovery Expert that you are running.

Additionally, DB2 Recovery Expert performs logging during the recovery task processing. This log file can be very helpful to customer support when trying to resolve issues.

The DB2 configuration parameter `diaglevel` controls the logging for Recovery Expert. The valid settings for `diaglevel` are:

- 0 - No diagnostic data (this level is not recommended).
- 1 - Severe errors only.
- 2 - All errors (severe and not severe).
- 3 - All errors and warnings (Default)
- 4 - All errors, warnings, informational messages, and other internal diagnostic information.

On Windows[®], the log messages are written to the `db2diag.log` file that is located in the `diagpath`. The default location is the instance directory. The messages are also written to the Windows Event Log. On all other platforms, the log file is `db2recex.log` and is located in the `diagpath`. The default location is `$INSTHOME/db2dump`.

For more information about `diaglevel` and `diagpath`, see the DB2 Troubleshooting Guide.

Chapter 1. Introduction

This chapter contains basic information about DB2 Recovery Expert for Multiplatforms.

Product description

IBM DB2 Recovery Expert provides powerful diagnosis, reporting, and automated recovery capabilities to maintain accurate data and high availability. Given an object or system to be recovered, DB2 Recovery Expert provides a menu of feasible recommendations and alternatives. Once you select a recovery recommendation or alternative, Recovery Expert carries out the recovery automatically and quickly using the specifics of the problem detected.

Highlights

- Provides targeted, flexible, and automated recovery of database assets, even as systems remain online.
- Allows DBAs to recover database objects precisely and quickly without having to resort to full disaster recovery processes.
- Offers precision recovery options to support database development and maintenance.
- Has features that provide intelligent analysis of altered, incorrect, or missing database assets including tables, indexes, or data and automates the process of rebuilding those assets to a correct point in time, all without disruption to normal database or business operations.

Required Privileges, Authorities, and Authorization

Privileges enable users to create or access database resources. Authority levels provide a method of grouping privileges. Together, these act to control access to the database manager and its database objects. Users can access only those objects for which they have the appropriate authorization; that is, the required privilege or authority.

You must have SYSADM, SYSCTRL, or SYSMAINT authority to use the Recovery Expert tool.

Hardware requirements

DB2 Recovery Expert for Multiplatforms has the following hardware requirements:

- Memory and disk space sufficient to fulfill DB2 minimum requirements.
- Approximately 25 MB of free disk space on any supported platform.
- Additional space for user data (varies depending on the size of your databases and the amount of activity your system encounters).

Software requirements

To use IBM DB2 Recovery Expert for Multiplatforms, you must have IBM DB2 V7 or higher and the IBM JDBC driver. You also must be able to connect to a DB2 database.

IBM DB2 V7 or higher running locally:

This release of DB2 Recovery Expert for Multiplatforms requires you to run DB2 locally. Future releases will allow for remote connections to DB2.

IBM JDBC driver:

To run Recovery Expert, you must install the IBM JDBC driver. This driver is included with the DB2 Run-Time Client. Recovery Expert does not include this component.

Required DB2 settings

These are the required settings in the DB2 database configuration parameters that ensure that Recovery Expert functions properly. Refer to the *DB2 High Availability Guide* for more information on how to set these parameters.

Logs:

There are different types of DB2 logging, each providing a different level of recovery capability. Recovery Expert requires the use of archive logging so that the log files will be retained and can be analyzed to determine possible recovery options. Archive logging is enabled for a database by turning on either the LOGRETAIN or USEREXIT configuration parameters.

- Setting LOGRETAIN to YES, ON, or RECOVERY causes DB2 to save archive log files in the directory specified by the database configuration NEWLOGPATH/Path. As active logs fill, DB2 rolls filled logs into archive log files. These archive log files are used by Recovery Expert during recovery analysis or log analysis. These files should be saved until they become obsolete, which typically occurs after a number of full backups have been done.
- Setting USEREXIT to YES causes DB2 to perform logging regardless of the LOGRETAIN setting. This also indicates that a user exit program must be used to archive and retrieve log files. You must also ensure that the DB2 Recovery Expert minilogs are updated before the USEREXIT removes the logs from availability.

Performance:

You must ensure that the Application heap size (APPLHEAPSZ) database configuration parameter is set to 128 to provide sufficient memory page space when Recovery Expert is running.

Constraints

Availability of backup images:

DB2 Recovery Expert requires access to backup images to interpret log records, to gather data for the recovery of a single dropped table, and when using the object translation tool. In these cases, the backup must be available on disk. If you are using a storage manager to manage your backups, the storage manager must put the backup on disk prior to using Recovery Expert.

Referential integrity (RI):

Recovery Expert cannot determine that any change was made because of an RI relationship with another object. Thus, the log analysis tool within Recovery Expert

can only report on the log records found for selected objects. The same is true for recovery processes, working only with the designated object.

You can resolve RI issues by recovering the multiple objects involved in the RI relationship. For example, if you want to roll a table space forward to a point in time, and a table in that table space participates in a referential integrity relationship with another table that is contained in another table space, you should roll both table spaces forward simultaneously to the same point in time.

Chapter 2. Installation and configuration

This chapter contains information about installing DB2 Recovery Expert for Multiplatforms.

To install DB2 Recovery Expert for Multiplatforms

1. Log on to your system:

For Windows, perform the following steps:

- a. Log in as the Administrator user if your Windows system is configured for multiple user accounts.
- b. Locate the directory containing the setup files and run `db2rewin32setup`.

On Non-Windows systems, perform the following steps:

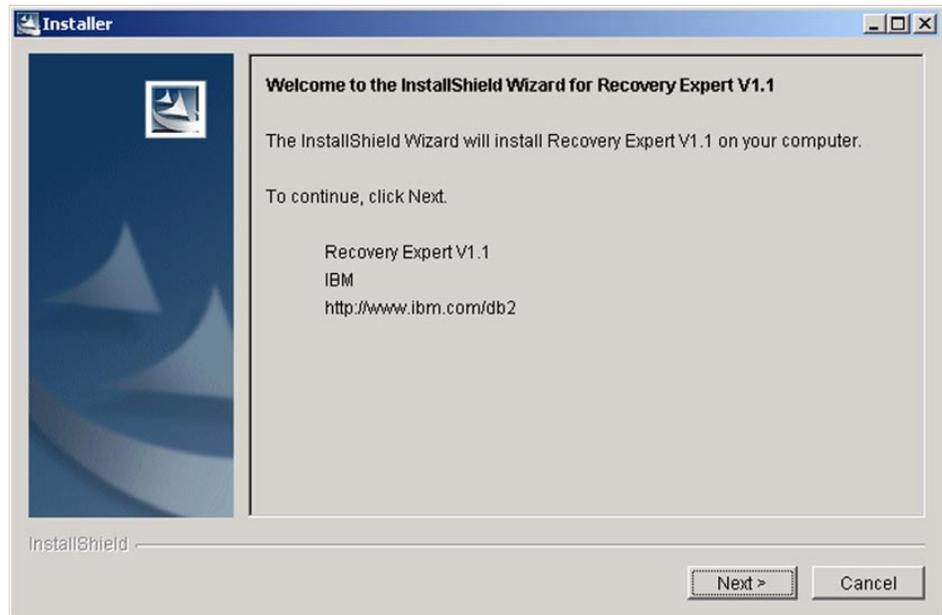
- a. Log in as the SuperUser ("root") user.
- b. Open a terminal window to activate a command-line shell.
- c. Change your current directory to the one containing the DB2 Recovery Expert installation package, as shown in the table below.
- d. Run the installation package from the shell. For example, `./db2reaixsetup`.

Note: If the program is not executable, you can make it executable with the command `chmod u+x ./db2re*setup`.

DB2 Recovery Expert install executable	Intended platform
<code>db2reaixsetup</code>	AIX®
<code>db2rehpuxsetup</code>	HPUX
<code>db2relinuxsetup</code>	Linux
<code>db2resolarissetup</code>	Solaris
<code>db2rewin32setup</code>	32-bit Windows

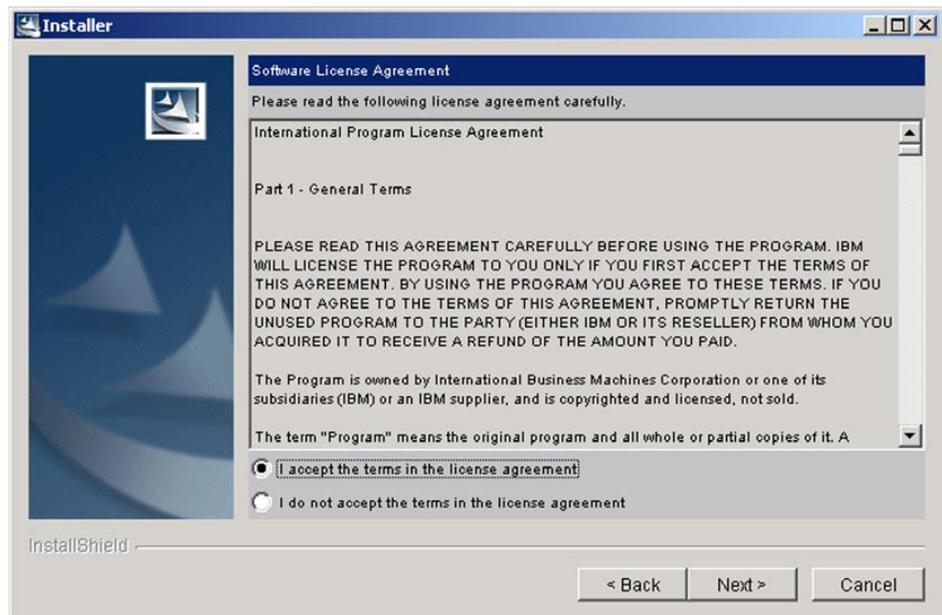
Note: The installation is the same for all environments, however, the Windows install is used as the example.

2. The DB2 Recovery Expert install starts.



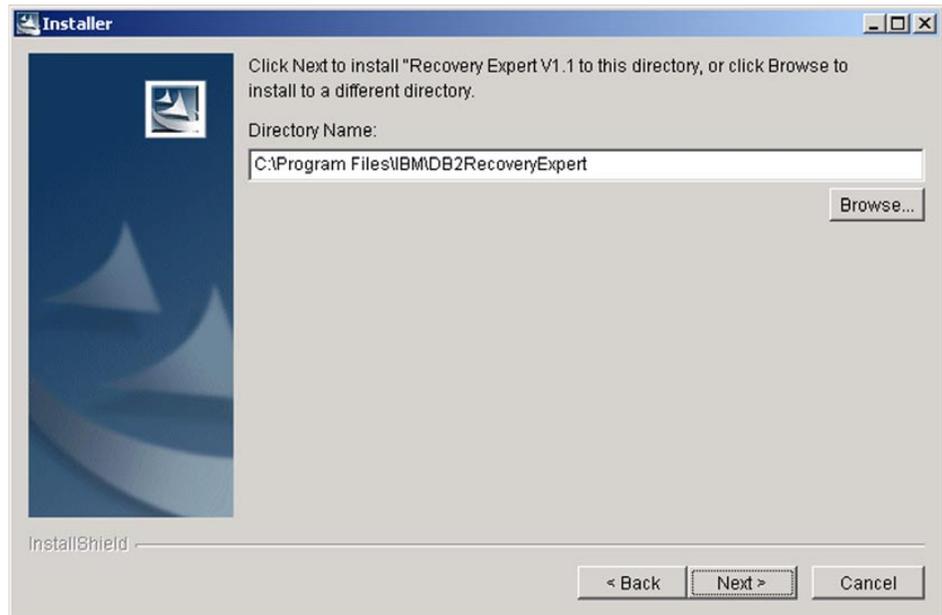
Click **Next** to continue.

3. In this step, you must read and accept the IBM Software License Agreement for DB2 Recovery Expert for Multiplatforms.



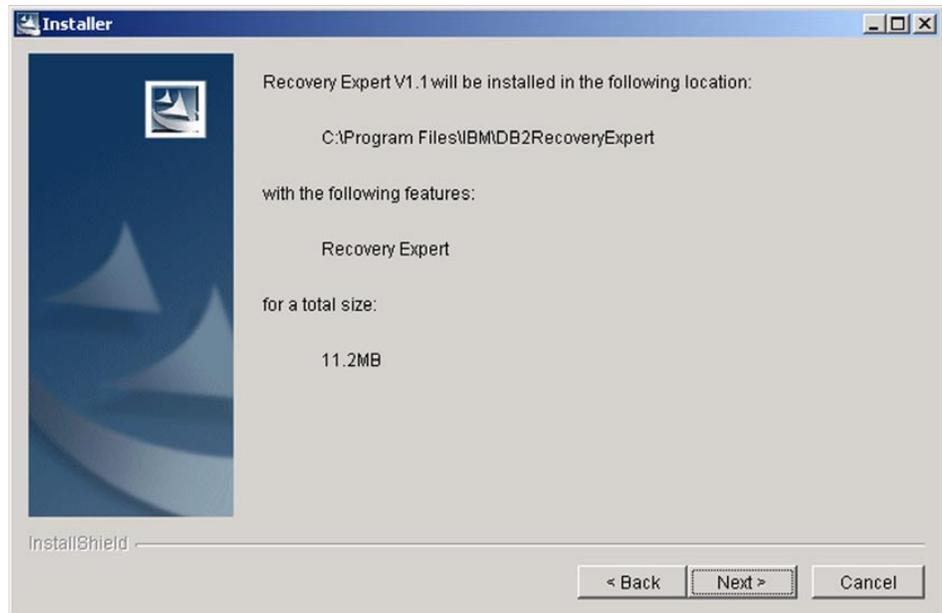
Click **Next** to continue.

- This step allows you to select a different install directory. You can type in a directory or use **Browse** to select a different location.

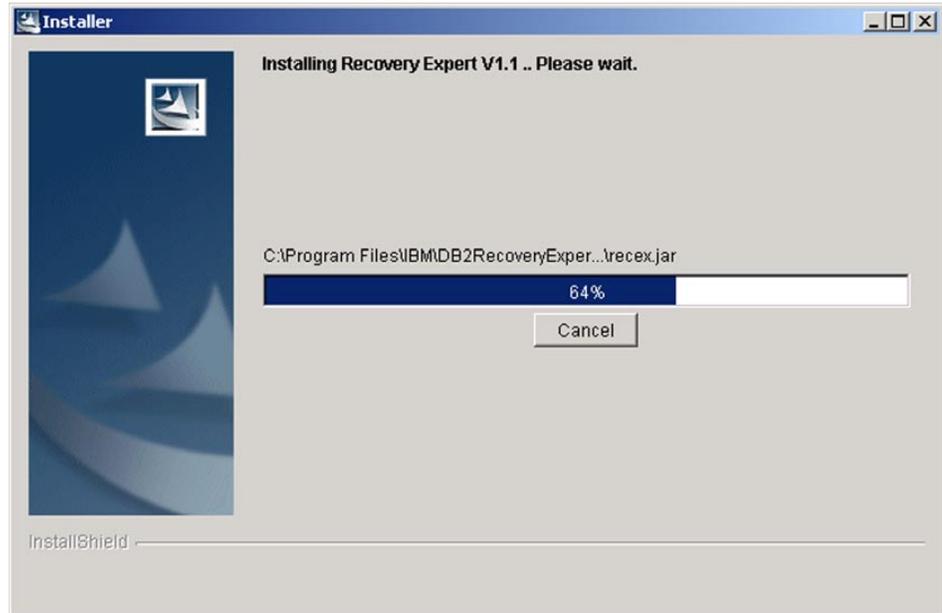


The default install directory for non-Windows systems is /opt/IBM/DB2RecoveryExpert. Click **Next** to continue.

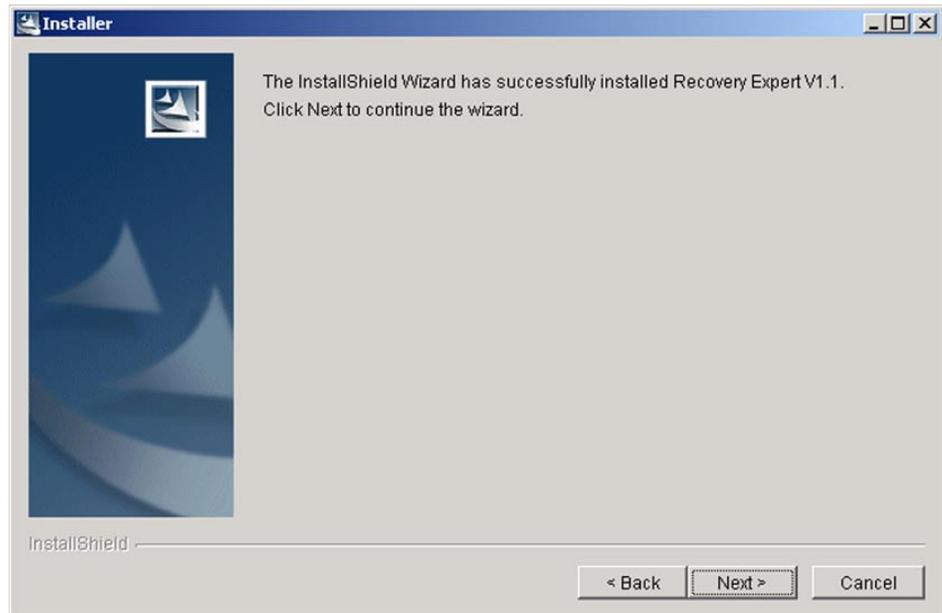
- This step provides you with summary information for the DB2 Recovery Expert install. If you want to make any changes, click **Back**, otherwise, click **Next** to continue.



6. This is the dialog that you will see while DB2 Recovery Expert is installed. To cancel at any time, click **Cancel**.

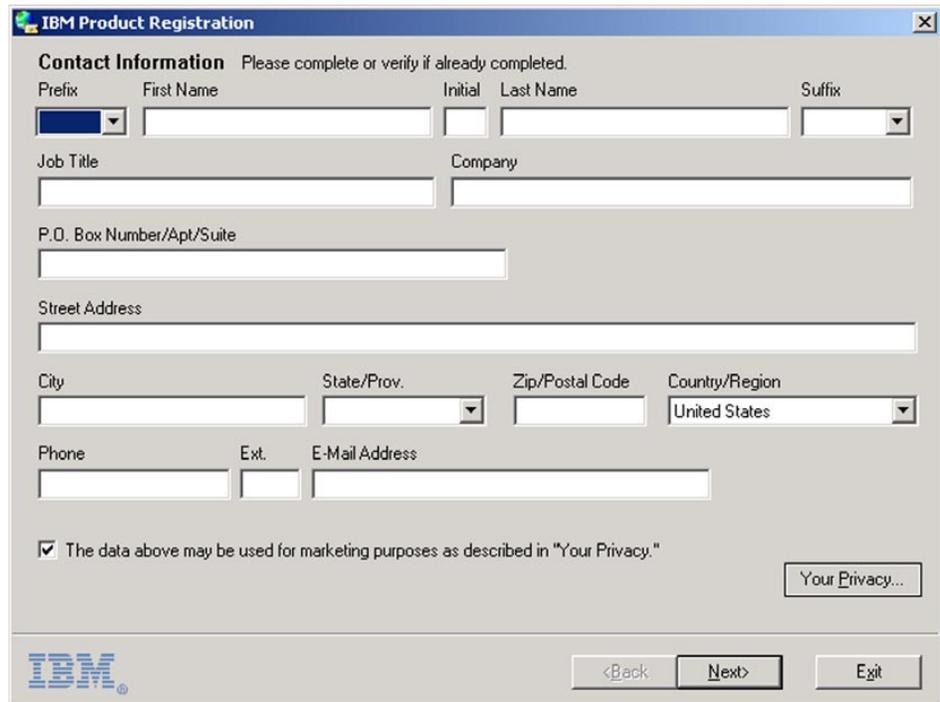


7. When DB2 Recovery Expert has been installed successfully, the following dialog appears.



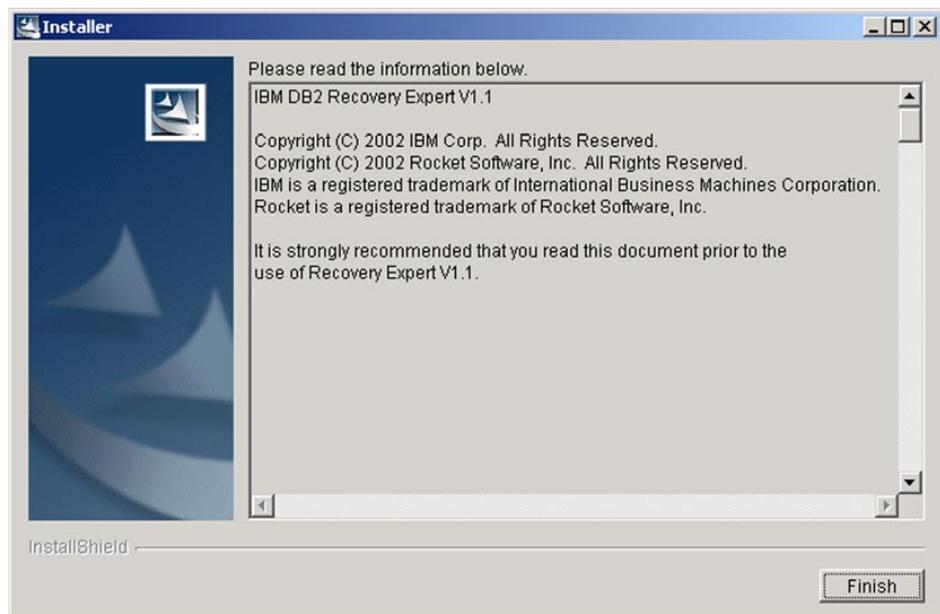
Click **Next** to proceed to the IBM Product Registration dialog.

8. Complete the registration so that IBM can notify you of product upgrades and future releases.



Click **Next** to continue.

9. Read the README information contained on the last dialog. When you have finished, click **Finish** to complete the installation of DB2 Recovery Expert.



10. When the installation is complete, run DB2 Recovery Expert by following the steps in Chapter 3, "Getting started with DB2 Recovery Expert" on page 13.

UNIX® installation information:

On UNIX systems, the installation process creates a program named db2recex which is used to start DB2 Recovery Expert and to run its various sub-functions

from the command line. This program grants users in the DB2 instance administration group permission to run DB2 Recovery Expert as the instance user so the program can access all available data recovery resources. To grant such permission, db2recex must have the “set-UID” permission set. This is done by default during the installation process.

Some security policies restrict the use of “set-UID” programs, or actively scan for such programs and remove that permission setting. The system administrator must ensure that the file system on which DB2 Recovery Expert is installed (/opt by default) allows “set-UID” programs to run. You can do this by checking the output of the mount(8) command for the presence of a “suited” flag (or absence of a “unsuited” flag) on the relevant file system. The administrator must also ensure that any automatic security policy enforcement programs are informed about db2recex’s “set-UID” permissions and do not remove or change them.

Configuration

There are three main directory paths in DB2 Recovery Expert:

1. **Installation Directory.** This is the directory where DB2 Recovery Expert was installed. The name for this path in the Configuration class is `INSTALL_DIR`. You cannot override this setting with a system property.
2. **Data Directory.** This is where the minilogs and other data generated by Recovery Expert reside. The directory contains a sub-directory for each instance, each of which contains the data for a single instance. The Configuration name for this path is `DATA_DIR`, it defaults to `{INSTALL_DIR}/{/}data`, and it can be overridden by the `recex.data` property.
3. **Configuration Directory.** This is where the `recex.properties` configuration file resides. System administrators can edit this file to make system-wide changes to the configuration. The name for this path is `CONF_DIR`, it defaults to `{INSTALL_DIR}/{/}conf`, and it can be overridden by the `recex.conf` property.

Configuration files

If you do not want to use the application defaults, then you can use instance overrides to configure DB2 Recovery Expert.

Application defaults

The application defaults ship as:

- `minilogs.path={DATA_DIR}/{/}minilogs`
- `sysstorage.path={DATA_DIR}/{/}sysstorage`
- `work.path={DATA_DIR}`

On UNIX systems, the installation process installs the `{INSTALL_DIR}/{/}conf{/}recex.properties` file as part of the archive unpacking. As a post-install procedure, it appends the following two lines to that file if the default `INSTALL_DIR` (beginning with /opt) was used and the installation was done as the superuser:

- `DATA_DIR={/}var{INSTALL_DIR}`
- `CONF_DIR={/}etc{INSTALL_DIR}`

This sets these directories to the appropriate locations in the UNIX file system, overriding any values for the same variables earlier in the file.

Instance overrides

You can make instance-specific changes by editing {CONF_DIR}/instance/recex.properties file. The CONF_DIR and DATA_DIR directories are set as described above, except at the instance overrides level. The DATA_DIR value changes to the DATA_DIR value from the application overrides or above appended with {/}{DB2INSTANCE}. There is also a DB2INSTANCE=<instance name> key/value pair added at this level. This makes the default instance DATA_DIR=/var/opt/IBM/DB2RecoveryExpert/DB2 for the DB2 instance.

Default Windows configuration values

For a Windows install for instance db2inst, the configuration values, at the instance level, are:

```
INSTALL_DIR=C:\Program Files\IBM\DB2RecoveryExpert
CONF_DIR=C:\Program Files\IBM\DB2RecoveryExpert
DATA_DIR=C:\Program Files\IBM\DB2RecoveryExpert
minilogs.path=C:\Program Files\IBM\DB2RecoveryExpert
sysstorage.path=C:\Program Files\IBM\DB2RecoveryExpert
work.path=C:\Program Files\IBM\DB2RecoveryExpert
```

Default non-Windows configuration values

For a non-Windows install for instance db2inst, the configuration values, at the instance level, are:

```
INSTALL_DIR=/opt/IBM/DB2RecoveryExpert
CONF_DIR=/etc/opt/IBM/DB2RecoveryExpert/db2inst
DATA_DIR=/var/opt/IBM/DB2RecoveryExpert/db2inst
minilogs.path=/var/opt/IBM/DB2RecoveryExpert/db2inst/minilogs
sysstorage.path=/var/opt/IBM/DB2RecoveryExpert/db2inst/sysstorage
work.path=/var/opt/IBM/DB2RecoveryExpert/db2inst
```

To configure DB2 Recovery Expert:

- To enable the command line tools from any directory, ensure that you add {INSTALL_DIR}/bin to your PATH.
- You can make system-wide changes by editing the application configuration file {INSTALL_DIR}/conf/recex.properties.
- You can make instance-specific changes by editing {CONF_DIR}/instance/recex.properties file.
The CONF_DIR and DATA_DIR directories are set as described above, except at the instance overrides level. The DATA_DIR value changes to the DATA_DIR value from the application overrides or above appended with {/}{DB2INSTANCE}. There is also a DB2INSTANCE=<instance name> key/value pair added at this level. This makes the default instance DATA_DIR=/var/opt/IBM/DB2RecoveryExpert/DB2 for the DB2 instance.

Editing recex.properties

The recex.properties file is a text file with key=value entries, one per line. Any line whose first character is a pound sign (#) is ignored. Pre-defined keys are:

Windows	Description	UNIX
\	File separator	/
;	Path separator	:

You can use key values for substitution by enclosing them in {}. The key values are:

- INSTALL_DIR=application installation directory
- CONF_DIR={INSTALL_DIR}/conf
- DATA_DIR={INSTALL_DIR}/data
- DB2INSTANCE=<default instance name>

Configuring for a DB2 UDB EEE environment

In a DB2 Universal Database™ Enterprise - Extended Edition (DB2 UDB EEE) environment, your database may be spread across several computers, which share a common directory area for each instance using the Network File System. DB2 Recovery Expert should be installed onto that shared filesystem so that each host participating in the partitioned database can use it. This is accomplished by specifying a path on the shared filesystem during the installation process.

DB2 Recovery Expert must also be configured so that its data directory is on the shared filesystem. This makes the information it collects about your database available from any host on which it is running. You configure the data directory by setting the DATA_DIR parameter in the configuration file. See “Configuration files” on page 10 for more information on configuration files. As an example, if your shared NFS filesystem is mounted on /net/db2, you could put the following line in the configuration file:

```
DATA_DIR={/}net{/}db2{/}DB2RecoveryExpert
```

This makes DB2 Recovery Expert maintain its data in a sub-directory on the shared filesystem. DB2 Recovery Expert creates the directory if it does not already exist.

When running DB2 Recovery Expert, either to perform Log Analysis or to recover data, you must run it on each partition (or node) in the database beginning with the catalog node. You must run it on the host that the partition's data and log files reside on. If you have more than one partition on a single host, you must connect to each partition on that host and perform the same operations on each partition. For example, to analyze the logs on a three-partition database, with nodes 0 and 1 on host A and node 2 on host B, where node 0 is the catalog node, you would do the following:

1. Run DB2 Recovery Expert on host A, connect to node 0 and perform the log analysis.
2. Connect to node 1 and perform the log analysis there.
3. Exit DB2 Recovery Expert on host A.
4. Run DB2 Recovery Expert again on host B, connecting to node 2.
5. Perform log analysis on node 2.

This method ensures that you have analyzed all the log files for your database. Starting with the catalog node is essential so that DB2 Recovery Expert can learn the structure of your database at the beginning of the process.

Chapter 3. Getting started with DB2 Recovery Expert

Once you have installed and configured IBM DB2 Recovery Expert, you can begin using DB2 Recovery Expert.

Before using DB2 Recovery Expert for Multiplatforms

DB2 Recovery Expert is a new product that offers advanced recovery facilities. Because recovering your data is of such critical importance, you must be experienced with the recovery process and with the recovery facilities included with DB2 Recovery Expert before deploying DB2 Recovery Expert to your production environment. Therefore, it is recommended that you thoroughly test DB2 Recovery Expert in a test environment before you deploy it to your production environment.

To test DB2 Recovery Expert:

1. Back up your data using the method of your choice. DB2 Recovery Expert depends on a quality backup of your data.
2. Analyze your logs using the log analysis tool.
3. Create and inspect Undo/Redo SQL.
4. Familiarize yourself with the versioning repository.

You must understand recovery facilities before using DB2 Recovery Expert in a production environment. Additionally, it is important that you stay current on DB2 Recovery Expert maintenance. Check the Data Management Tools Support page for the latest maintenance:

<http://www.ibm.com/software/data/db2imstools/support.html>

Launching DB2 Recovery Expert

After you launch DB2 Recovery Expert, you connect to the DB2 database you want to work with.

To launch DB2 Recovery Expert in Windows:

- Select **Start-->Programs-->DB2 Recovery Expert-->DB2 Recovery Expert**.

To launch DB2 Recovery Expert on a non-Windows platform:

- Using the GUI, you can click the DB2 Recovery Expert icon.
- Using the command line, type

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex
```

assuming that you installed Recovery Expert in the default location. If the installation's bin directory is in your PATH, you can type db2recex.

Database status

The Database status page is where you start the recovery process by selecting the database with which you want to work. Once you select a database, the status page provides you with information specific to the selected database.

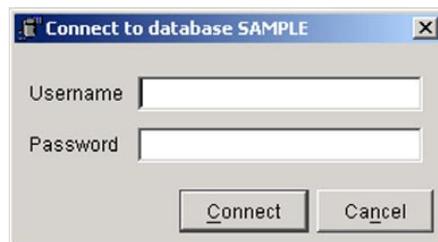
Database status fields

Field	Description
Database	Click the arrow to drop down a list of available databases.
Database state	Identifies whether the database is available.
Restore a dropped database	Select this box to restore a database that has been dropped. You must then type in the name of the database that you want to restore. A dropped database can only be restored from a backup image, thus, many of the options available for other object recoveries do not apply and are unavailable.
Log retain	Identifies the log retain setting in the database configuration. See the IBM DB2 <i>Data Recovery and High Availability Guide and Reference</i> for more information.
User exit	Identifies the user exit setting in the database configuration. See the IBM DB2 <i>Data Recovery and High Availability Guide and Reference</i> for more information.

Field	Description
Roll-forward recovery enabled	Identifies whether you can roll forward the data in the database. You can roll forward the data when logging has been enabled, either through the Log retain setting or the User exit setting in the database configuration.
First catalog log	The name of the first catalog log file created and used by log analysis.
Last catalog log	The name of the last catalog log file created and used by log analysis.
Last catalog update	Identifies the timestamp of the most recent catalog update. A catalog update occurs to the catalog log files that are created by the log analysis tool when it scans the logs for activity.
Log file directories or tapes	Lists the locations of the log file directories or tapes as recorded in the database history. If you have moved those log files using a mechanism outside of DB2, their original location is still provided here.

Selecting a database

You use the Database field drop-down menu to select the database that you want. When you select the database, you are prompted for a username and password that has the authority to connect to the selected database.



Type in a valid username and password that has the authority to perform recovery operations on the database and click **Connect**. The other fields on the Database Status page now contain the information for the selected database.

Note: In a EEE environment, there is a third field on the connection dialog for the node number. The node number must be the node number (0-999) to which you want to connect which resides on the machine that you are using.

To continue with restoring a DB2 object, see Chapter 4, “Recovering a database” on page 19.

Exiting DB2 Recovery Expert

Exiting from DB2 Recovery Expert closes the connection that you had with the database and closes Recovery Expert.

To exit DB2 Recovery Analyzer:

1. Click **File**. A drop-down menu appears.
2. Click **Exit**. Recovery Expert closes.

Other operations available in DB2 Recovery Expert

There are two other operations available in DB2 Recovery Expert from the **Tools** menu. The first feature allows you to perform log analysis. You can quickly search the DB2 logs to identify and recover data that is lost or corrupted. The other menu choice, Versioning Repository, allows you to create or update the Versioning Repository that DB2 Recovery Expert uses to recover DB2 objects.

Accessing log analysis

You can use the log analysis functionality in Recovery Expert for greater flexibility in recovering or restoring lost or corrupted data. You do not have to recover an entire database or table if only a portion of the data from a table must be recovered. You can simply use the log analysis functionality to generate REDO or UNDO SQL for the required data.

To access log analysis:

1. Click **Tools**. A drop-down menu appears.
2. Click **Log Analysis**. The log analysis facility opens.

For more information, refer to the documentation for log analysis in Chapter 7, “Log analysis” on page 39.

Creating or updating the versioning repository

The versioning repository is a database that is used to capture DB2 system catalog information. This information allows Recovery Expert to recreate object and authorization DDL. The versioning repository is not a shadow catalog. The information contained in the versioning repository consists of DB2 system catalog information and Recovery Expert object recovery information. Only DB2 system catalog information that is required to recreate object DDL is maintained in the versioning repository.

You must create a versioning repository for every database against which you want to run Recovery Expert. You must also periodically update the versioning repository as tables within the database change. DB2 Recovery Expert uses the versioning repository to recover DB2 databases or tables that cannot be recovered in any other way, for example, from a backup image.

Note: You can update the DB2 Recovery Expert versioning repository multiple times a day. It is recommended that you update the versioning repository at least once a day, especially after any ALTERs, DROPs, or CREATEs and before any other activities like INSERTs, UPDATEs, DELETEs, or backups are performed. Failure to keep the versioning repository up to date may result in invalid recovery paths.

Creating the versioning repository

To create the versioning repository:

1. Click **Tools**. A drop-down menu appears.
2. Select **Versioning Repository**. Another drop-down menu appears.
3. Click **Create**. This option is unavailable if the versioning repository already exists.

A message displays stating that the versioning repository has been created.

Note: Creating the Versioning Repository creates the ORVRTS table space and the many tables that comprise the versioning repository. This table space and tables are managed by DB2 Recovery Expert and the versioning repository update process. As such, this is an integral part of Recovery Expert and you should avoid performing any direct operations in this table space.

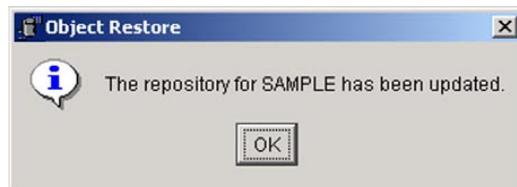
Updating the versioning repository

To update the versioning repository:

1. Click **Tools**. A drop-down menu appears.
2. Click **Versioning Repository**. Another drop-down menu appears.
3. Click **Update**. A confirmation dialog appears.



4. Click **Update**. Recovery Expert updates the versioning repository. A message appears to inform you when the process completes.



It is also possible to create and update the versioning repository through a command line interface, for example, if you wanted to use a scheduler to update the versioning repository at regular intervals. See Chapter 8, "Recovery Expert Versioning Repository command line interface" on page 49.

Chapter 4. Recovering a database

Once you have selected a database to use in recovery (as discussed in “Selecting a database” on page 15), you can make more specific recovery choices.

Database object selection

1. Select the database option on the Object Selection page.
2. Select the type of database recovery you want:
 - a. Recover the selected database to the original location.
 - b. Recover to another existing database. You must select the database name from the drop-down menu.
 - c. Recover to a new database. You must type in a new name and select the database node from the drop-down menu.
3. Click **Next** to continue.

DB2 Recovery Expert - Database SAMPLE

File Tools Help

Database Status | Object Selection | Point-in-time Selection | Recovery Path | Recovery Options | Summary

Select the object to restore.
Choose the type of object that you would like to restore, and then determine the exact object to recover.

Database

Recover to SAMPLE in node DB2

Recover to SATCTLDB in node DB2CTLSV

Recover to [] in node DB2

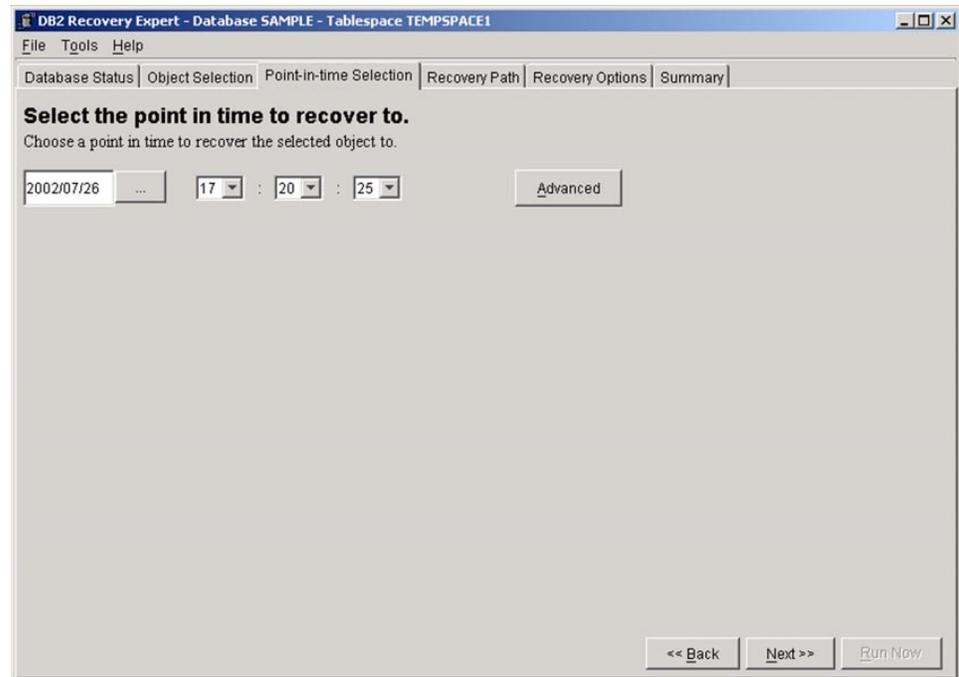
Tablespace []

Table []

<< Back Next >> Run Now

Database point-in-time selection

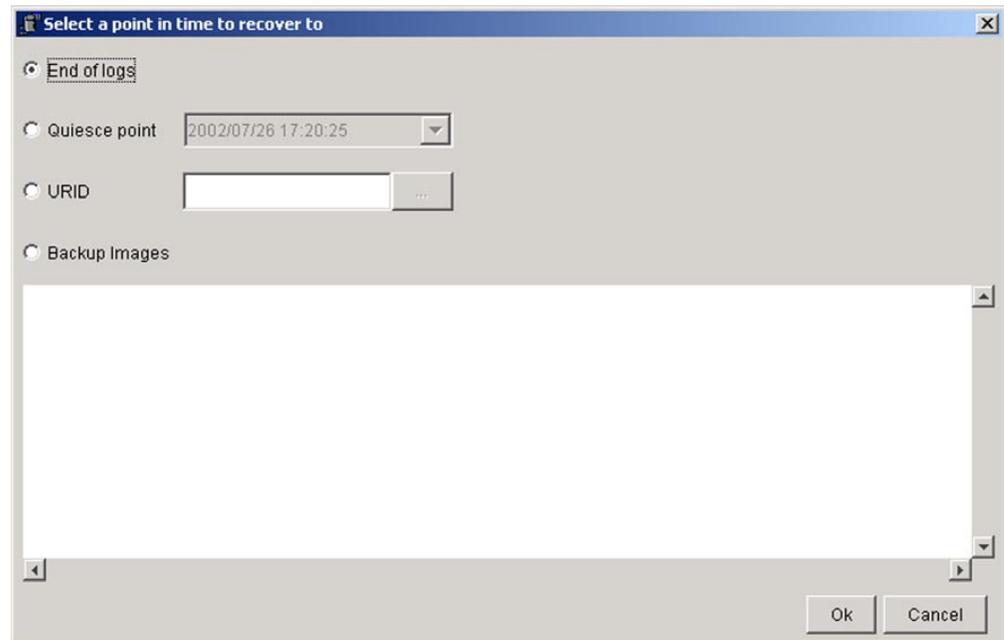
1. Use the drop-down lists to select the point in time to which you want to restore. If you want to select a point in time other than a date and time, such as the end of the logs, a quiesce point, a URID, or a specific backup image, see “To select other points in time:” on page 21.



2. When you have finished selecting the point to which you want to recover, Click **Next**.

To select other points in time:

When you click **Advanced** on the Point-in-time Selection page, another dialog opens offering you other possible points to restore to:



This dialog allows you to select:

- **End of logs** - Select this option to recover the object up to the end of the logs.
- **Quiesce point** - Select this option to recover the object up to the selected quiesce point. You select the quiesce point using the drop-down list.

Note: All quiesce points appear in the drop-down list. When recovering a table, however, some of the quiesce points may have been taken for tables other than the table you selected. The quiesce points for the table that you selected appear in blue. The quiesce points for other tables appear in black.

- **URID** - Select this option to recover the object up to a selected transaction. See “To select a URID:” on page 22.
- **Backup Images** - Select this option to recover a backup image of the object. The columns in the display area are as follows:

Table 1.

Column	Description
Object Type	Identifies the type of object.
Operation	Identifies the operation that created the image.
Operation Type	Identifies the type of backup.
Start timestamp	Identifies the starting timestamp on the backup.
End timestamp	Identifies the ending timestamp on the backup.
Table spaces	Identifies the number of table spaces on the backup.
Media Type	Identifies the media on which the backup resides.
Location	Identifies the location of the backup.

To select a URID:

1. Click the button to the right of the URID field. The URID Finder panel appears.

2. Use the date and time fields and drop-down menus to select the appropriate start and end dates and times.
3. Select the filters for those log items for which you want to search.
4. Click **Start**. After processing, the DB2 Log Analysis Transactions panel appears.

URID	Date	Time	Authid	Nodegroup	Tablespace	Table O...	Table Name	Updates	Insert:
0000000002...	2002/07/...	16:32:34		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	32	
000000000103	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000332	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000333	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000334	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000335	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000336	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000337	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000338	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000339	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
00000000033A	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
00000000033B	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
0000000003...	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
0000000003...	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
00000000033E	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
00000000033F	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000340	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	
000000000341	2002/07/...	16:48:40		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOYEE	1	

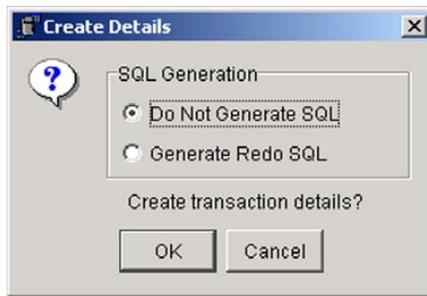
5. Select the URID to which you want to recover.
6. Click **Select URID**.

To get detailed transaction information:

Continuing from the previous section,

1. Double-click the row for which you want to generate a detail report. In this way, you will be able to pinpoint a specific transaction to which you want to recover.

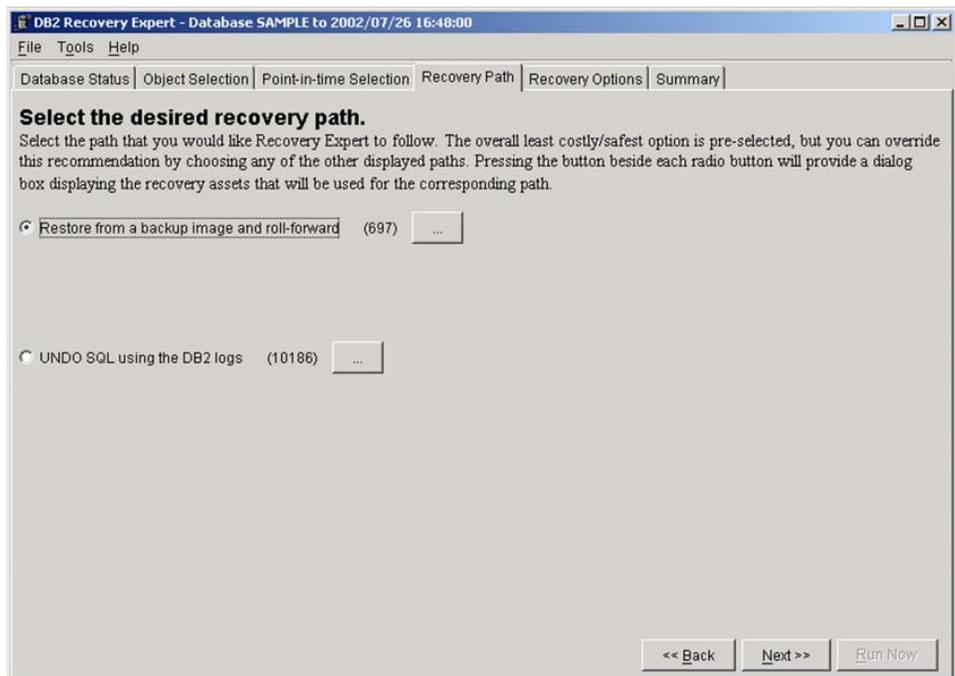
A confirmation dialog opens.



2. Indicate whether you want to generate Redo SQL or no SQL with the detail report.
3. Click **OK**. The log analysis facility generates a report detailing the transaction details as well as the SQL if you selected to do so. You can review the specific transaction details to locate the URID to which you want to recover.

Database recovery path

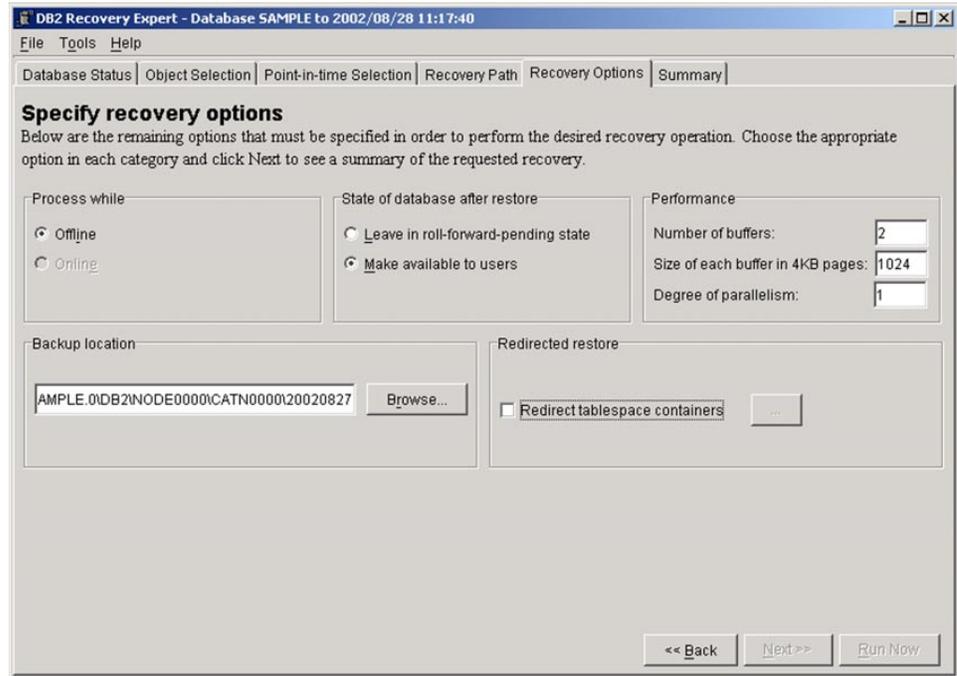
1. Select the desired recovery path.
 - The number in parentheses corresponds to the number of operations required for recovery. This is Recovery Expert's best estimate of the recovery cost. Generally speaking, the lower a number is, the lower the recovery cost for that option.
 - The recommended recovery path is preselected, but you can override the selection by selecting another option.
 - Clicking the button to the right of the recovery options displays the recovery assets that will be used in the recovery process.



2. Click **Next** to continue.

Database recovery options

1. Specify any additional recovery options that you want.



Note: All options may not be visible on the Recovery Options page. These options are dependent upon your prior selections.

The remaining options are:

Table 2.

Option	Description
Process while	Indicate how you want to process the database recovery, either offline or online.
State of database after restore	Indicate the state that you want the database left in after the restore completes. If you leave the database in roll-forward-pending state, the database will not be available to users.
Number of buffers	Specify the number of buffers to use when restoring the database.
Size of each buffer in 4KB pages	Specify the size of the buffers to use when restoring the database.
Degree of parallelism	Specify the number of concurrently executing processes for the recovery. While parallel processes can perform the recover faster than a single process, each degree of parallelism adds to the system memory and CPU overhead.
Commit Scope	Specify how often a commit should be performed during the recovery process. This option is available only when UNDO or REDO SQL is used in the recovery.

Table 2. (continued)

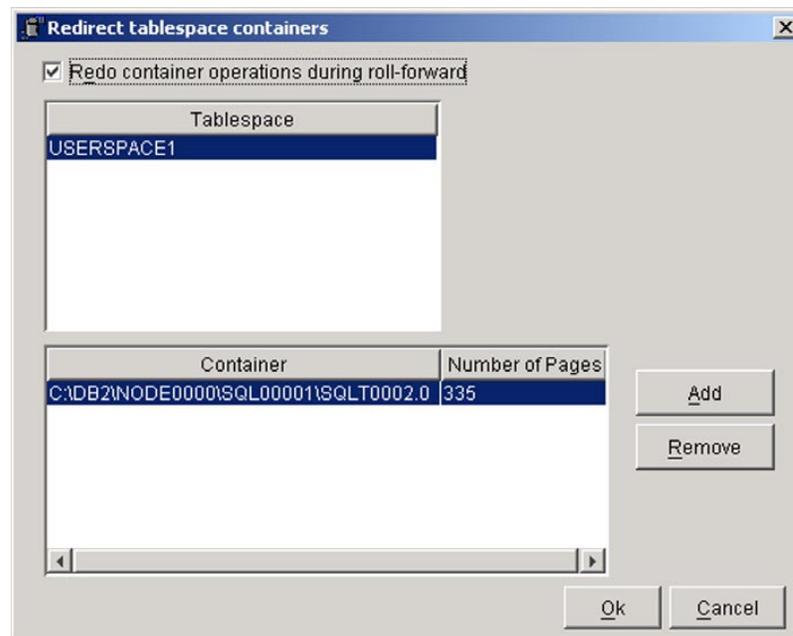
Option	Description
Backup location	If you have moved the database backup image files, you must specify an alternate directory for the backup image and identify the directory path by using Add to browse to the directory and Remove to delete any unwanted paths.
Backup location and file name	This option is available only when you select Restore a dropped database on the Database Status page. See "Database status" on page 14. You must specify the backup location and file name for the backup image of the dropped database by using Add to browse to the directory and file and Remove to delete any unwanted paths. files.
Redirected restore	This option applies to a table space recovery from a backup image. Selecting it and clicking the button to the right of the option allows you to select other containers for the table space. See "Redirected restore" for more information.

2. Click **Next** to continue.

Redirected restore

During a backup of a database, a record is kept of all the table space containers in use by the table spaces that are backed up. During a recovery, all containers listed in the backup are checked to see if they currently exist and are accessible. Redirected restore allows you to add, change, or remove table space containers.

Selecting **Redirected restore** and clicking the button to the right of the option displays the Redirect tablespace containers dialog that allows you to select other containers for the table space.



The current containers are set as the defaults. This means that if you select **Redirect table space containers** and change nothing on the dialog for the table space, that table space will have its containers redirected to the current ones. Therefore, you must ensure that when you choose to redirect, you must be certain that all table spaces have the appropriate containers. If you want to redirect for some table spaces but not for others, you should remove all containers for the table spaces that you do not want to redirect. An empty container list results in the table spaces being recovered to their original containers.

Use **Add** to browse to the redirect location. Depending on the type of database you are recovering, you will encounter different dialogs:

- If you are recovering an SMS database, when you click **Add** a path browser appears so that you can browse to the correct directory.
- If you are recovering a DMS database, when you click **Add** you then have the option to select either **File** or **Raw device**:
 - When you select **File** and click **Add**, you will then use a path browser to select a location and file name.
 - When you select **Raw device** and click **Add**, you will then type in a partition, drive or machine designation (for example, C:).

Once you select the redirect location, you must type in the number of pages that you want in the container. If you have selected a redirect location, but do not specify a number of pages, you will not be able to continue.

Database summary

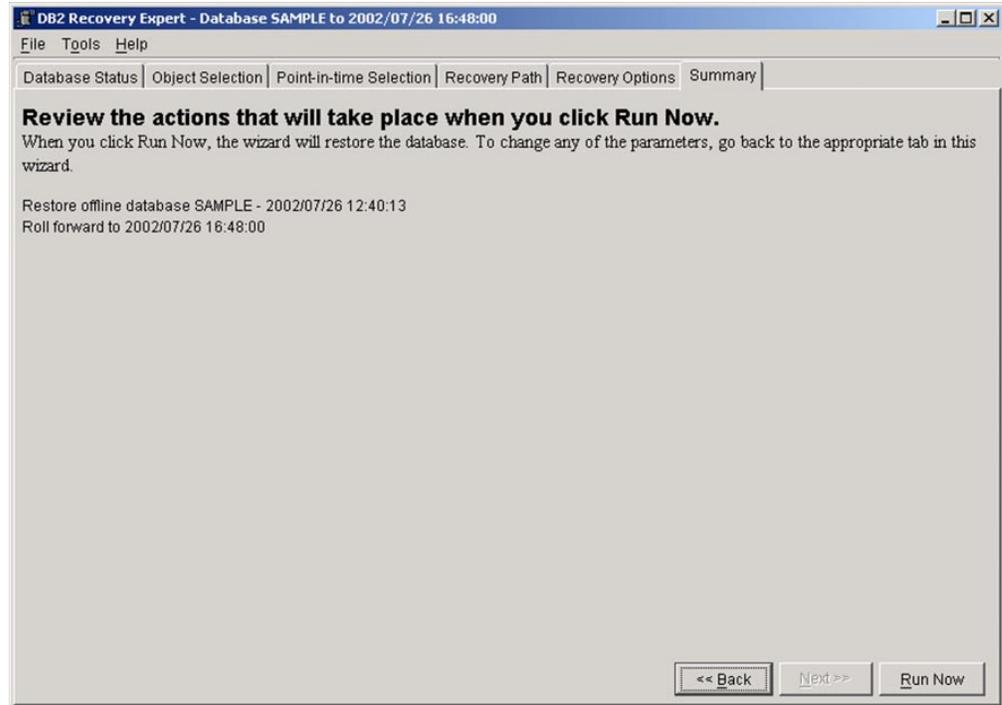
The Summary page provides you with a final review of all the recovery options that you have selected.

If the recovery options are correct:

- Click **Run Now** to restore the database.

If you want to make any changes to the recovery options:

- Use **Back** or the tabs across the top to select the page containing the options that you want to change.

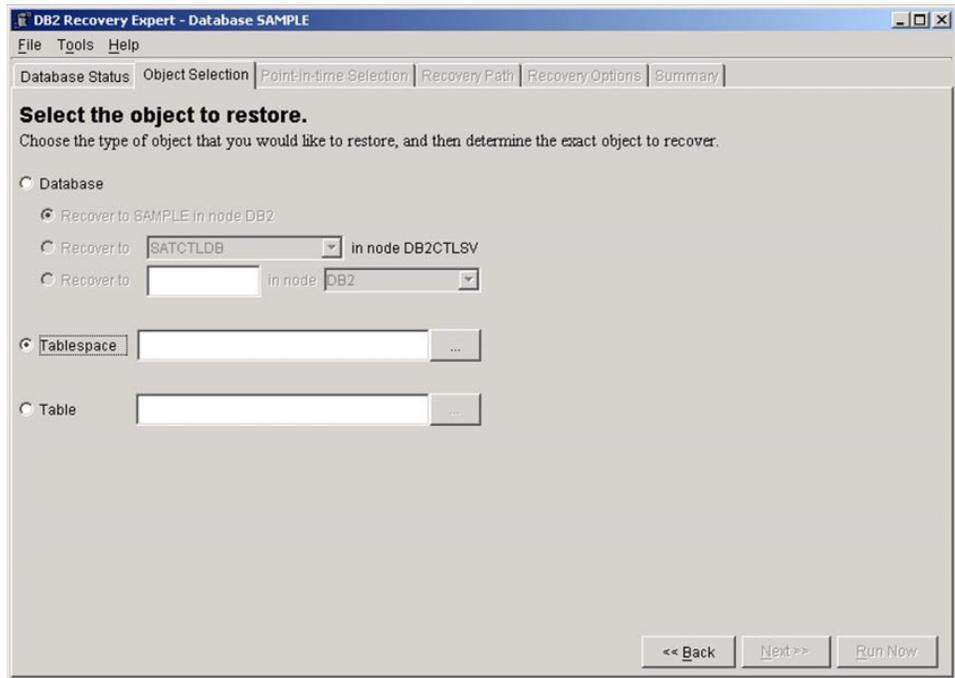


Chapter 5. Recovering a table space

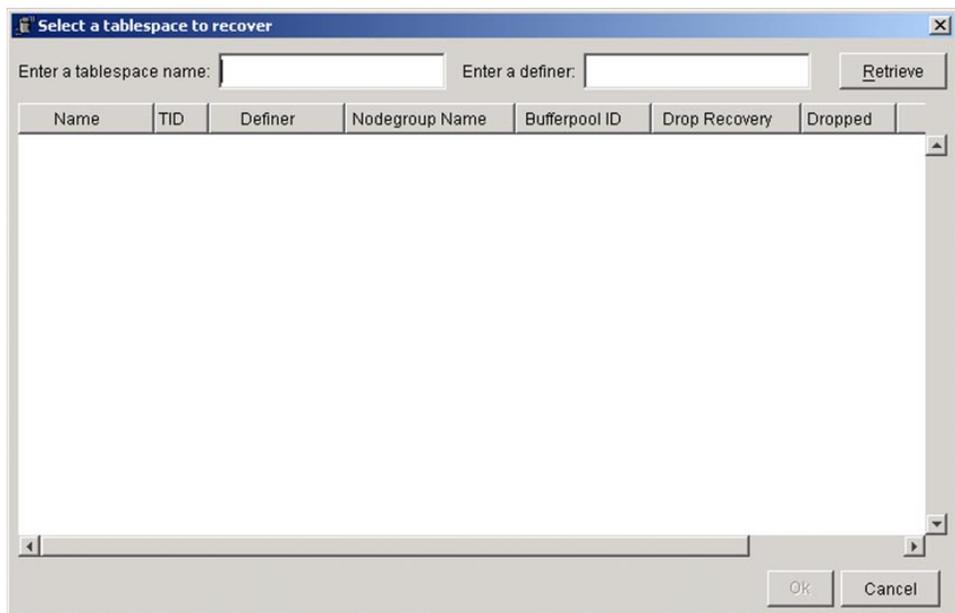
Once you have selected a database to use in recovery (as discussed in “Selecting a database” on page 15), you can make more specific recovery choices.

Table space object selection

1. Select the Table spaces option on the Object Selection page.



2. Type in the table space name or click the button to the right of the table space field to browse for a table space. The table space selection dialog appears.



3. Type the table space name, the definer's name, or a portion of either in the filter fields. You can use the star (*) or the percent sign (%) as a wildcard in the filter fields.
4. Click **Retrieve**. The display area is populated with the table spaces that match your selection criteria. When you make any changes to the filter fields, you must click **Retrieve** to refresh the list using the changed filter criteria.

Note: Dropped objects appear in red.

5. Select the table space that you want to recover.
6. Click **Ok**.
7. Click **Next** to continue.

Table space selection display columns

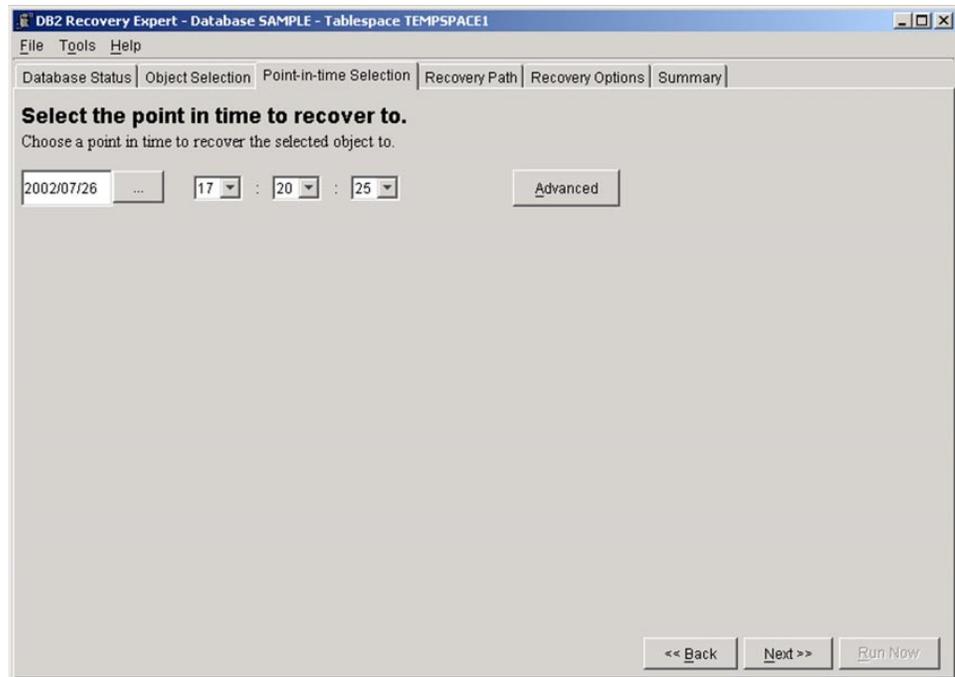
Table 3.

Column name	Description
Name	Table space name.
TID	Table space ID.
Definer	The ID of the person or entity that created the table space.
Nodegroup Name	The name of group of database partitions to which the table space belongs.
Bufferpool ID	The ID of the bufferpool to which the table space belongs.
Drop Recovery	Identifies whether dropped tables in the specified table space may be recovered using the RECOVER DROPPED TABLE ON option of the ROLLFORWARD command.
Dropped	Indicates whether the table space is currently dropped.
Object Create Time	The timestamp when the table space was created.

Note: The Version End Time is often the same timestamp as the Version Create Time of the next version.

Table space point-in-time selection

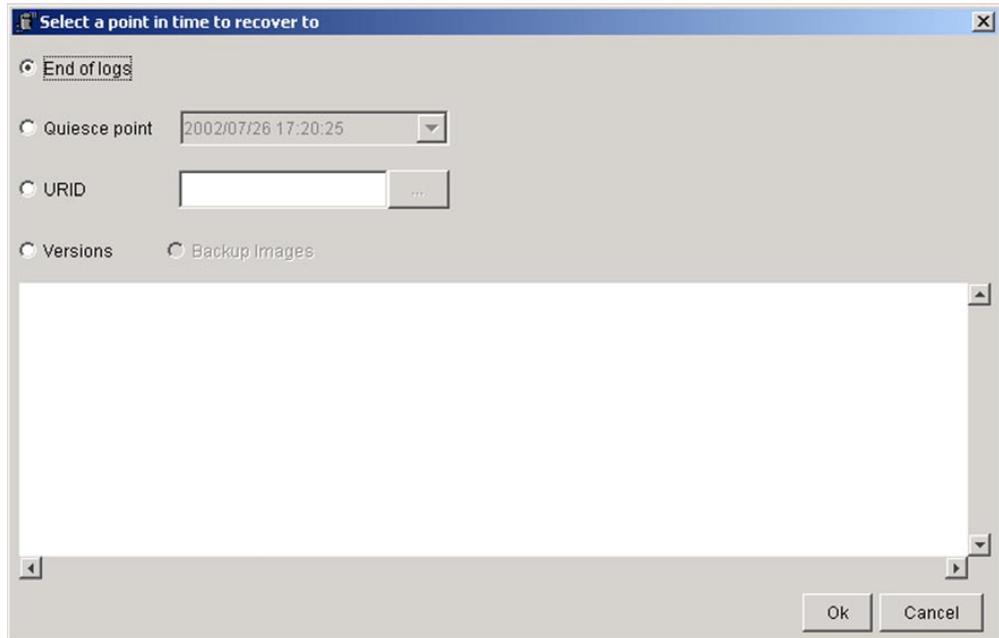
1. Use the drop-down lists to select the point in time to which you want to restore. If you want to select a point in time other than a date and time, such as the end of the logs, a quiesce point, a URID, a different version of the table space or a specific backup image, see “To select other points in time:” on page 21.



2. When you have finished selecting the point to which you want to recover, Click **Next**.

To select other points in time:

When you click the button to the far right on the Point-in-time Selection page, another dialog opens offering you other possible points to restore to:



This dialog is the same as the dialog discussed in “To select other points in time:” on page 21, with the following additional option:

- **Versions** - Select this option to recover a specific version of the object. The columns in the display area are as follows:

Column	Description
Owner	Identifies the owner of the object.
Name	Identifies the name of the object.
Version Create Time	Identifies the time that the object was created.
Version End Time	Identifies the time that the object ended.

Table space recovery path

Refer to “Database recovery path” on page 23 for information on this page, substituting table space for database.

Table space recovery options

Refer to “Database recovery options” on page 24 for information on this page, substituting table space for database.

Table space summary

Refer to “Database summary” on page 26 for information on this page, substituting table space for database.

DMS considerations

DB2 Recovery Expert provides the ability to recover DMS table spaces as well as tables that existed in a DMS table space.

Recovering DMS table spaces:

When you recover a DMS table space, Recovery Expert restores the table space as an SMS table space, yet it has the *same* name and container locations that it did when it existed.

Recovering tables that existed in a DMS table space:

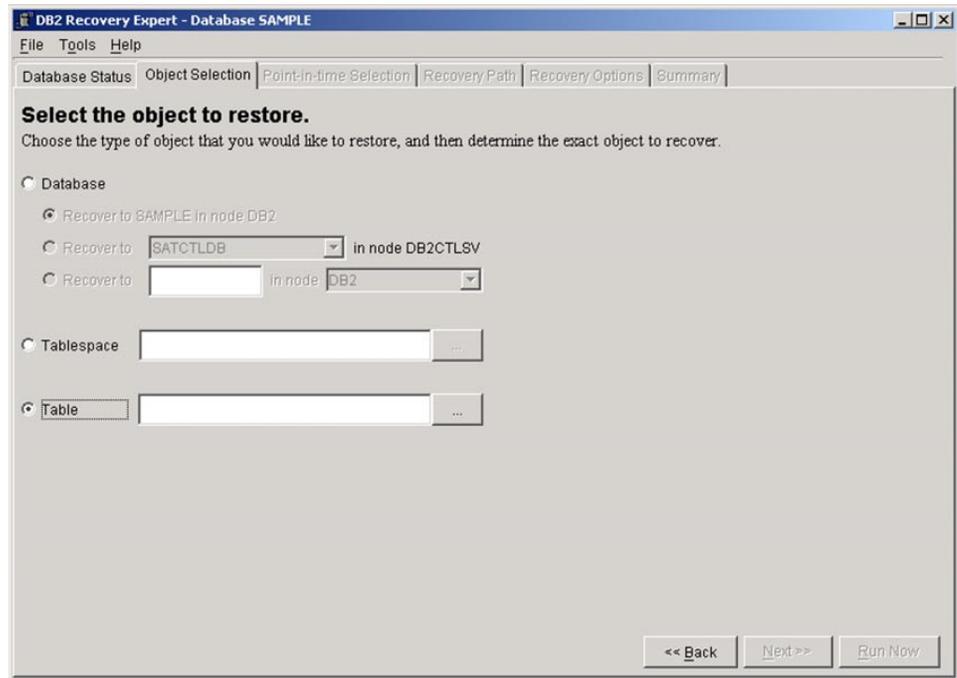
When you recover tables that existed in a DMS table space, Recovery Expert restores the table space as an SMS table space with a *similar* name and container locations of the DMS table space when the table did exist. The table space and container names contain most of their original names, but has “_SMS#” appended to it, where # is an integer.

Chapter 6. Recovering a table

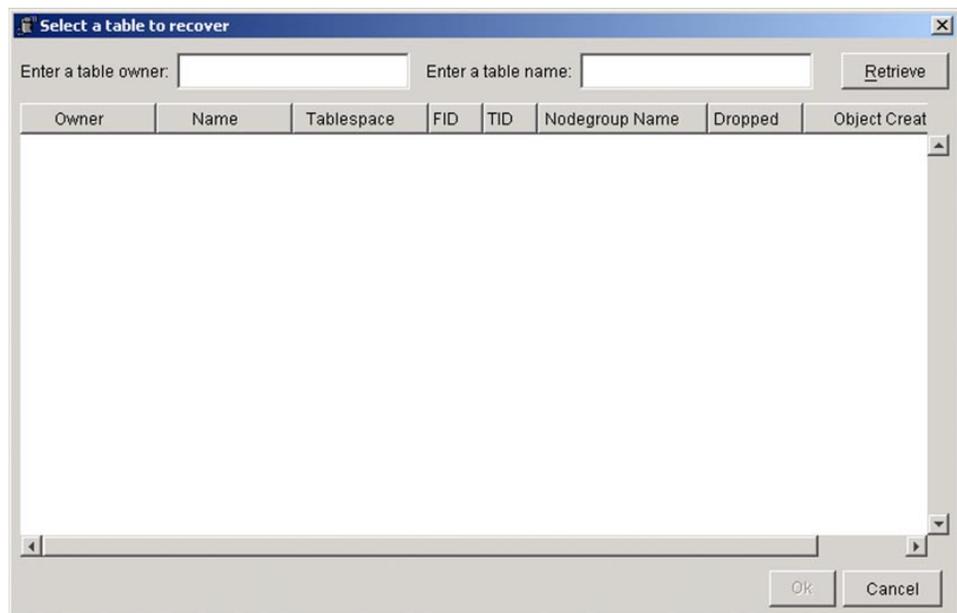
Once you have selected a database to use in recovery (as discussed in “Selecting a database” on page 15), you can make more specific recovery choices.

Table object selection

1. Select the Table option on the Object Selection page.



2. Type in the table name or click the button to the right of the table field to browse for a table. The table selection dialog appears.



3. Type the table owner, the table name, or a portion of either in the filter fields. You can use the star (*) or the percent sign (%) as a wildcard in the filter fields.
4. Click **Retrieve**. The display area is populated with the tables that match your selection criteria. When you make any changes to the filter fields, you must click **Retrieve** to refresh the list using the changed filter criteria.

Note: Dropped objects appear in red.

5. Select the table that you want to recover.
6. Click **Ok**.
7. Click **Next** to continue.

Table selection display columns

Table 4.

Column name	Description
Owner	The table owner.
Name	Table name.
Table space	Table space name to which the table belongs.
FID	Table ID.
TID	Table space ID.
Nodegroup Name	The name of group of database partitions to which the table belongs.
Dropped	Indicates whether the table is currently dropped.
Object Create Time	The timestamp when the table was created.
Version Create Time	The timestamp when this version of the table was created using the repository update.
Version End Time	The timestamp when this version of the table ceased to exist as determined by the repository update.

Note: The Version End Time is often the same timestamp as the Version Create Time of the next version.

Table point-in-time selection

Refer to “Table space point-in-time selection” on page 31 for information on this page, substituting table for table space.

Table recovery path

Refer to “Database recovery path” on page 23 for information on this page, substituting table for database.

Table recovery options

Refer to “Database recovery options” on page 24 for information on this page, substituting table for database.

Table summary

Refer to “Database summary” on page 26 for information on this page, substituting table for database.

Chapter 7. Log analysis

Log analysis in Recovery Expert provides you with greater flexibility in recovering or restoring lost or corrupted data. You can use the DB2 logs to recover a portion of the data from an entire table. You use log analysis to search the DB2 logs, find the lost or corrupted data before it became lost or corrupted, and generate REDO or UNDO SQL to recover that data.

Getting started with log analysis

To access the log analysis facility:

1. Click **Tools**. A drop-down menu appears.
2. Click **Log Analysis**. The log analysis facility opens.

The screenshot shows the 'DB2 Log Analysis Tool' dialog box. It is titled 'DB2 Log Analysis Tool' and has a standard Windows-style window with a close button. The dialog is divided into several sections. On the left, there is a 'Database' field with the text 'SAMPLE'. Below that is a 'Report options' section with two radio buttons: 'Specify report options' (which is selected) and 'Run complete report'. Underneath is a 'Log date/time range' section with four text boxes: 'Start Date' (2002/08/28), 'Start Time' (00:00:00), 'End Date' (2002/08/28), and 'End Time' (22:55:50). On the right side, there is a 'Filters for log data' section with five checkboxes: 'Show UPDATES', 'Show DELETES', 'Show INSERTS', 'Catalog Data', and 'Scan Archive Logs' (which is checked). Below the filters are four text boxes for 'Tablespace Name', 'Table Owner', 'Table Name', and 'User ID'. At the bottom of the dialog are four buttons: 'Start', 'Stop', 'Help', and 'Exit'.

Complete the fields on the main page as follows:

Table 5.

Field	Description
Database	Type the name of the database on which you want to run log analysis.

Table 5. (continued)

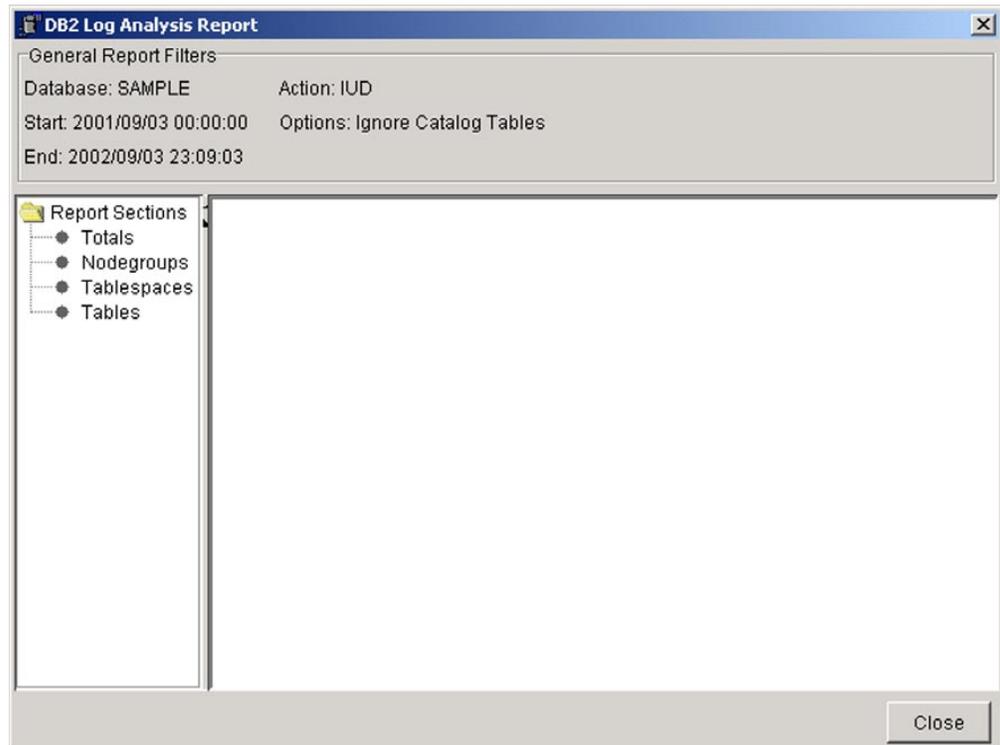
Field	Description
Report options	<p>Indicate what type of report you want to generate. The options are:</p> <ul style="list-style-type: none"> • Specify report options—indicates that you want to specify a time range and filters for the log analysis to reduce the volume of data in the report. • Run complete report—indicates that you want a complete report on all of the UPDATES, DELETES, and INSERTs that occurred in all tables in the selected database, using all available log files. If you select this option, all other parameters on the page become unavailable.
Start Date	Specify a start date for log analysis.
Start Time	Specify a start time for log analysis.
End Date	Specify an end date for log analysis.
End Time	Specify an end time for log analysis.
Filters for log data (can be used together in any combination)	
Show UPDATES	Indicates that you want to include UPDATE statements in the log analysis report.
Show DELETES	Indicates that you want to include DELETE statements in the log analysis report.
Show INSERTs	Indicates that you want to include INSERT statements in the log analysis report.
Catalog Data	Indicates that you want to include DB2 catalog data in the log analysis report.
Scan Archive Logs	<p>Indicates that you want log analysis to include archived logs in its analysis. The log analysis tool will read all previously unread archive logs, update the archive log meta data for these new archive logs, and then use that new information in determining recovery paths.</p> <p>If this is option not selected, all data about the tables comes from the current system catalog only, and create/drop/alter events for tables are ignored. You should remove this default selection only if you want reports on most recent log activity or for tables that are never dropped or altered.</p>
Tablespace Name	Type the name of the table space for which you want a log analysis report. This is the only table space for which a report will be generated. You can also type a few letters and use an asterisk (*) as a wildcard character.
Table Owner	Type the name of the table owner for which you want a log analysis report. This is the only table owner for which a report will be generated. You can also type a few letters and use an asterisk (*) as a wildcard character.

Table 5. (continued)

Field	Description
Table Name	Type the name of the table for which you want a log analysis report. This is the only table for which a report will be generated. You can also type a few letters and use an asterisk (*) as a wildcard character.
User ID	Type the user ID for which you want a log analysis report. This is the only user ID for which a report will be generated. You can also type a few letters and use an asterisk (*) as a wildcard character. Note: To filter on User ID, you must have had CHANGE DATA CAPTURE enabled in the DB2 tableconfiguration when the changes occurred. This captures the AUTHID in the log, allowing log analysis to filter on that column.

Generating the report

Once you have selected and entered all the criteria for your log analysis, click **Start** to generate the report. You can click **Stop** at any time while the log analysis facility is running to stop generating the report. When the log analysis facility finishes processing, a summary level report appears.



Across the top of the report is the basic information for the contents of the report. This is essentially the information that selected on the main dialog for log analysis. For more information on these options, see “Getting started with log analysis” on page 39. The summary report provides summary information on the number of UPDATES, INSERTs, and DELETES. The data is sorted in different sections of the report, and you can review this data as:

- **Totals**--the total number of UPDATES, INSERTs, and DELETES found.

- **Nodegroups**--the total number of UPDATES, INSERTS, and DELETES found in each nodegroup.
- **Tablespaces**--the total number of UPDATES, INSERTS, and DELETES found in each table space.
- **Tables**--the total number of UPDATES, INSERTS, and DELETES found in each table.

By clicking on the word in the left frame, the information in the right frame changes according to your selection. For example, clicking **Totals** provides you with this view of the data:

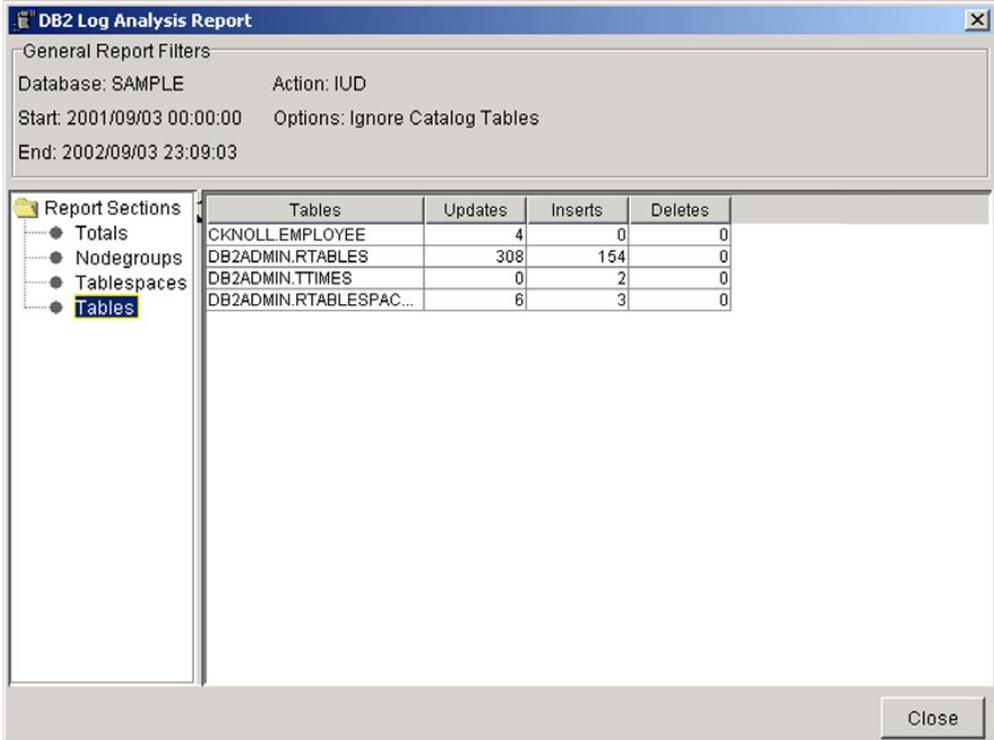
The screenshot shows a window titled "DB2 Log Analysis Report" with a "General Report Filters" section. The filters are: Database: SAMPLE, Action: IUD, Start: 2001/09/03 00:00:00, Options: Ignore Catalog Tables, and End: 2002/09/03 23:09:03. On the left, under "Report Sections", "Totals" is selected. The main area displays a table with the following data:

Action	Total
Updates	318
Inserts	159
Deletes	0

A "Close" button is located in the bottom right corner of the window.

Log analysis in the report

Once you generate the summary report, you can continue the analysis of the report by double-clicking a row in the right frame. For example, selecting the Tables section of the report provides you with this view of the data:

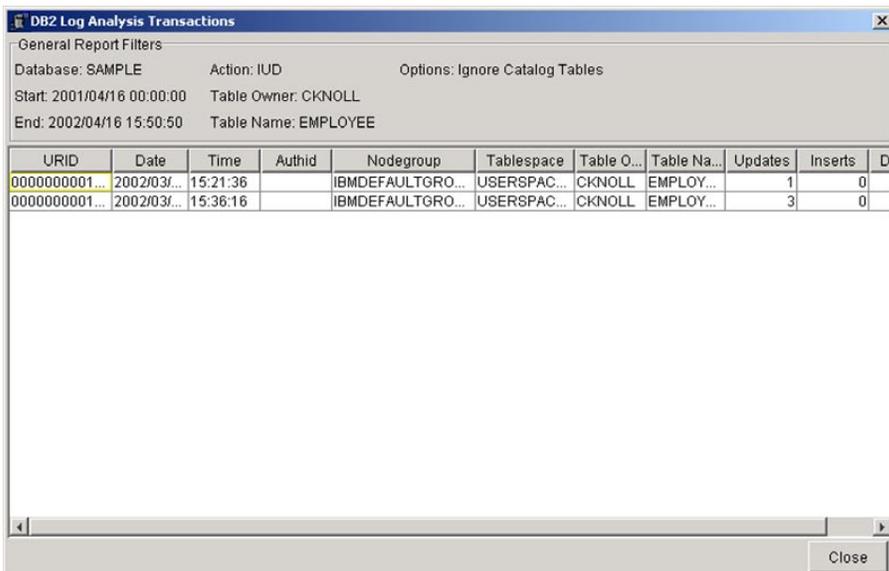


The screenshot shows a window titled "DB2 Log Analysis Report". It has a "General Report Filters" section with the following information: Database: SAMPLE, Action: IUD, Start: 2001/09/03 00:00:00, End: 2002/09/03 23:09:03, and Options: Ignore Catalog Tables. Below the filters is a "Report Sections" tree with "Tables" selected. The main area displays a table with the following data:

Tables	Updates	Inserts	Deletes
CKNOLL.EMPLOYEE	4	0	0
DB2ADMIN.RTABLES	308	154	0
DB2ADMIN.TTIMES	0	2	0
DB2ADMIN.RTABLESPAC...	6	3	0

A "Close" button is located at the bottom right of the window.

The log analysis facility sorts the data according to table, and if you know which table has the data you are looking for, you can double-click that table name in the right frame. By doing so, you move to the transaction-level view of the general report:



The screenshot shows a window titled "DB2 Log Analysis Transactions". It has a "General Report Filters" section with the following information: Database: SAMPLE, Action: IUD, Options: Ignore Catalog Tables, Start: 2001/04/16 00:00:00, End: 2002/04/16 15:50:50, Table Owner: CKNOLL, and Table Name: EMPLOYEE. Below the filters is a table with the following data:

URID	Date	Time	Authid	Nodegroup	Tablespace	Table O...	Table Na...	Updates	Inserts	D
0000000001...	2002/03/...	15:21:36		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOY...	1	0	
0000000001...	2002/03/...	15:36:16		IBMDEFAULTGRO...	USERSPAC...	CKNOLL	EMPLOY...	3	0	

A "Close" button is located at the bottom right of the window.

In this view you can start to gain a more detailed view of the individual transactions that have occurred on the specified table, yet you still do not know what the actual data change was. The columns on this report view are detailed below:

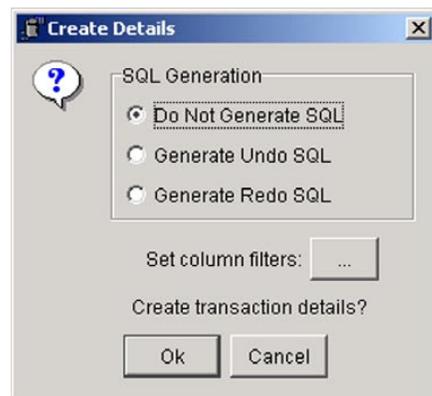
Table 6.

Column	Description
URID	The unit of recovery ID.
Date	The date on which the transactions occurred.
Time	The time at which the transactions occurred.
Authid	The authorization ID that made the transactions.
Nodegroup	The name of group of database partitions to which the table belongs.
Tablespace	The name of the table space to which the table belongs.
Table Owner	The name of the table owner.
Table Name	The table name.
Updates	The number of updates that occurred in that unit of recovery.
Inserts	The number of inserts that occurred in that unit of recovery.
Deletes	The number of deletes that occurred in that unit of recovery.

Generating the details

To generate the details report:

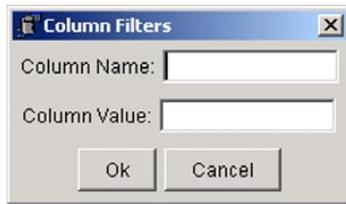
1. Double-click the row in the right frame for which you want to generate a details report.
2. When the confirmation dialog opens, select whether you want to generate Undo SQL, Redo SQL, or no SQL.



Note: When generating Undo or Redo SQL, ensure that there is a primary, or unique, key on the table for which you are generating SQL. Without a primary key, you may not get the desired results because of the inability to uniquely identify rows that were changed by the original SQL.

3. *Optional.* If you want to filter the details report and the generated SQL, click the button to the right of the text: **Set column filters**. The Column Filters dialog

appears.

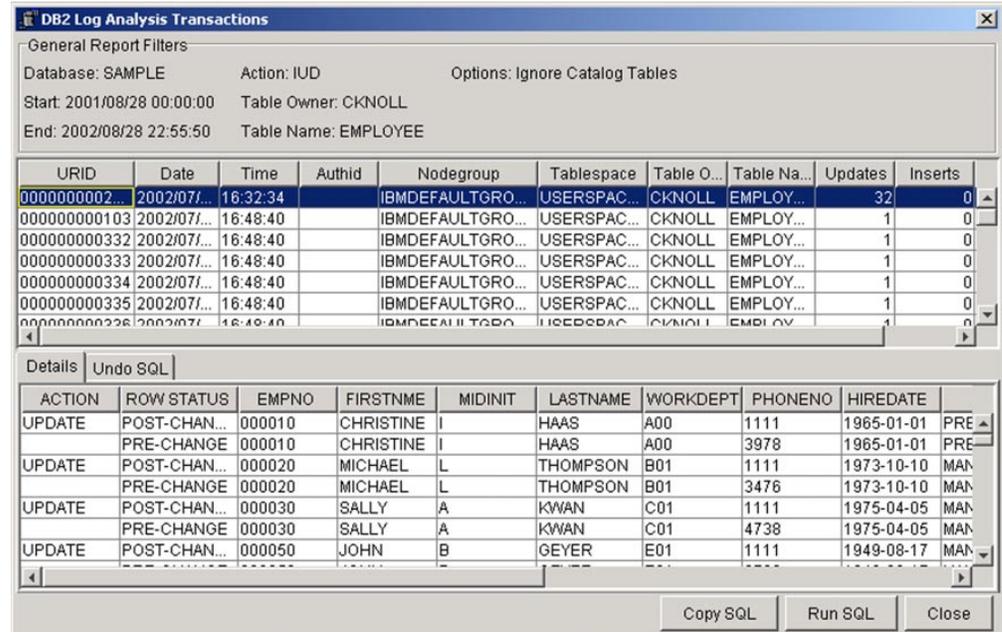


By specifying a column name and a column value, you can filter the details report and any generated SQL. Click **OK** when finished.

4. Click **OK** on the Create Details dialog to create the details report and generate the SQL, if specified.

Reviewing the details

Having generated the details, you can review the actual data changes contained in the log. If you also generated Undo or Redo SQL, you can view the SQL that you requested. Some sample details information is shown below:



The first two columns are always the same on the details page:

Table 7.

Column	Description
ACTION	Identifies what type of transaction occurred. Values are: <ul style="list-style-type: none"> • Updates • Inserts • Deletes
ROW STATUS	Identifies the status of the data in that row in the log. Values are: <ul style="list-style-type: none"> • POST-CHANGE • PRE-CHANGE

The rest of the columns in the details view are the columns from the specific table, so they will vary from table to table. By selecting different rows in the top display area, you can change the details that are shown in the details area of the display.

Reviewing the SQL

If you specified that you wanted to generate Undo or Redo SQL when you generated the details, you will see a tab next to the Details tab in the lower display area titled either Redo or Undo SQL, depending on which one you selected. Clicking that tab displays the SQL that the log analysis feature generated.

The screenshot shows the 'DB2 Log Analysis Transactions' window. The 'General Report Filters' section is configured with Database: SAMPLE, Action: IUD, and Options: Ignore Catalog Tables. The time range is from 2001/08/28 00:00:00 to 2002/08/28 22:55:50, and the Table Name is EMPLOYEE. The main table lists transactions with columns for URID, Date, Time, Authid, Nodegroup, Tablespace, Table O..., Table Na..., Updates, and Inserts. The selected row (URID 0000000002) is highlighted. Below the table, the 'Undo SQL' tab is active, displaying the following SQL:

```
-- UNDO UPDATE CKNOLL.EMPLOYEE IN URID 00000000025C
UPDATE CKNOLL.EMPLOYEE
  SET PHONENO = '5698'
 WHERE EMPNO = '000340' AND
        FIRSTNAME = 'JASON' AND
        MIDINIT = 'R' AND
        LASTNAME = 'GOUNOT' AND
        WORKDEPT = 'E21' AND
        PHONENO = '1111' AND
```

Buttons for 'Copy SQL', 'Run SQL', and 'Close' are visible at the bottom right of the window.

As before, by selecting different rows in the top display area, you can change the SQL that is shown in the SQL area of the display. If you want to, you can select and copy this SQL to a file for later use.

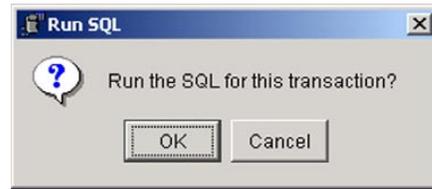
To copy the SQL:

Click **Copy SQL** and a dialog appears informing you that the SQL has been copied to the clipboard. You can now paste the SQL into any application you wish.



To run the SQL:

Click Run SQL and a confirmation dialog appears to confirm that you want to run the SQL for the selected transaction:



Click OK to run the SQL. A dialog confirms that the SQL has finished processing.

Chapter 8. Recovery Expert Versioning Repository command line interface

This chapter covers the command line options and functionality of the command line interface to the Versioning Repository of Recovery Expert. This allows you to build the object repository that holds all of the changes made to the objects in the selected database and to update the Versioning Repository with any object changes that you have made since you last ran the update job. The Recovery Expert Versioning Repository currently holds table space, table, and index database objects. See “Creating or updating the versioning repository” on page 16 for more information about the versioning repository.

Versioning Repository

The Versioning Repository tracks changes that you make to objects within your databases. Before you begin tracking those changes, you must first create the repository tables that will hold all of your object changes. There are two ways to create the versioning repository from the command line:

1.

Use CreateVersionRepository: You can type a command on the command line to create the Versioning Repository for a specified database. The command is

```
javaw -classpath location;recex.jar CreateVersionRepository
-d dbname -u username -p password
```

The options are detailed in the following table.

Option	Description
-classpath	The variable <i>location</i> identifies the full path location of the IBM DB2 JDBC driver that was installed as part of DB2. This required software is discussed in “Software requirements” on page 1.
recex.jar	The text <i>recex.jar</i> specifies the Recovery Expert Versioning Repository classes that were installed with Recovery Expert.
-d	The variable <i>dbname</i> specifies the database name. This option is required.
-u	The variable <i>username</i> specifies the user ID to establish a connection to the database. This option is required.
-p	The variable <i>password</i> specifies the password to connect to DB2. This option is required.

2.

Submit the DDL:

A DDL file to create the repository tables and table space is provided with Recovery Expert. This DDL is in a file named *objres.ddl* and can be found in the directory where you installed Recovery Expert.

To create the versioning repository:

- a. Locate *objres.ddl* in the directory where you installed Recovery Expert.
- b. Locate the following statement in *objres.ddl*:

```
CREATE NODEGROUP ORVRNG1 ON NODE (0);
```

and change the node number to the node on which you want to create the versioning repository.

Note: Ensure that you do not change the table names or the schema name.

- c. Submit `objres.ddl` using any method that you prefer. For example, you can use the DB2 command line processor to submit the DDL as a batch job or you can use QMF™ for Windows to submit it one CREATE statement at a time.

Note: Creating the Versioning Repository creates the ORVRTS table space and the many tables that comprise the versioning repository. This table space and tables are managed by DB2 Recovery Expert and the versioning repository update process. As such, this is an integral part of Recovery Expert and you should avoid performing any direct operations in this table space.

Versioning Repository update command

Once you create the Versioning Repository, you can populate it by running the repository update job. You should run the job at whatever frequency you feel is necessary to record your object changes. You run the update job using the command:

```
javaw -classpath location;recex.jar UpdateRepository
-d dbname -u username -p password
```

The options are detailed in the following table.

Option	Description
-classpath	The variable <i>location</i> identifies the full path location of the IBM DB2 JDBC driver that was installed as part of DB2. This required software is discussed in “Software requirements” on page 1.
recex.jar	The text <code>recex.jar</code> specifies the Recovery Expert Versioning Repository classes that were installed with Recovery Expert.
-d	The variable <i>dbname</i> specifies the database name. The tables in this database will be checked for changes. This option is required.
-u	The variable <i>username</i> specifies the user ID to establish a connection to the database. This option is required.
-p	The variable <i>password</i> specifies the password to connect to DB2. This option is required.

The first time that you run the repository update job, you are populating the repository tables with the catalog table information from the specified database. Subsequent runs of this update job capture only the changes that have been made to the catalog tables since the last time that you ran the update.

In addition to using the command line to submit the repository update job, you can use any viable scheduler in your environment (for example, cron or at) to allow you to run the update at intervals that you deem appropriate or overnight when there is little activity on your systems.

Saving DDL

You may encounter a situation where you want to save to a file the DDL required to restore a table or table space. This functionality is available through the command:

```
javaw -classpath location;recec.jar GenerateDDL
-d dbname -u username -p password -o objectType
-n objectSchema.objectName -f fileLocation
```

The options are detailed in the following table.

Option	Description
-classpath	The variable <i>location</i> identifies the full path location of the IBM DB2 JDBC driver that was installed as part of DB2. This required software is discussed in “Software requirements” on page 1.
recec.jar	The text <i>recec.jar</i> specifies the Recovery Expert Versioning Repository classes that were installed with Recovery Expert.
-d	The variable <i>dbname</i> specifies the database name. The tables in this database will be checked for changes. This option is required.
-u	The variable <i>username</i> specifies the user ID to establish a connection to the database. This option is required.
-p	The variable <i>password</i> specifies the password to connect to DB2. This option is required.
-o	The variable <i>objectType</i> specifies the DB2 object. Valid values are <ul style="list-style-type: none"> • table • tablespace This option is required.
-n	The variable <i>objectSchema.objectName</i> specifies the schema and name for the DDL generated for this file. This option is required.
-f	The variable <i>fileLocation</i> specifies the full windows pathname for where the DDL should be stored. If -f is omitted, the DDL will be saved in file <i>objectSchema_objectName.ddl</i> .

Chapter 9. DB2 minilog management tool command line interface

This chapter covers the command line options and functionality of the DB2 minilog management tool portion of Recovery Expert. This allows you to extract log records from the DB2 log for a certain table or set of tables (filtered by owner ID, tablespace, nodegroup or the entire database), save those records into separate files (also called the minilog), manage, and update these files.

Note: It is important that you update the minilog files regularly. It is recommended that you update the minilogs at least once a day, especially if there is heavy activity on your system.

Minilog management tool command

The minilog management tool command is db2mm followed by the options. The options are detailed in the following table.

Note: With the exception of the -d option, all other options that accept a name value are case sensitive.

Option	Description
-d	Specifies the database name. The tables in this database will be used for minilog generation. This option is required.
-O	Specifies the node number for a EEE environment. By specifying a number (0-999), you are identifying in which node you want this operation to run. If this option is omitted, the node number is 0.
-b	Specifies the start date and time for minilog generation in the format "YYYY/MM/DD HH:MM:SS". If you specify this option, any existing minilogs are overwritten.
-e	Specifies the end date and time for minilog generation in the format "YYYY/MM/DD HH:MM:SS". This value must be greater than or equal to the value of -b.
-u	Specifies the user ID to establish a connection to the database. If you omit this option, Recovery Expert uses the ID of user who is currently logged on.
-p	If this option contains a value, this value is used as a password to connect to DB2. If this option contains no value, you are prompted for a password to connect to DB2. If you omit this option, option -u uses an empty password to connect to DB2.
-N	Specifies the node group. The tables in this node group will be used for minilog generation.
-t	Specifies the table space. The tables in this tablespace will be used for minilog generation.
-o	Specifies the owner ID for the tables. The tables owned by this ID will be used for minilog generation.
-n	Specifies the table name. This table will be used for minilog generation.
-l	Specifies that the minilogs must be built for a list of tables (full table names separated with spaces).
-T	Specifies the numeric identifier of tablespace. The tables in this tablespace will be used for minilog generation.

Option	Description
-F	Specifies the numeric identifier of the table. This table will be used in minilog generation.
-S	Specifies to exclude the system tables from minilog generation.
-m	Specifies the path of the generated minilogs. If you omit this option, the current directory is used.
-P	Specifies that log records should be prepared before they get written into the minilog. Preparation means the reconstruction of data details for masked update records.
-B	Specifies base path of backup images for the analyzed database. If omitted, the tool attempts to locate the backup images automatically.
-s	Specifies the base path of system minilogs created and used by Recovery Expert. The default path is <code>./sysstorage/</code>
-L	Specifies a list of additional paths to log files for the database. If omitted, the tool only uses the path to the log files from the database configuration.
-w	Specifies paths to the intermediate work files for a session. The default work path is the current directory.
-q	Specifies that the SQL generation should be performed for minilogs that satisfy the set of filters. During SQL generation, the minilog set is not updated. The SQL code is saved in the same directory as the corresponding minilog. The values are: <ul style="list-style-type: none"> • r - Generate redo SQL code • u - Generate undo SQL code The values are mutually exclusive.
-f	Specifies the file name for the generated Redo/Undo SQL. If omitted, the SQL is generated separately for each minilog, placed in the same directory as the corresponding minilog, and has the same file name as the minilog with the addition of <code>.REDO.sql</code> or <code>.UNDO.sql</code> .
-G	Specifies direction for masked update reconstruction for SQL code generation. The values are: <ul style="list-style-type: none"> • b - always use rollback • f - always use roll forward This option must be used in conjunction with option -q, above. If you do not specify this option, Recovery Expert will choose the best available direction.
-c	Specifies commit scope for generated redo or undo SQL code. If you omit this option, the commit scope is one.
-r	Specifies that the minilogs must be regenerated, not just updated.
-k	Specifies that no new minilogs are to be created. Only existing minilogs are to be updated or rewritten.
-l	Prints a report about the available minilogs in the base location directory specified by option -m. If you use -l, you do not have to specify database name.
-h	Displays help.
-v	Displays the version.

Command line examples

Below are two examples of command line statements and a description of what occurs.

Example 1

```
db2mm -d SAMPLE -o Q -n STAFF -m /home/db2inst/minilogs
```

This command generates (or updates) the minilog for table Q.STAFF of database SAMPLE and saves it in /home/db2inst/minilogs.

Example 2

```
db2mm -d SAMPLE -t DSQTSTBT  
-m /home/db2inst/minilogs/ -r
```

This command regenerates the minilogs for all tables in tablespace DSQTSTBT of database SAMPLE, saving them in /home/db2inst/minilogs.

Note: To run db2mm in a non-Windows environment, you must use db2recex to run the minilog management tool by typing:

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex
```

before the db2mm command. For example,

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex db2mm  
-d SAMPLE -o Q -n STAFF
```

Minilog management tool output

When you run the minilog management tool, there is output that the tool writes to the screen as well as the output that gets written to files.

Screen output

The output that the minilog management tool writes to the screen is a series of messages detailing what the tool is creating. The format of this output is:

```
tableschematablename (tid.fid)
```

where tid is the tablespace ID and fid is the table ID (also known as the file ID). For example, the statement

```
c:>db2mm -d resample -t RETSSMP
```

produces the following screen output:

```
New table RE.DEPARTMENT (3.14) found. Minilog was created.  
New table RE.ORG (3.2) found. Minilog was created.  
New table RE.STAFF (3.3) found. Minilog was created.  
New table RE.EMPLOYEE (3.13) found. Minilog was created.  
New table RE.INTERVIEW (3.7) found. Minilog was created.  
New table RE.SUPPLIER (3.9) found. Minilog was created.
```

File output

After you run the minilog management tool, the output directory contains a subtree with the structure: ./DBNAME/NODEGROUPNAME/TABLESPACE/NAME. In that directory is a distributed set of minilogs for the tables (each minilog is actually a pair of files with

the extensions .mlg and .mlx). Additionally, there is a pair of files Catalog.mlg and Catalog.mlx that contain meta information about the minilogs stored in the output directory.

Catalog.mlg format

The file Catalog.mlg consists of a 12-byte header and a number of blocks (one block per minilog).

Format of header:

Description	Length	Comment
Tune-up sequence	4	4-byte int, equal to 1. This is used to determine whether the file contains little- or big-endian data.
Number of blocks	4	
Reserved	4	

Catalog.mlx format

The file Catalog.mlx contains 4-byte offsets of blocks as listed in the following table.

Format of blocks:

Description	Length
Timestamp of end of coverage	4
Reserved	4
Length of database name	4
Database name	variable
Length of node group name	4
Node group name	variable
Length of tablespace name	4
Tablespace name	variable
Length of table owner	4
Table owner	variable
Length of table name	4
Table name	variable
Length of TID	4
TID	variable
Length of FID	4
FID	variable
Length of URID	4
URID of first recorded transaction	variable
Length of timestamp	4
Timestamp of first recorded transaction	variable
Length of URID	4
URID of last recorded transaction	variable
Length of timestamp	4
Timestamp of last recorded transaction	variable

Description	Length
Length of timestamp	4
Timestamp of beginning of coverage	variable
Length of timestamp	4
Timestamp of end of coverage	variable
Length of number of INSERTs	4
Number of INSERTs	variable
Length of number of DELETEs	4
Number of DELETEs	variable
Length of number of UPDATEs	4
Number of UPDATEs	variable
Length of number of transactions	4
Number of transactions	variable

Scenarios for use

In the following scenarios the system minilogs are stored in /home/db2inst/sysstorage and could be created by any of tools of object restore suite, using library db2ors.

Scenario 1

You want to create and maintain the full set of minilogs for the entire database DBNAME in the directory /home/db2inst/minilogs, except for the minilogs for system tables. To create or to update the minilogs at any time, use the following command:

```
db2mm -d DBNAME -s /home/db2inst/sysstorage
-m /home/db2inst/minilogs -S
```

Scenario 2

You want to create and maintain the full set of minilogs for entire tablespace TBSPACE of database DBNAME in the directory /home/db2inst/minilogs. To create or to update the minilogs at any time, use the following command: db2mm -d DBNAME -t TBSPACE -s /home/db2inst/sysstorage -m /home/db2inst/minilogs

Scenario 3

You want to create and maintain the partial set of minilogs for several essential tables (OWNER1.TABLE1, OWNER2.TABLE2, ..., OWNERN.TABLEN) of database DBNAME in the directory /home/db2inst/minilogs/Essential. You would use to following commands:

To create the minilogs:

```
db2mm -d DBNAME -o OWNER1 -n TABLE1 -s /home/db2inst/sysstorage
-m /home/db2inst/minilogs
db2mm -d DBNAME -o OWNER2 -n TABLE2 -s /home/db2inst/sysstorage
-m /home/db2inst/minilogs
...
db2mm -d DBNAME -o OWNERN -n TABLEN -s /home/db2inst/sysstorage
-m /home/db2inst/minilogs
```

To update any single table's minilog from the set:

```
db2mm -d DBNAME -o OWNERX -n TABLEX -s /home/db2inst/sysstorage  
-m /home/db2inst/minilogs
```

To update the entire set of minilogs:

```
db2mm -d DBNAME -k -s /home/db2inst/sysstorage  
-m/home/db2inst/minilogs
```

Chapter 10. Log analysis command line interface

This chapter covers the command line options and functionality of the log analysis portion of Recovery Expert. This allows you to retrieve information about database activity. Specifically, you can extract the INSERT, UPDATE, and DELETE operations and create a report detailing that information. The data for report are retrieved from the active and archive log files.

Log analysis command

The log analysis command is `db21a` followed by the options. The options are detailed in the following table.

Note: With the exception of the `-d` option, all other options that accept a name value are case sensitive.

Option	Description
<code>-d</code>	Specifies the database name. Operations in this database will be included in the generated report. This option is required.
<code>-O</code>	Specifies the node number for a EEE environment. By specifying a number (0-999), you are identifying in which node you want this operation to run. If this option is omitted, the node number is 0.
<code>-b</code>	Specifies the start date and time for log analysis in the format "YYYY/MM/DD HH:MM:SS". If you omit this option, the analysis begins at the beginning of the entire log.
<code>-e</code>	Specifies the end date and time for log analysis in the format "YYYY/MM/DD HH:MM:SS". This value must be greater than or equal to the value of <code>-b</code> . If you omit this option, the analysis runs to the end of the entire log.
<code>-a</code>	Specifies the operations to show in the report. The values are: <ul style="list-style-type: none">• <code>i</code> - to show INSERTs• <code>u</code> - to show UPDATEs• <code>d</code> - to show DELETEs The values listed above can be combined into any set like <code>iud</code> , <code>du</code> , <code>id</code> , etc. If you omit this option, all operations are shown in the report.
<code>-V</code>	Specifies the content level for reports. The values are: <ul style="list-style-type: none">• 0 - Maximum content. All operations within a transaction listed in their order and not grouped in the general report.• 1 - High content. All repeated operations (featuring same object and same action) are grouped within a transaction in the general report.• 2 - Medium content. The general report contains numbers of operations within a transaction grouped by table names.• 3 - Minimum content. The general report shows grand totals only.• 4 - Silent mode. The reports produce no text output; they generate raw storages only. If you omit this option, the default content level is 0. For detailed reports, modes 0 to 3 produce the same report.

Option	Description
-U	Specifies the authorization ID of operations to be included into the generated report. If you omit this option, the report will contain the data about the operations performed under all authorization IDs. Note: To filter on authorization ID, you must have had CHANGE DATA CAPTURE enabled in the DB2 tableconfiguration when the changes occurred. This captures the AUTHID in the log, allowing log analysis to filter on that column.
-u	Specifies the user ID to establish a connection to the database.If you omit this option, Recovery Expert uses the ID of user who is currently logged on.
-p	If this option contains a value, this value is used as a password to connect to DB2. If this option contains no value, the user is prompted for a password to connect to DB2. If you omit this option, option -u uses an empty password to connect to DB2.
-N	Specifies the node group. Operations in this node group will be included in the generated report. If you omit this option and do not specify a table name or table owner, the report will contain the data about the operations in all node groups.
-t	Specifies the table space. The operations in this tablespace will be included in the generated report. If you omit this option and do not specify a table name or table owner, the report will contain the data about the operations in all table spaces.
-o	Specifies the owner ID for the tables. The operations performed by this owner ID will be included in the generated report. If you omit this option, the report will contain the data about the operations performed on the tables by all owner IDs.
-n	Specifies the table name. The operations performed on this table will be included in the generated report. If you omit this option, the report will contain the data about operations performed on all tables.
-l	Specifies that the report must be built for a list of tables (full table names separated with spaces).
-R	Specifies the transaction ID of operations that you want included in the generated report. If you omit this option, the report will contain the data about the operations performed in all transactions.
-T	Specifies the numeric identifier of tablespace. The operations in this tablespace will be included in the generated report.
-F	Specifies the numeric identifier of the table. The operations in this table will be included in the generated report.
-S	Specifies to exclude the system tables from the generated report.
-D	Specifies that the application run a detailed report based on the data collected in the general report.
-B	Specifies the base path of backup images for an analyzed database. If you omit this option, Recovery Expert will attempt to locate backup images automatically.
-s	Specifies the base path of system minilogs created and used by Recovery Expert. The default path is ./sysstorage/
-A	Specifies that log analysis runs in catalog only mode. This means that all data about the tables comes from the current system catalog only, and create/drop/alter events for tables are ignored. Only use this option to get reports on most recent log activity or for tables that are never dropped or altered.

Option	Description
-C	Specifies name of the column (COL_NAME) for filtering condition COL_NAME=COL_VALUE. This option must be used in conjunction with option -E.
-E	Specifies value of the column (COL_VALUE) for filtering condition COL_NAME=COL_VALUE. This option must be used in conjunction with option -C.
-G	Specifies direction for masked update reconstruction. The values are: <ul style="list-style-type: none"> • b - always use rollback • f - always use roll forward This option must be used in conjunction with option -D. If you do not specify this option, Recovery Expert will choose the best available direction.
-i	Specifies that the report must be built only for a specific operations' range within a single transaction. Option can only be used together with options -D and -R. The values are: <ul style="list-style-type: none"> • n - show details on n-th operation of transaction only • :n - show details on first n operations of transaction • n: - show details on operations from n-th to the end of transaction • n:m - show details on operations from n-th to m-th
-q	Specifies that the SQL generation should be performed while the detailed report is being generated. The values are: <ul style="list-style-type: none"> • r - Generate redo SQL code • u - Generate undo SQL code The values are mutually exclusive.
-f	Specifies the file name for the generated redo or undo SQL code. If you omit this option, the file name is redo.sql or undo.sql, respectively. The redo or undo SQL is placed in the data/<instance_name> directory.
-c	Specifies commit scope for generated redo or undo SQL code. If you omit this option, the commit scope is one.
-h	Displays help.
-v	Displays the version.

Command line examples

Below are two examples of command line statements and a description of what they will produce.

Example 1

```
db2lta -b "2001/01/01 00:00:00"
-e "2001/02/01 00:00:00" -a iu -U DB2ADMIN -d SAMPLE
-o Q -n STAFF
```

This command generates a report containing all INSERTs and UPDATEs made by user DB2ADMIN on database SAMPLE to table Q.STAFF between midnight of January 1st, 2001, and midnight of February 1st, 2001.

Example 2

```
db2la -b "2001/01/01 11:55:00"  
-e "2001/01/01 12:00:00" -a d -U DB2GUEST  
-d TEST -o Q -n ORG
```

This command generates a report containing all DELETES made by user DB2GUEST on database TEST to table Q.ORG between 11:55am and 12:00pm of January 1st, 2001.

Note: To run db2la in a non-Windows environment, you must use db2recex to run the log analysis tool by typing:

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex
```

before the db2la command. For example,

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex db2la  
-b "2001/01/01 00:00:00" -e "2001/02/01 00:00:00"  
-a iu -U DB2ADMIN -d SAMPLE -o STAFF
```

Sample reports

Maximum content report:

```
DB2 LOG ANALYSIS- GENERAL REPORT  
BUILT ON YYYY/MM/DD AT HH:MM:SS  
*****
```

FILTERS

```
DATABASE : <DBNAME>  
{START DATE : YYYY/MM/DD}  
{START TIME : HH:MM:SS}  
{END DATE : YYYY/MM/DD}  
{END TIME : HH:MM:SS}  
{ACTION : <ACTION>}  
{NODEGROUP : <NGNAME>}  
{TABLESPACE : <TSNAME>}  
{OWNER : <OWNER>}  
{TABLENAME : <TNAME>}  
{TABLE LIST : <NAME1> <NAME2> ... <NAMEN>}  
{AUTHID : <AUTHID>}  
{URID : <URID>}  
{TID : <TID>}  
{FID : <FID>}  
{OPTIONS : IGNORE SYSTEM TABLES}
```

=====

```
URID      DATE      TIME      AUTHID  
-----  
<URID1>   YYYY-MM-DD HH:MM:SS <USERID>  
-----
```

```
NGNAME    TSNAME    TABLE OWNER  TABLE NAME  ACTION  
-----  
<NGNAME1> <TSNAME1> <TOWNER1>    <TNAME1>    <ACT1>  
<NGNAME2> <TSNAME2> <TOWNER2>    <TNAME2>    <ACT2>  
<NGNAME3> <TSNAME3> <TOWNER3>    <TNAME3>    <ACT3>
```

=====

```
URID      DATE      TIME      AUTHID  
-----  
<URID2>   YYYY-MM-DD HH:MM:SS <USERID>  
-----
```

NGNAME	TSNAME	TABLE OWNER	TABLE NAME	ACTION
<NGNAME1>	<TSNAME1>	<TOWNER1>	<TNAME1>	<ACT1>
<NGNAME2>	<TSNAME2>	<TOWNER2>	<TNAME2>	<ACT2>
<NGNAME3>	<TSNAME3>	<TOWNER3>	<TNAME3>	<ACT3>

SUMMARY

NODEGROUP	UPDATES	INSERTS	DELETES
<NGNAME1>	<NU1>	<NI1>	<ND1>
<NGNAME2>	<NU2>	<NI2>	<ND2>
<NGNAME3>	<NU3>	<NI3>	<ND3>

TABLESPACE	UPDATES	INSERTS	DELETES
<TSNAME1>	<NU1>	<NI1>	<ND1>
<TSNAME2>	<NU2>	<NI2>	<ND2>
<TSNAME3>	<NU3>	<NI3>	<ND3>

TABLE	UPDATES	INSERTS	DELETES
<TNAME1>	<NU1>	<NI1>	<ND1>
<TNAME2>	<NU2>	<NI2>	<ND2>
<TNAME3>	<NU3>	<NI3>	<ND3>

TOTAL SUMMARY

TOTAL UPDATES: <NUT>
TOTAL INSERTS: <NIT>
TOTAL DELETES: <NDT>

High content report:

DB2 LOG ANALYSIS- GENERAL REPORT
BUILT ON YYYY/MM/DD AT HH:MM:SS

FILTERS

DATABASE : <DBNAME>
{START DATE : YYYY/MM/DD}
{START TIME : HH:MM:SS}
{END DATE : YYYY/MM/DD}
{END TIME : HH:MM:SS}
{ACTION : <ACTION>}
{NODEGROUP : <NGNAME>}
{TABLESPACE : <TSNAME>}
{OWNER : <OWNER>}
{TABLENAME : <TNAME>}
{TABLE LIST : <NAME1> <NAME2> ... <NAMEN>}
{AUTHID : <AUTHID>}
{URID : <URID>}
{TID : <TID>}
{FID : <FID>}
{OPTIONS : IGNORE SYSTEM TABLES}

URID	DATE	TIME	AUTHID
<URID1>	YYYY-MM-DD	HH:MM:SS	<USERID>

NGNAME	TSNAME	TABLE OWNER	TABLE NAME	ACTION
--------	--------	-------------	------------	--------

```

<NGNAME1> <TSNAME1> <TOWNER1> <TNAME1> <ACT1> [N1]
<NGNAME2> <TSNAME2> <TOWNER2> <TNAME2> <ACT2> [N2]
<NGNAME3> <TSNAME3> <TOWNER3> <TNAME3> <ACT3> [N3]

```

```

=====
URID          DATE          TIME          AUTHID
-----
<URID2>      YYYY-MM-DD HH:MM:SS <USERID>
-----

```

```

NGNAME      TSNAME      TABLE OWNER  TABLE NAME  ACTION
-----
<NGNAME1> <TSNAME1> <TOWNER1> <TNAME1> <ACT1> [N1]
<NGNAME2> <TSNAME2> <TOWNER2> <TNAME2> <ACT2> [N2]
<NGNAME3> <TSNAME3> <TOWNER3> <TNAME3> <ACT3> [N3]

```

SUMMARY

```

-----
NODEGROUP          UPDATES  INSERTS  DELETES
-----
<NGNAME1>          <NU1>   <NI1>   <ND1>
<NGNAME2>          <NU2>   <NI2>   <ND2>
<NGNAME3>          <NU3>   <NI3>   <ND3>

```

```

-----
TABLESPACE         UPDATES  INSERTS  DELETES
-----
<TSNAME1>          <NU1>   <NI1>   <ND1>
<TSNAME2>          <NU2>   <NI2>   <ND2>
<TSNAME3>          <NU3>   <NI3>   <ND3>

```

```

-----
TABLE              UPDATES  INSERTS  DELETES
-----
<TNAME1>           <NU1>   <NI1>   <ND1>
<TNAME2>           <NU2>   <NI2>   <ND2>
<TNAME3>           <NU3>   <NI3>   <ND3>

```

TOTAL SUMMARY

```

-----
TOTAL UPDATES: <NUT>
TOTAL INSERTS: <NIT>
TOTAL DELETES: <NDT>

```

Medium content report:

```

DB2 LOG ANALYSIS- GENERAL REPORT
                BUILT ON YYYY/MM/DD AT HH:MM:SS
*****

```

FILTERS

```

-----
DATABASE   : <DBNAME>
{START DATE : YYYY/MM/DD}
{START TIME : HH:MM:SS}
{END DATE   : YYYY/MM/DD}
{END TIME   : HH:MM:SS}
{ACTION     : <ACTION>}
{NODEGROUP  : <NGNAME>}
{TABLESPACE : <TSNAME>}
{OWNER      : <OWNER>}
{TABLENAME  : <TNAME>}
{TABLE LIST : <NAME1> <NAME2> ... <NAMEN>}
{AUTHID     : <AUTHID>}
{URID       : <URID>}
{TID        : <TID>}
{FID       : <FID>}

```

{OPTIONS : IGNORE SYSTEM TABLES}

```
=====
URID          DATE          TIME          AUTHID
-----
<URID1>      YYYY-MM-DD HH:MM:SS <USERID>
=====
```

```
NGNAME      TSNAME      TABLE OWNER  TABLE NAME  UPDATES  INSERTS  DELETES
-----
<NGNAME1> <TSNAME1> <TOWNER1> <TNAME1> <NU1> <NI1> <ND1>
<NGNAME2> <TSNAME2> <TOWNER2> <TNAME2> <NU2> <NI2> <ND2>
<NGNAME3> <TSNAME3> <TOWNER3> <TNAME3> <NU3> <NI3> <ND3>
```

```
=====
URID          DATE          TIME          AUTHID
-----
<URID2>      YYYY-MM-DD HH:MM:SS <USERID>
=====
```

```
NGNAME      TSNAME      TABLE OWNER  TABLE NAME  UPDATES  INSERTS  DELETES
-----
<NGNAME1> <TSNAME1> <TOWNER1> <TNAME1> <NU1> <NI1> <ND1>
<NGNAME2> <TSNAME2> <TOWNER2> <TNAME2> <NU2> <NI2> <ND2>
<NGNAME3> <TSNAME3> <TOWNER3> <TNAME3> <NU3> <NI3> <ND3>
```

SUMMARY

```
NODEGROUP          UPDATES  INSERTS  DELETES
-----
<NGNAME1>          <NU1> <NI1> <ND1>
<NGNAME2>          <NU2> <NI2> <ND2>
<NGNAME3>          <NU3> <NI3> <ND3>
```

```
TABLESPACE          UPDATES  INSERTS  DELETES
-----
<TSNAME1>          <NU1> <NI1> <ND1>
<TSNAME2>          <NU2> <NI2> <ND2>
<TSNAME3>          <NU3> <NI3> <ND3>
```

```
TABLE              UPDATES  INSERTS  DELETES
-----
<TNAME1>          <NU1> <NI1> <ND1>
<TNAME2>          <NU2> <NI2> <ND2>
<TNAME3>          <NU3> <NI3> <ND3>
```

TOTAL SUMMARY

TOTAL UPDATES: <NUT>
TOTAL INSERTS: <NIT>
TOTAL DELETES: <NDT>

Minimum content report:

DB2 LOG ANALYSIS- GENERAL REPORT
BUILT ON YYYY/MM/DD AT HH:MM:SS

FILTERS

DATABASE : <DBNAME>
{START DATE : YYYY/MM/DD}
{START TIME : HH:MM:SS}
{END DATE : YYYY/MM/DD}
{END TIME : HH:MM:SS}

```

{ACTION      : <ACTION>}
{NODEGROUP   : <NGNAME>}
{TABLESPACE  : <TSNAME>}
{OWNER       : <OWNER>}
{TABLENAME   : <TNAME>}
{TABLE LIST  : <NAME1> <NAME2> ... <NAMEN>}
{AUTHID      : <AUTHID>}
{URID        : <URID>}
{TID         : <TID>}
{FID         : <FID>}
{OPTIONS     : IGNORE SYSTEM TABLES}

```

SUMMARY

NODEGROUP	UPDATES	INSERTS	DELETES
<NGNAME1>	<NU1>	<NI1>	<ND1>
<NGNAME2>	<NU2>	<NI2>	<ND2>
<NGNAME3>	<NU3>	<NI3>	<ND3>

TABLESPACE	UPDATES	INSERTS	DELETES
<TSNAME1>	<NU1>	<NI1>	<ND1>
<TSNAME2>	<NU2>	<NI2>	<ND2>
<TSNAME3>	<NU3>	<NI3>	<ND3>

TABLE	UPDATES	INSERTS	DELETES
<TNAME1>	<NU1>	<NI1>	<ND1>
<TNAME2>	<NU2>	<NI2>	<ND2>
<TNAME3>	<NU3>	<NI3>	<ND3>

TOTAL SUMMARY

```

TOTAL UPDATES: <NUT>
TOTAL INSERTS: <NIT>
TOTAL DELETES: <NDT>

```

Detailed report:

```

DB2 LOG ANALYSIS - DETAILED REPORT
BUILT ON YYYY/MM/DD AT HH:MM:SS
*****
BASED ON PREVIOUSLY BUILT GENERAL REPORT

```

ADDITIONAL FILTERS

```

DATABASE      : <DBNAME>
{START DATE   : YYYY/MM/DD}
{START TIME   : HH:MM:SS}
{END DATE     : YYYY/MM/DD}
{END TIME     : HH:MM:SS}
{ACTION       : <ACTION>}
{NODEGROUP    : <NGNAME>}
{TABLESPACE   : <TSNAME>}
{OWNER        : <OWNER>}
{TABLENAME    : <TNAME>}
{TABLE LIST   : <NAME1> <NAME2> ... <NAMEN>}
{AUTHID       : <AUTHID>}
{URID         : <URID> {[N:M] |[N] |[BEGIN:N] |[N|END] }}
{TID          : <TID>}
{FID          : <FID>}
{COLUMN       : COLNAME = COLVALUE}
OPTIONS       : {IGNORE SYSTEM TABLES}

```

```

-----
ACTION DATE      TIME      URID      AUTHID
-----
<ACT1> YYYY/MM/DD HH:MM:SS <URID1>      <AUTHID>

NGNAME          TSNAME          TABLE OWNER    TABLE NAME
-----
<NGNAME1>      <TSNAME1>      <TOWNER1>      <TNAME1>

ROW STATUS <COL1>    <COL2>    <COL3>    <COLN>
-----
POST-CHANGE <POSTVAL1> <POSTVAL2> <POSTVAL3> ... <POSTVALN>
PRE-CHANGE <PREVAL1> <PREVAL2> <PREVAL3>    <PREVALN>

...

```

```

-----
ACTION DATE      TIME      URID      AUTHID
-----
<ACT2> YYYY/MM/DD HH:MM:SS <URID1>      <AUTHID>

NGNAME          TSNAME          TABLE OWNER    TABLE NAME
-----
<NGNAME2>      <TSNAME2>      <TOWNER2>      <TNAME2>

ROW STATUS <COL1>    <COL2>    <COL3>    <COLN>
-----
POST-CHANGE <POSTVAL1> <POSTVAL2> <POSTVAL3> ... <POSTVALN>
PRE-CHANGE <PREVAL1> <PREVAL2> <PREVAL3>    <PREVALN>

```

Chapter 11. DB2 object translation tool command line interface

This chapter covers the command line options and functionality of the DB2 object translation tool portion of Recovery Expert. This tool allows you to extract a single object from a DB2 backup file (single table mode) or to clone the backup file with a translation of object identifiers to make a copy of the backup image that would allow you to perform a complete database recovery (full mode).

DB2 object translation tool command

The DB2 object translation tool command is `db2ox` followed by the options. The options are detailed in the following table.

Note: With the exception of the `-d` option, all other options that accept a name value are case sensitive.

Option	Description
-d	Specifies the database name. The backup images of this database will be used as the data source. This option is required.
-O	Specifies the node number for a EEE environment. By specifying a number (0-999), you are identifying in which node you want this operation to run. If this option is omitted, the node number is 0.
-b	Specifies the timestamp of the backup image that should be used as the data source (the format is "YYYY/MM/DD HH:MM:SS"). This option is required.
-f	Specifies the fully qualified output file name without the extension. This option is required in full mode. If you omit this option in single table mode, the output file is written into current directory, and its name is determined by the new table identifier.
-u	Specifies the user ID to establish a connection to the database. If you omit this option, Recovery Expert uses the ID of user who is currently logged on.
-p	If this option contains a value, this value is used as a password to connect to DB2. If this option contains no value, you are prompted for a password to connect to DB2. If you omit this option, option <code>-u</code> uses an empty password to connect to DB2.
-o	Specifies the owner ID for the table in single table mode. You must use this option in conjunction with option <code>-n</code> .
-n	Specifies the table name in single table mode. You must use this option in conjunction with option <code>-o</code> .
-T	Specifies the numeric identifier of the tablespace for a table in single table mode. You must use this option in conjunction with option <code>-F</code> .
-F	Specifies the numeric identifier of a table in single table mode. You must use this option in conjunction with option <code>-T</code> .
-B	Specifies the base path of backup images for a database. If you omit this option, Recovery Expert will attempt to locate backup images automatically.

Option	Description
-X	Specifies the object ID translation map in the following format: OLDTID1:OLDFID1-NEWTID1:NEWFID1 OLDID2:OLDFID2-NEWTID2:NEWFID2 . . . If you omit this option, no object ID translation occurs. If the old TID and FID in single table mode are both equal to zero, the translation occurs regardless of the actual table identifiers.
-h	Displays help.
-v	Displays the version.

Object translation constraints

- The object translation tool does not support foreign keys or any referential integrity between tables.
- The object translation tool always extracts a single object at a time from a backup image. Use multiple runs of db2ox to get more than one object.
- The object translation tool does not support multiple instances or tables spread across multiple containers.

Command line examples

Below are two examples of command line statements and a description of what they will produce.

Example 1

```
db2ox -d SAMPLE -o Q -n STAFF -b "2001/01/20 12:20:23"
-f SQL00002
```

This command Extracts the image of table Q.STAFF from the DB2 backup file of database SAMPLE created at 12:20:23 on January 20, 2001 and puts the table image in file SQL00002.DAT.

Example 2

```
db2ox -d SAMPLE -b "2001/01/20 11:20:00" -f 112000 -X
"12:23-6:12 12:24-6:13"
```

This command does several things:

- Creates a copy of the backup image of database SAMPLE created at 11:20:00 on January 20, 2001.
- Changes object identifiers for object with TID=12 and FID=23 to TID=6 and FID=12.
- Changes object identifiers for object with TID=12 and FID=24 to TID=6 and FID=13.
- Writes a result backup image to file(s) 112000.001, (112000.002, etc.).

Note: To run db2ox in a non-Windows environment, you must use db2recex to run the object translation tool by typing:

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex
```

before the db2ox command. For example,

```
/opt/IBM/DB2RecoveryExpert/bin/db2recex db2ox -d SAMPLE
-o Q -n STAFF -b "2001/01/20 12:20:23" -f SQL00002
```

This is because db2recex is a “set-UID” program that changes its Real and Effective User and Group IDs to that of the instance user, then runs the requested program.

Appendix. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other 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 10594-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country 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 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 publication 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 as your own risk.

Licensees of this program who want 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 which has been exchanged, should contact:

IBM Corporation

J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

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 information 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 measurement 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.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains 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 contains sample application programs in source language, which illustrates 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.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

AIX
DB2
DB2 Universal Database
IBM
QMF

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

Index

A

- accessing 16
- accessing log analysis 16
- AIX 5
- APPLHEAPSZ 2
- application heap size 2
- authority 1
- availability of backup images 2

B

- backup images 2
- Backup images 21
- Backup location 25
- Backup location and file name 25

C

- Catalog Data 40
- Column Filters 44
- command line
 - enabling 11
 - log analysis 59
 - log analysis reports 62, 63, 64, 65, 66
 - minilogs 53
 - object translation tool 69
 - versioning repository 49
- comments, how to send v
- Commit Scope 24
- CONF_DIR 10
- configuration directory 10
- configuration file 10
- configuration values, default 11
- connecting to a database 15
- containers 25
- copy SQL 46
- cost for recovery 23
- CreateVersionRepository 49
- creating the versioning repository 16, 49
- customer support v

D

- data directory 10
- DATA_DIR 10
- Database 14, 39
- Database state 14
- Database Status 14
- DB2 object translation tool 69
- DB2 Recovery Expert
 - exiting 15
 - installing 5
 - launching 13
- DB2 UDB EEE 12, 15, 53, 59, 69
- DB2 Universal Database Enterprise - Extended Edition 12
- db2diag.log vi

- db2la 59
- db2mm 53
- db2ox 69
- db2reaixsetup 5
- db2recex.log vi
- db2rehpuxsetup 5
- db2relinuxsetup 5
- db2resolarissetup 5
- db2rewin32setup 5
- default configuration values 11
- Degree of parallelism 24
- details report 45
- diaglevel vi
- DMS 26, 33

E

- editing configuration file 11
- EEE 12, 15, 53, 59, 69
- enabling command line tools 11
- End Date 40
- End of Logs 21
- End Time 40
- error
 - log file vi
 - logging vi
- exiting 15

F

- filtering data 44
- First catalog log 15

G

- generate SQL 44
- GenerateDDL command 50

H

- HPUX 5

I

- install directory 7, 10
- INSTALL_DIR 10
- installing 5

L

- Last catalog log 15
- Last catalog update 15
- launching DB2 Recovery Expert 13
- Linux 5
- Log analysis 16, 39
 - command line 59
 - Detailed report 66

- Log analysis (*continued*)
 - details report 44
 - filtering data 44
 - general report 41
 - generate SQL 44
 - High content report 63
 - Maximum content report 62
 - Medium content report 64
 - Minimum content report 65
 - transaction-level view 43
- Log file directories or tapes 15
- log files 2
- Log retain 14
- logging vi
- logging settings 2
- LOGRETAIN 2

M

- minilog 53
 - output 55
- multiple objects 3

N

- node number 15, 53, 59, 69
- Notices 73
- Number of buffers 24

O

- Object Selection 29, 35
 - database 19
- objres.ddl 49
- online support v
- ORVRTS 17, 50
- other points-in-time 21

P

- Point-in-time Selection 20
- primary key 44
- privileges 1
- Process while 24

Q

- Quiesce point 21

R

- recex.properties 10, 11
- recovering to transaction 21
- recovery assets 23
- recovery cost 23
- Recovery Options 24
- Recovery Path 23
- Redirect table space containers 25
- Redirected restore 25
- redo SQL 44

- referential integrity 3
- Report options 40
- required settings 2
- Restore a dropped database 14
- reviewing SQL 46
- RI 3
- Roll-forward recovery enabled 15
- run SQL 46

S

- saving DDL 50
- Scan Archive Logs 40
- Select a table space to recover 30
- Select a table to recover 36
- sending comments to IBM v
- set-UID 10
- Show DELETES 40
- Show INSERTS 40
- Show UPDATES 40
- Size of each buffer in 4KB pages 24
- SMS 26
- Solaris 5
- SQL 46
- Start Date 40
- Start Time 40
- State of database after restore 24
- Summary 26
- support v

T

- table 35
- Table Name 41
- Table Owner 40
- table space 29
 - DMS 26, 33
 - SMS 26
- table space containers 25
- table space selection 30
- Tablespace Name 40
- Tools
 - Log Analysis 16
 - Versioning Repository 16, 17
- transaction-level view 43
- transaction, recovering to 21

U

- undo SQL 44
- unique key 44
- UNIX 5
- updating the versioning repository 17, 50
- URID 21
- URID Finder 22
- User exit 14
- User ID 41
- USEREXIT 2
- using a storage manager 2

V

Versioning Repository 16
 command line 49
 create command 49
 creating 16
 GUI 16
 update command 50
 updating 17
Versions 32

W

Windows 5

Readers' Comments — We'd Like to Hear from You

IBM DB2 Recovery
Expert for Multiplatforms
User's Guide
Version 1 Release 1

Publication No. SC27-1624-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



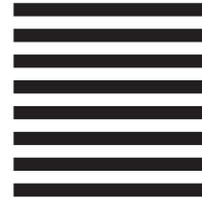
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Department 090/H150
555 Bailey Avenue
San Jose, CA 95141-9989
U.S.A.



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

SC27-1624-00

