

December 2010

IBM Information Management software



DB2 Compression Estimation Tool

*George Baklarz
Information Management Technical Pre-sales Support
IBM Software Group*

Introduction¹

As the amount of data being managed continues to rise, the cost of storage is becoming an increasing burden to CIOs and IT executives. Although the cost of disk drives as measured in dollars per gigabyte of storage continues to decline, the overall IT expense on storage infrastructure continues to climb. This is as a result of the explosive growth in storage requirements of both structured and unstructured data that needs to be managed. Looking at the cost of database storage along with the cost of backup media and disaster recovery storage, a significant amount of money can be saved if the database were to consume substantially less space.

DB2 (Version 9, 9.5, and 9.7 for distributed platforms) brings to the table a significant advancement with its new compression feature. This feature allows databases to be compressed by 45% on average thus reducing storage costs and improving performance in many cases. In addition, backup images are now half the size and therefore backup/restore operations now finish in almost half the time and consume half the space. Multiply that by the number of backup images you keep and the savings grow even larger. In addition, disaster recovery sites now require almost half the amount of storage as well. Looking at the cost of enclosures, controllers, power, cooling, floor space, etc., you can quickly see that DB2 Compression can save a significant amount of money for the enterprise.

What is row compression

DB2 uses a dictionary based Lempel-Ziv form of compression. Repeating patterns within a table are extracted from the rows and stored once in a dictionary. These patterns are then replaced with a symbol within the row itself. As shown in Figure 1, repeating column values, as well as strings that repeat across column boundaries, can be replaced.

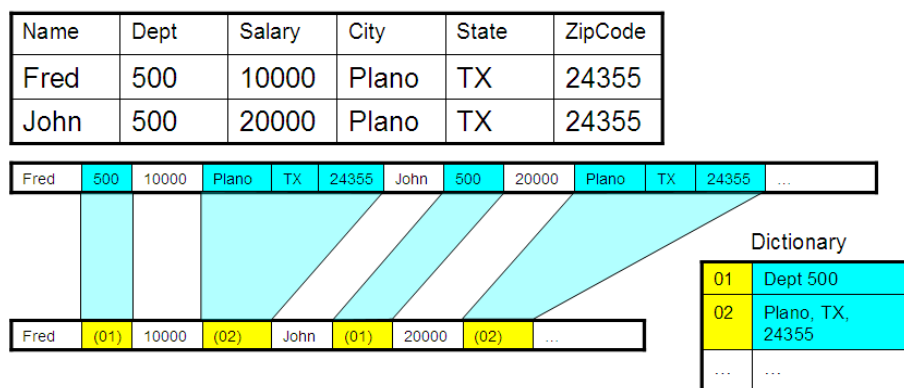


Figure 1: Compressing data in a row

Using this form of compression, DB2 can scan the entire table, find repeating strings within the data and replace those strings with compression symbols. The amount of compression achieved in this manner on real customer data during the DB2 9 test drive has been reported at between 50% and 80%. That is, the size of the table after compression is between 50% and 80% smaller than when it started. Thus the storage consumption for these tables is

1. Compression in DB2 Viper, Chris Eaton, Toronto Laboratory, IBM Canada Limited, May, 2006

significantly less. In addition, DB2 keeps this data compressed in the buffer pool so that you effectively have twice the amount of memory available. That is, since rows remain compressed in the buffer pool, you can expect to see twice as many rows being stored in memory compared to non-compressed data. In one test of an SAP table with real customer data, the number of rows packed into the bufferpool was in fact four times the number of rows achieved without compression.

How compression works and how to use it

Depending on which release of DB2 you currently have, there are a number of ways to compress a table:

- Manually run the reorganization command against tables that have compression enabled
- Automatically have the tables compressed as they reach a threshold size

To compress a table in DB2 there are two steps involved. The first is to make the table eligible for compression. Simply create the table or alter the table and specify COMPRESS YES. This tells DB2 that you want to use compression algorithms on this table.

The second step is to build a dictionary of repeating values that exist within the table. To perform this task, simply run

```
REORG TABLE table_name
```

The reorg will create the dictionary of common values (if one does not exist), and then compress those values out of the data as the table is being reorganized. Future reorgs give you the option to keep the dictionary or rebuild a new one if data patterns change. From this point onward, any insert, update, import or load of data into the table will respect the dictionary information and compress those rows as well.

DB2 also has a compression estimation utility. If you are using DB2 9 and you want to know how much compression will benefit a given table you can use the INSPECT command. Running the following command will provide compression estimates for the given table:

```
INSPECT ROWCOMPESTIMATE TABLE NAME table_name  
RESULTS KEEP file_name
```

In a test of this feature, the compression estimate for two SAP tables was 75% for one and 74% for the second. After compressing the two tables the actual savings were 75.4% and 74.5% respectively.

Estimating Compression using DB2 Compression Estimation Tool

The most accurate way of determining the amount of compression in a table is to use the INSPECT command within DB2 9 or the ADMIN_GET_INDEX_COMPRESS_INFO and ADMIN_GET_TAB_COMPRESS_INFO_V97 table functions. However, this technique requires that the data be transferred or migrated to a DB2 9/9.5/9.7 database. If this is not possible, then the DB2 Compression Estimate Tool may be a useful alternative.

The DB2 Compression Estimation Tool (“Estimator”) determines the potential compression of a table or index by looking at the DDL of the object, the current statistics on the object, and then using heuristics based on extensive compression testing to determine what the potential compression could be for the object.

The tool does not read any production data within a table or index. The estimate is based solely on the statistics available to the tool. The DB2 compression algorithm can compress across (and within) columns while the Estimator can only base its results on individual column compression values. For this reason, the results determined by the tool will tend to underestimate the actual compression that can be achieved.

Requirements

The DB2 Compression tool is distributed as a single executable file called `compressestimate.exe`. The requirements for this program are:

- Installation on a Windows platform (Windows XP, Windows Vista, Windows 7)
- Administrator privileges on Windows to allow the installation of the program
- DB2 CLI client code
- Access to a DB2 system

The actual database that is being analyzed can be on any of the supported DB2 platforms. The user connecting to the DB2 database must have:

- Connect capability to the database (V8+)
- Access to SYSCAT tables

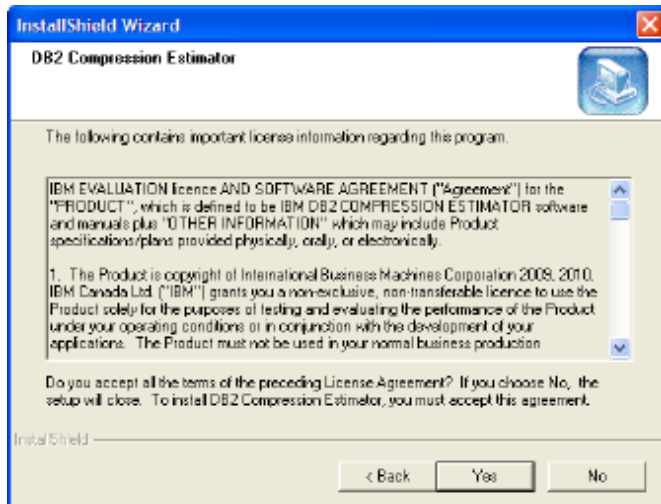
This program does not update or read any user data in the database. The only access the program has is to the SYSCAT tables described above.

Installation

To install the tool, the executable file needs to be run. The installation program will bring up a number of panels that will determine where the program will be installed. The first panel is just an introductory screen:

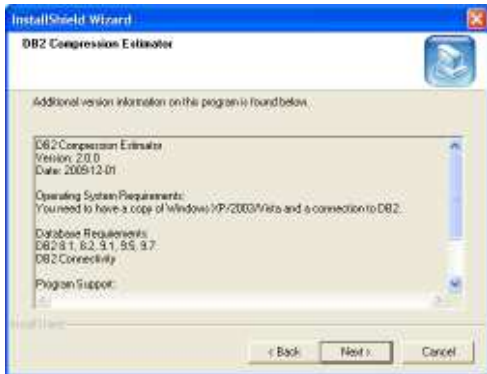


After pressing the NEXT button, the disclaimer for the program is displayed:



You must select YES in order for the installation to continue.

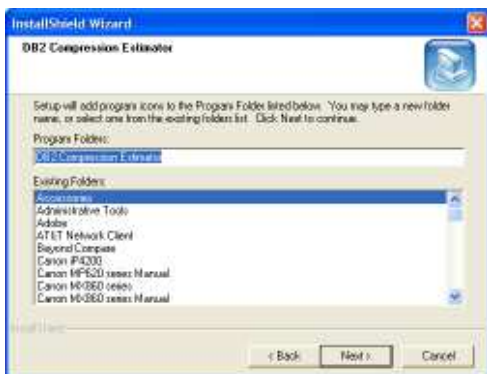
Once the YES button is pressed, information on the version and any additional requirements are shown:




After pressing the NEXT button, the path for the program is displayed. By default, the path is in the standard Program Files directory, under the DB2 Compression Estimate folder:



You can modify this path to any valid location on your system. Once the path has been defined and the NEXT statement pressed, the Program Folder is shown (DB2 Compression Estimator is the default).



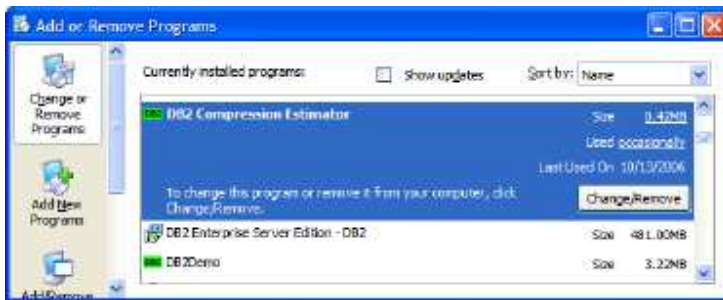
After pressing the NEXT button, the program installation continues until the final screen is displayed:




The program is now installed on the workstation.

Removing the Program

The program can be uninstalled by either running the `compressionestimate.exe` again, or using the Add or Remove Program option in the System Folder:



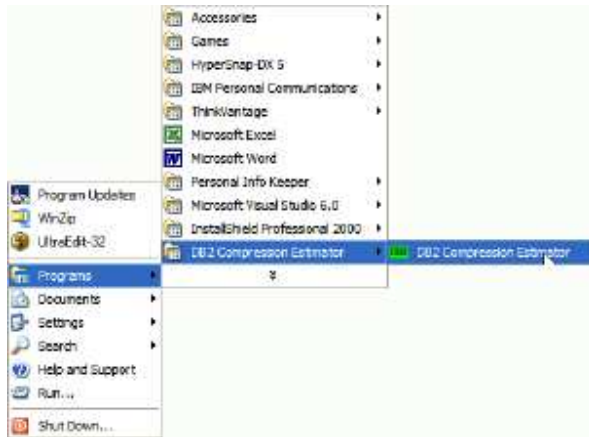
The uninstaller will give you the option to keep it or remove it:



If you have already installed the program and do a subsequent installation, the program will remove itself.

Running the Program

To run the program, select the DB2 Compression Estimator icon in the DB2 Compression Estimator folder:

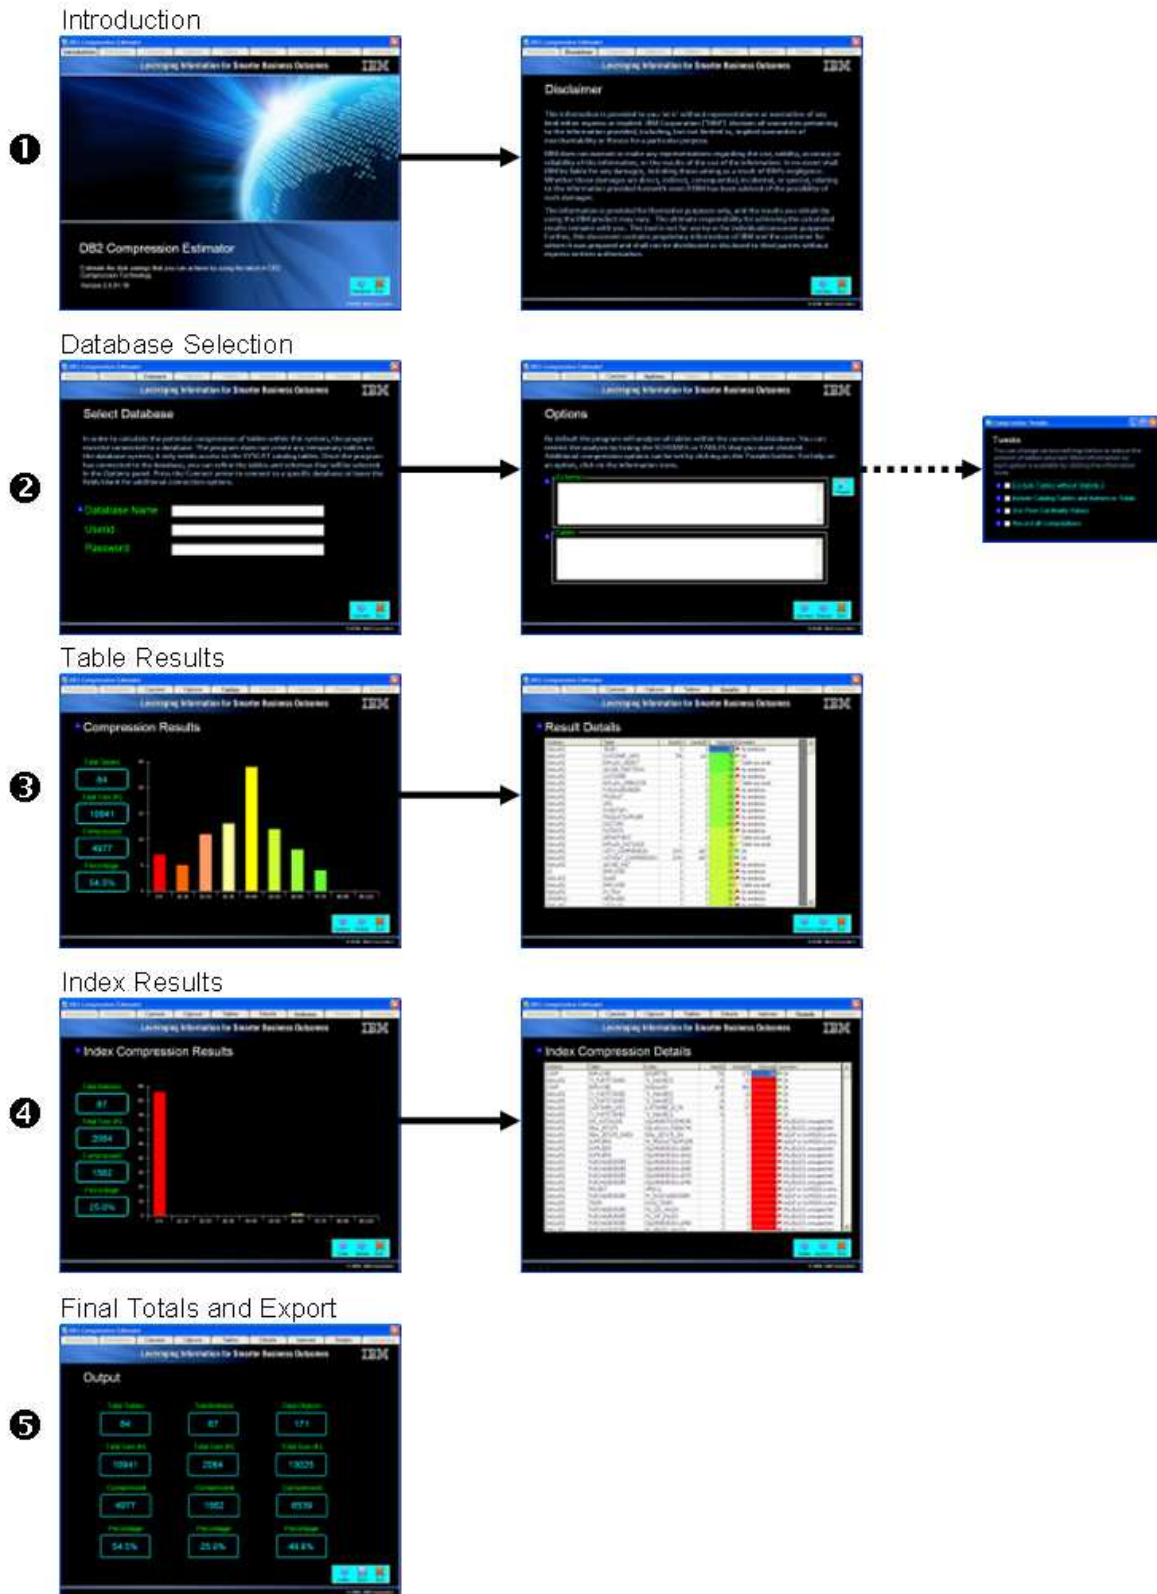


There is also an icon on the desktop that can be used to start the program:



The panels and options are described below.

Program Flow



The program will take you through five distinct steps in the estimation process. These steps include:

1. Introduction and Disclaimer

These two panels are shown whenever the program is started. You must page through these two panels before starting any of the compression analysis. Once you have read the disclaimer you will not be able to go back to view it during the session.

2. Database Selection

The two panels that follow the disclaimer allow you to enter the name of the database directly, or select it using the DB2 CLI/ODBC selection applet. Once a successful connection is made to the database, you can refine your analysis by entering certain schemas or tables that you want to analyze. An optional step is available (tweaks) that lets you modify some of the assumptions that the program makes.

You can always go back to a previous panel by using the BACK found in the bottom right-hand corner of any panel. Pressing the forward arrow will cause the program to start its analysis.

3. Table Results

A graph that displays the numbers of tables analyzed, and their compression ratio, will dynamically update as the program analyzes your database. Once the analysis is complete you can click the forward arrow to see a detailed view of all of the tables.

4. Index Results

Pressing the forward arrow on the table details screen will cause the program to analyze the indexes in the database. Note that the indexes that are analyzed are based on the same criteria as the tables. A graph similar to that for tables is shown during index analysis. Pressing the forward key on the index graph will then display more details on the indexes that were analyzed.

5. Final Totals and Export

Pressing the final forward key on the index graph will display the compression summary for tables and indexes. You have the option on this screen to export the results to a file for further analysis.

If you need to go back to a previous panel, you can always click on the back arrows found on any panel, or use the tabs that are found at the top of the screen. Note that some tabs will be grayed out if they are currently not available or the analysis has not been completed.

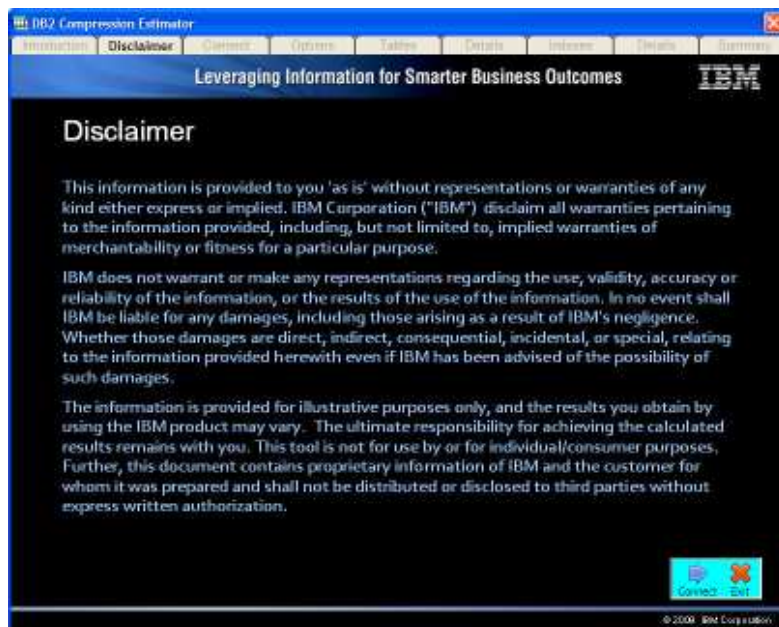
The next section will take you into more details on each one of these panels.

Initial Panel



The first panel gives you the current version of the program. Clicking the Disclaimer button will take you to the legal page. The Exit button will close the program.

Disclaimer



Remember that this tool does estimates and that the values could differ significantly from what you get in production. Click on the CONNECT button to connect to a database.

Connection Panel



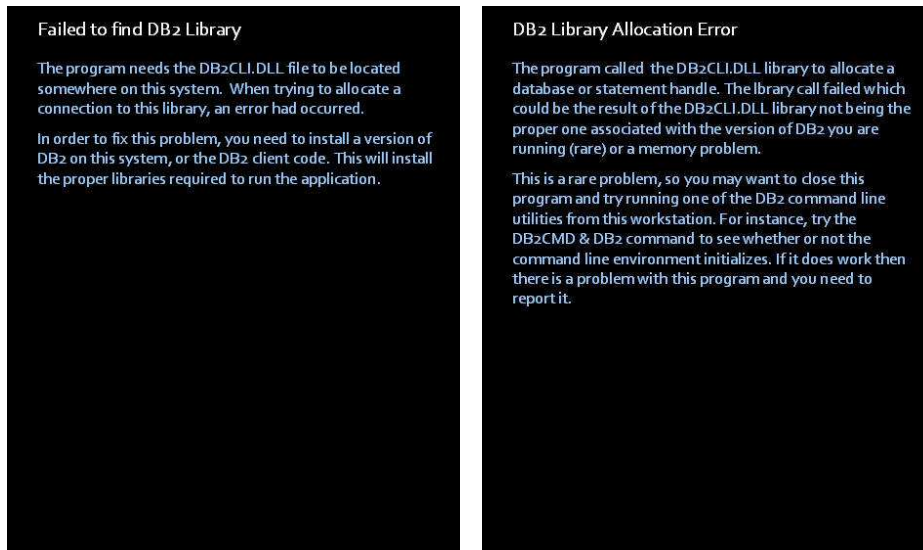
The Estimator needs to connect to a DB2 database on one of the supported platforms. There are three fields that can be filled in:

- Database – This is the name of the database that you want to analyze. The database must be catalogued and available to the client that you are connecting from.
- Userid – This is the userid that you will be using to connect to the database. If this value is left blank, it will use your local userid on Windows as the connected user.
- Password – This is the password associated with the userid. If this value is left blank, it will use the local userid information to perform the connection.

This program does not require any temporary tables in order to analyze the database. All select statements are done against the SYSCAT tables.

Pressing the Connect button will cause the program to connect to the database. If there is any error in the connection, a message will be displayed and any corrections will need to be made.

Aside from invalid passwords, userids, and database names, the only other error that you may receive is a DB2CLI library or DB2 Allocation error.




Both of these problems are related to the DB2CLI library either being unavailable or the library being at an incorrect level. This could be caused by:

- The DB2 client code not being installed on your system
- The DB2 client is from a very old version of DB2 (prior to DB2 Version 8)

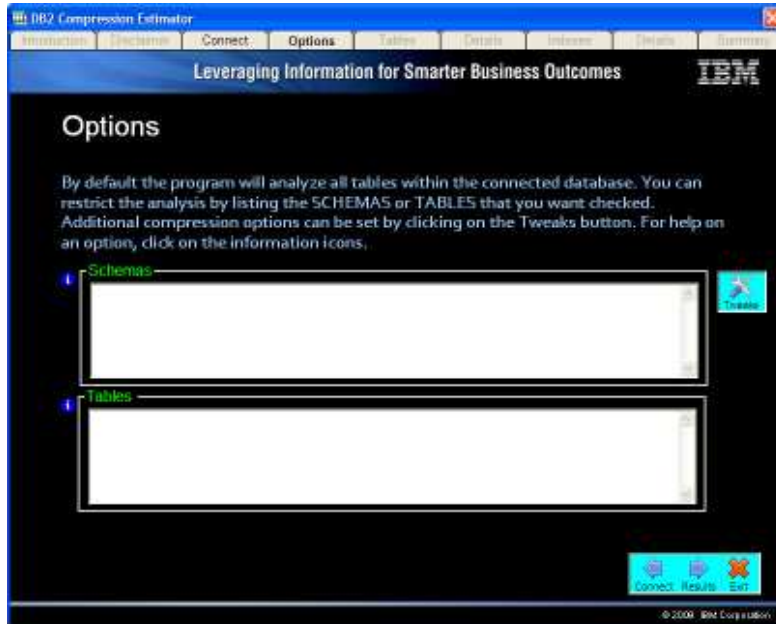
The program cannot continue with these errors since it uses the library to communicate with the DB2 database. You can test whether or not your client is working by issuing the DB2CMD command in a CMD shell on Windows, and then if successful, the DB2 command. If a command line environment for DB2 successfully starts then there may be a problem with the program.

If you left the connection fields blank, a CLI connection panel will be displayed where you can select from a database in a list.



Pressing OK will then cause the program to connect to the database you want to analyze.

Options



The options screen allows you to modify the types of tables and schemas that are analyzed. The options that are available are described below. If you have a large number of tables in your database, you may want to restrict them by SCHEMA or table name attribute before analyzing the results.

The compression estimator will list all schemas available in the database that is selected. If this results in too many schemas, you can list the schemas in this input box. Type in each schema name separated by a comma. For instance, the following list will only show the schemas SYSCAT and BAKLARZ:



A screenshot of a text input box with a black border. The word "Schemas" is written in green at the top left. The text "syscat,baklarz" is entered in the box.

Note that the schema name can include the SQL "%" character to represent any combination of characters. The following list would find any schema that starts with the letters "SYS" or ends with "CAT":



A screenshot of a text input box with a black border. The word "Schemas" is written in green at the top left. The text "sys%,%cat" is entered in the box.

Any schema name that does not include this special character will need to be exactly matched, so "syscat" by itself will only match SYSCAT.

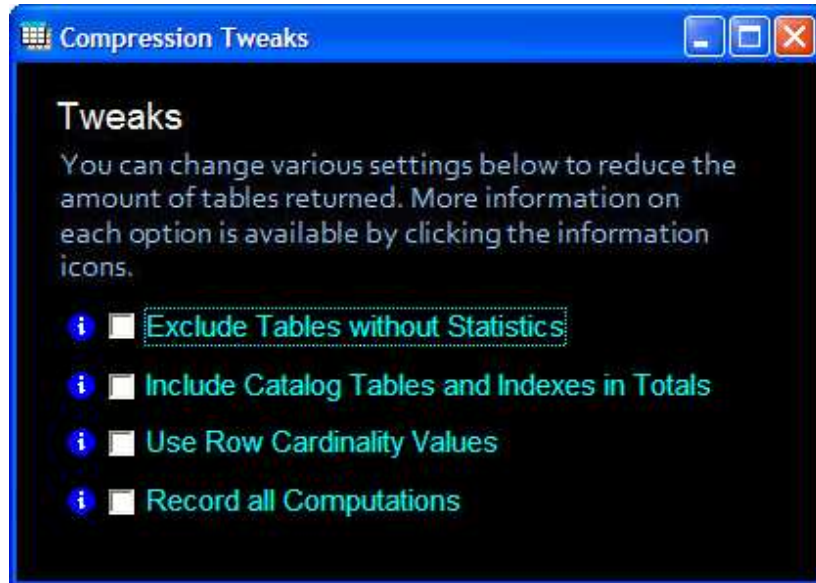
Schemas that include special characters or lowercase letters need to be placed in double quotes. The program will fold all table names to uppercase, except for those in quotes:



A screenshot of a text input box with a black border. The word "Schemas" is written in green at the top left. The text ""lowercase"" is entered in the box.

The table input box behaves in exactly the same manner as the schema box. You can limit the tables that are retrieved by using patterns. Note that the pattern is applied to all tables that are retrieved from all of the schemas. This means that if you type in "EMPLOYEE" in the table box, and various schemas in the schema box, you will get the EMPLOYEE table from all schemas that contain that table.

Tweaks



The Options panel has an additional button called "Tweaks" which allows you to change a few of the parameters that affect how the estimator will run.

- Exclude Tables without Statistics

By default, all tables are considered, regardless of statistics on the table. The compression algorithm can determine what the potential compression of a table will be, using "worst-case" assumptions. One problem with the lack of statistics is that varchar columns are assumed to be their full length (similar to character columns). This will result in exaggerated row sizes and a possible skew in the compression values.

If statistics are available for a table, the algorithm can calculate potential space savings and graph the results. Tables with no statistics will only show the potential savings and no additional information.

Selecting "Ignore tables" will result in tables without statistics being eliminated from the compression algorithm.

- Include Catalog Tables and Indexes in Totals

The estimator does not include SYSTEM catalog tables and indexes in the compression estimates. The compression algorithms that DB2 uses are not applied to the catalog tables, so no compression would be shown for these tables. The graphs produced by the program will still exclude the system tables and indexes, but the final object count and size will include them.

- Use Row Cardinality Values

The compression estimator bases its estimates on the cardinality of values on each column of a row. The actual compression within DB2 does not rely on cardinality values, but rather on patterns across the entire row. For this reason, the values

returned by the DB2 INSPECT command will be typically higher than that found by this tool.

One way to adjust the compression estimate is to change the cardinality assumption to use either the maximum cardinality of "similar" columns (row cardinality) or use the cardinality of the individual column.

The default setting uses cardinality values on each column. This method will result in a more optimistic compression estimate, and may be more realistic if you believe that patterns are independent across similar data types.

This setting estimates the compression across the entire row and will assume that columns of a similar type (i.e. all DECIMAL columns, all DATE columns) have a higher potential for repeating patterns across them.

- Record all Computations

The compression estimator can record all of the SQL it generates as well the calculations it makes for tables and indexes. This option will create a file called:

DB2COMPRESS.DBG

The DB2COMPRESS.DBG file will contain database information, status information, and details on the tables and indexes that were analyzed.

This file can be read with a text editor or with a spreadsheet program. Data is separated with tabs so it may be easier to view some of the output in a spreadsheet format. Note that if you are analyzing 1,000's of tables that the output could be quite large. Make sure you have enough space in the directory that the compression estimator was installed in.

Sample out is shown below.

```
DB2 Compression Estimator Version 2.0    2.0.1.18
[Environment]
Database          SAMPLE
User
Date              1/19/2010
Time              10:34:47 AM
Version Logic     97

[Settings]
Stats Required    NO
Cardinality       COLUMN
SYS* Tables       IGNORED

[Start SQL]
SELECT T.TABSCHEMA, T.TABNAME,
       T.NPAGES, T.CARD, C.COLNAME, C.TYPENAME, C.LENGTH, C.COLCARD,
       C.AVGCOLLEN, D.PAGESIZE
FROM SYSCAT.TABLES T, SYSCAT.COLUMNS C, SYSCAT.TABLESPACES D
WHERE
  T.TABSCHEMA = C.TABSCHEMA AND T.TABNAME = C.TABNAME
  AND (T.TYPE = 'T' OR T.TYPE = 'S')
  AND (D.TBSPACEID = T.TBSPACEID)
  AND (T.TBSPACEID >= 0)
  AND (T.TABSCHEMA NOT LIKE 'SYS%')
UNION ALL
```

```

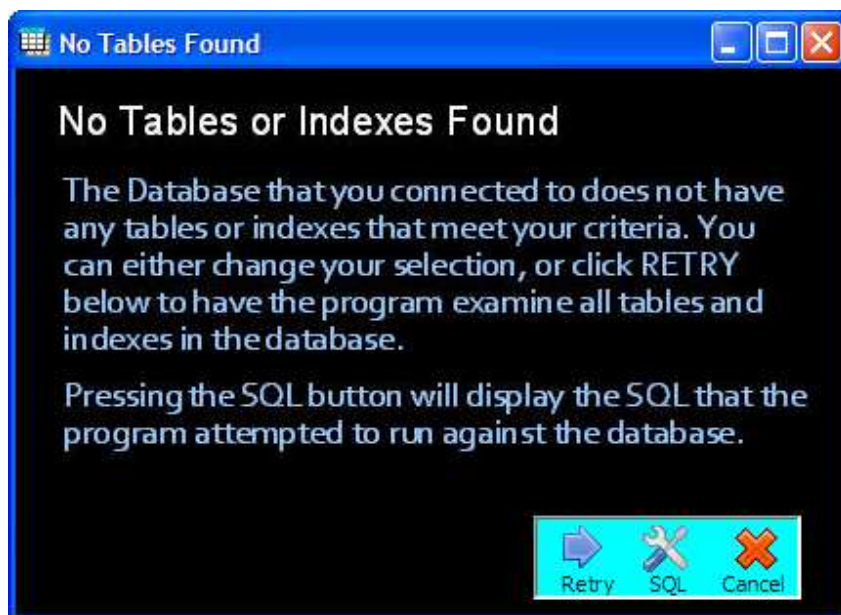
SELECT T.TABSCHEMA, T.TABNAME || '[' || P.SEQNO || ']', P.NPAGES,
P.CARD, C.COLNAME, C.TYPENAME, C.LENGTH, C.COLCARD, C.AVGCOLLEN, D.PAGESIZE
FROM SYSCAT.TABLES T, SYSCAT.COLUMNS C, SYSCAT.TABLESPACES D,
SYSCAT.DATAPARTITIONS P
WHERE
T.TABSCHEMA = C.TABSCHEMA AND T.TABNAME = C.TABNAME
AND (T.TYPE = 'T' OR T.TYPE = 'S')
AND (P.TABSCHEMA = T.TABSCHEMA AND P.TABNAME = T.TABNAME)
AND (D.TBSPACEID = P.TBSPACEID)
AND (T.TBSPACEID < 0)
AND (T.TABSCHEMA NOT LIKE 'SYS%')
ORDER BY 1, 2
FOR READ ONLY WITH UR
[End SQL]

```

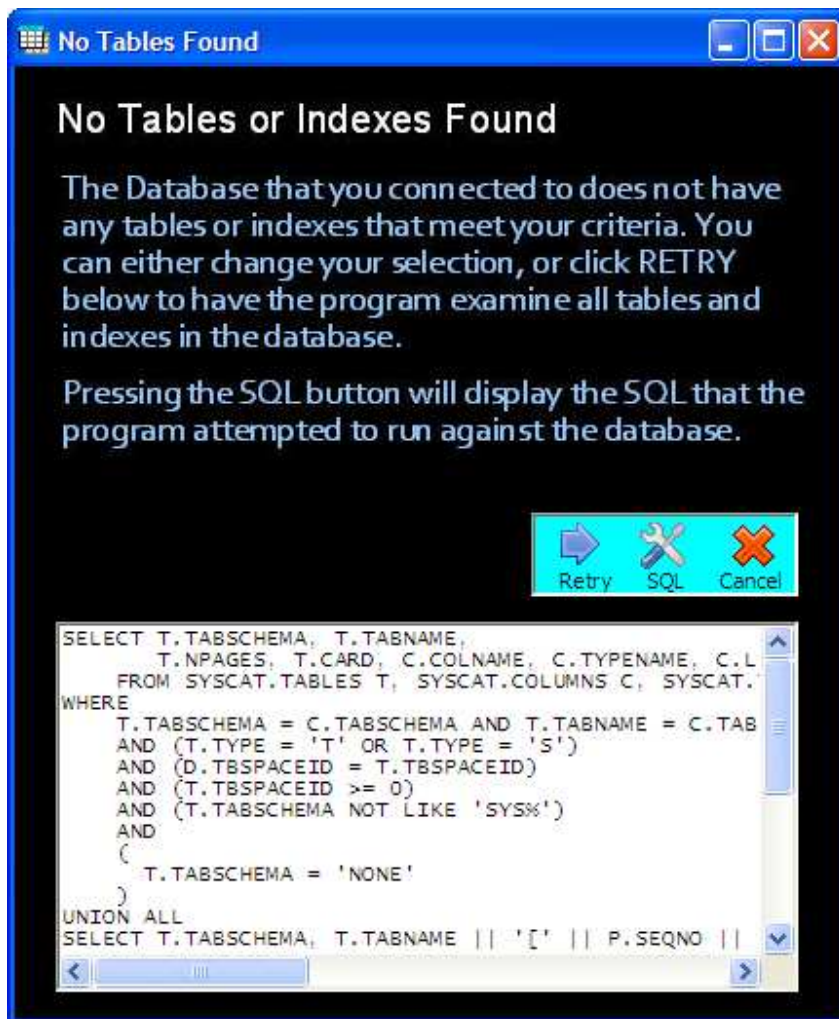
When you have completed updating your options, you can close the Tweaks window and press the RESULTS button.

Table Results

In the event that your table or schema selection does not result in any tables, you will see the following panel displayed.




Pressing the RETRY button will have the program examine all tables in the current database and ignore your original settings. Pressing CANCEL will return you to the TABLE and SCHEMA selection screen. In addition there is the option of pressing the SQL button. This button will display the SQL that the program attempted to run. This will give you the chance to copy the SQL and try it manually using a different SQL interface to see if any results are returned. This is sometimes useful when debugging connection issues.



You can return to the option panel at any time to modify the tables that are returned. If any of the options are modified, the results panel will be automatically updated when you return to that panel. If no changes are made, the results do not get updated. This is important if you want to go to the options panel to change the delimiter option which only affects the output, not the actual tables retrieved.

Support

For questions regarding the DB2 Compression Estimator, please contact George Baklarz at baklarz@ca.ibm.com.



© Copyright IBM Corporation 2010
IBM Canada
8200 Warden Avenue
Markham, ON
L6G 1C7
Canada

Printed in United States of America
10-2010
All Rights Reserved.

IBM, DB2, DB2 Universal Database, OS/390, z/OS, S/390, and the ebusiness logo are trademarks of the International Business Machines Corporation in the United States, other countries or both.
UNIX and Unix-based trademarks and logos are trademarks or registered trademarks of The Open Group. Intel and Intel-based trademarks and logos are trademarks or registered trademarks of Intel Corp. Other company, product or service names may be the trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

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 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.

The information in this white paper is provided AS IS without warranty.
Such information was obtained from publicly available sources, is current as of 10/01/2006, and is subject to change. Any performance data included in the paper was obtained in the specific operating environment and is provided as an illustration. Performance in other operating environments may vary. More specific information about the capabilities of products described should be obtained from the suppliers of those products.