Syncsort Optimize DB2
(EZ-DB2)

Stats User Guide
Release 9.10

Version 6.14

July 2017

Syncsort DB2 Stats User Guide
(formerly EZ-DB2 EZ-Stats User Guide Guide)

NOTICE

Syncsort DB2 Stats User Guide contains proprietary and confidential material, and is only for use by licensees of the Syncsort Optimize DB2 (EZ-DB2) proprietary software system. This publication may not be reproduced in whole or in part, in any form, without written permission from Syncsort Incorporated.

IBM, z/OS and Db2 are trademarks of International Business Machines Corporation. The names of other products or brands in this document are trademarks of their respective holders.

Syncsort Optimize DB2, formerly called EZ-DB2 is designed to be used by licensed users of IBM Db2 for z/OS, as a supplement to this product.

PREFACE:

EZ-DB2 is a program product of Cogito providing an integrated suite of *workload-centric* analysis and optimization tools for DB2 on the IBM z/OS and OS/390 compatible environments.

EZ-Stats is a component of **EZ-DB2** that allows the user to manipulate DB2 system catalog statistics. In particular, **EZ-Stats** can be used to copy statistics from one DB2 system to another; view and edit statistics, and archive and restore statistics on a given DB2 subsystem.

RELATED DOCUMENTATION:

EZ-DB2 Installation Reference
EZ-DB2 Getting Started and IVP Guide
EZ-DB2 Tracer & Cache User Guide
EZ-DB2 SQL Warehouse User Guide
EZ-DB2 Reports Guide
EZ-DB2 Commands Reference
EZ-DB2 Alerts User Guide
EZ-DB2 Impact Analyzer User Guide
EZ-DB2 Index Analyzer User Guide
EZ-DB2 Index Optimizer (EZ-XOP) User Guide
EZ-DB2 Stats User Guide (this document)
EZ-DB2 Housekeeping and Controls User Guide
EZ-DB2 Tables Reference

TABLE OF CONTENTS

FIGURES

New Features

Enhanced DB2 V9 Statistics Captured

DB2 Version 9 introduced some additional statistics including DATAREPEATFACTORF in SYSINDEXES and HISTOGRAM statistics in SYSCOLDIST. These enhanced statistics are now captured when extracting statistics to the EZ-DB2 statistics repository, and are included when copying statistics between DB2 subsystems or updating a DB2 subsystem catalog from the statistics repository.

The following figure shows an example of the SYSCOLDIST histogram statistics. Refer to the *DB2 V9 Performance and Tuning Guide* for further information about the HISTOGRAM statistics and DATAREPEATFACTF.

Note that another enhancement when displaying the SYSCOLDIST data is the SHOW option. Instead of displaying all of the SYSCOLDIST data types together in one display you can now select which SYSCOLDIST TYPE you wish to see, namely:-

- C – Cardinality
- F – Frequent value
- H – Histogram Statistics
- N – Non-padded frequent value

The following figure shows TYPE H (Histogram statistics).

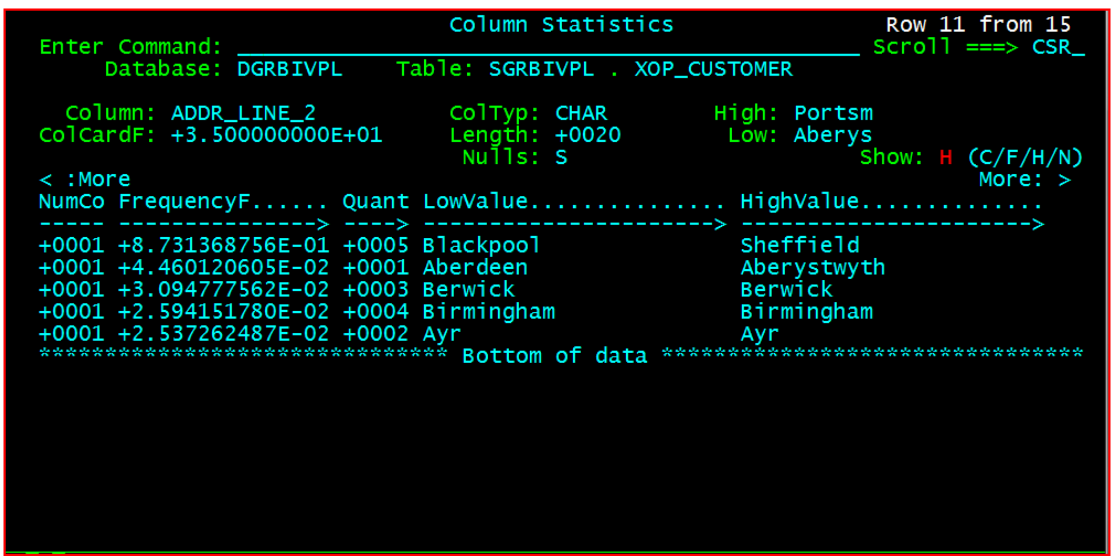

Figure 1 SYSCOLDIST Histogram Statistics

Select all databases

In order to make it easier to refresh ALL of the statistics in the EZ-DB2 statistics repository, a new option has been added to EZ-Stats option 5.5 to allow the user to select ALL of the databases.

See Select All Databases in the following figure.

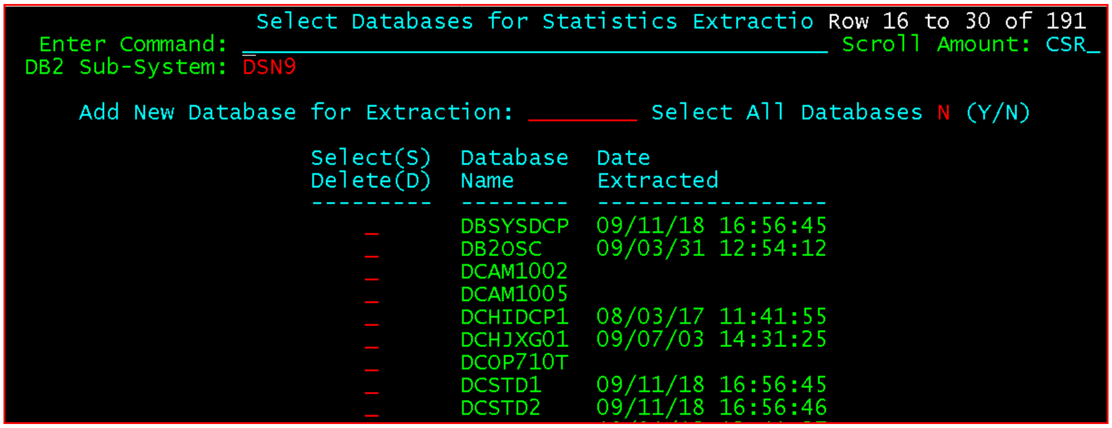

Figure 2 Add New database to Statistics Repository

Index and table space name translation

When copying statistics between subsystems, EZ-Stats may need to translate the names of indexes and/or table spaces. This may be necessary, for example, if the corresponding indexes have been created with different names in a test system compared to the names used in the production system. This is typically the case in SAP environments, where the index and table space names are automatically generated during the install process.

In prior releases, name translation is activated by specifying:-

Is this an SAP Database(Y/N): Y

Effective with APAR CX910123, name translation will automatically occur if you specify:-

Ignore Missing Objects(Y/N): Y

Irrespective of whether this is a SAP database.

Select Multiple Table Names and Table Space names

When selecting a single database for update, you can now drill on the database name to display a list of table names or table space names for the applicable database. You are then able to select one or more tables or table spaces to be updated.

You can filter the list of tables by creator or name using the wildcard character '%'

Refer to Section Selecting by Tablespace and Table for more information.

Update Database DSNDB06

It is now possible to select database DSNDB06 to be copied or updated from the Stats repository. Previously, update of DSNDB06 was disallowed. If you select DSNDB06 you will see a Pop-Up with a warning message.

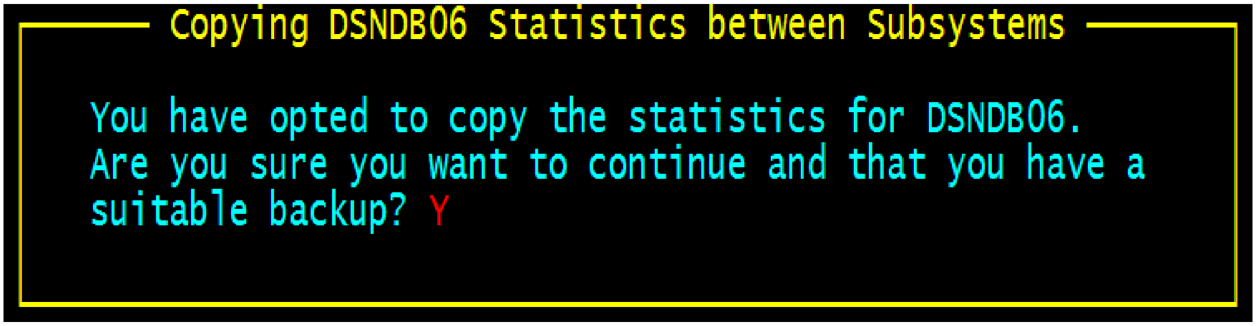

Enter Y if you wish to proceed with the update.

Process additional Statistics

EZ-Stats now includes the following additional statistics when extracting, updating and copying statistics:-

- SYSCOLSTATS
- SYSKEYTARGETS
- SYSKEYTGTDIST
- SYSROUTINES

EZ-Stats now includes ALL of the statistics that are used for access path selection as defined in the IBM Manual “Managing Performance” (SC19-2978-06).

Refer to Appendix “Statistics used for access path selection” for further details.

Enhanced filtering options and performance improvements

When filtering on selection panels, for example to reduce the number of database names displayed, you can use the characters “%” and “_” (underscore).

% is used to match all trailing characters. For example, ABC% will match on any character string that’s starts with the characters ABC.

_ is used to match on any character in that position. For example, AB_DE will match on any character in the third character position, such as AB1DE, AB2DE etc.

It is possible to combine both filtering options, such as AB_CD%.

Also, on panels which permit filtering options, the process has been changed to list no database names when the panel is first invoked. This is for performance reasons as previously the panel would by default list ALL names, which may have resulted in a delay before the initial list was displayed. Now the user must press <**enter**> if no filters are required, or specify the filters and then press <**enter**> before the list of names is displayed.

Furthermore, the process will only read as many names as required to fill the first screen, before the screen is displayed.

Multiple Archive Versions

It is now possible to retain multiple versions of archived catalog statistics through the use of Archive Statistic PROFILES, For example, you may wish to retain statistics relating to specific events, such as - Before RUNSTATS on a particular date, or corresponding to the BIND/REBIND of particular applications.

Refer to sections 5.2 - Archive Current Catalog Statistics, 5.3 - Display Catalog Statistics in Archive and 5.4 - Restore Catalog Statistics from Archive for further information.

Sort on Date extracted

User can now sort on date extracted when displaying database names currently in statistics warehouse. Enter A or D in the command line and press <enter> with the cursor in the Date Extracted column to sort ascending or descending respectively.

Default sort order is descending.

Generate Statistics Using RUNSTATS REPORT ONLY

User can now generate statistics from RUNSTATS REPORTS ONLY. This may be useful if it is required to copy statistics from a subsystem where it is not permissible or desirable to run RUNSTATS and physically update the catalog statistics. The user can now run RUNSTATS REPORT ONLY in that environment, and use EZ-Stats to build the statistics dataset to be used to update the target system.

Refer to section ?? for further information.

Generate RUNSTATS parameters from Statistics Feedback

Effective with DB2 11, the DB2 optimizer provides information about missing or conflicting statistics. Refer to the IBM DB2 11 Managing Performance Guide for further information.

EZ-Stats now includes an option to generate the global RUNSTATS parameters based upon the information in the SYSIBM.SYSSTATFEEDBACK table.

Refer to section ?? for further information.

EZ-STATS

5 EZ-Stats

- Model Production systems by Copying Catalog Statistics from Production to QA and Test environments, with automatic conversion of names and qualifiers where applicable.
- Tweak Catalog Statistics to influence access path selection.
- Archive current DB2 Statistics before making any changes and Restore them as needed.

It is not always possible or desirable to Explain SQL against the Production system. However, if a DBA were to create a test database AClone@ using a copy of the production DDL, but without the production Catalog Statistics then, when the SQL is explained against this AClone@, the access paths selected will be meaningless. This is because, when DB2 evaluates the SQL without production Catalog Statistics being present, DB2 will use default statistics, which will bear no relationship to the live environment at all. By migrating Catalog Statistics from a Live system to the Test system, a DBA is able to create an Optimizer Identical Clone of a production system, so that DBAs could do Awhat if@ type work on SQL and get the same access paths as on production.

There are other factors that mean that access paths may still not be identical even when using the same statistics in a Test environment as Production. Refer to AModelling your Production System@ section A.2 for further details.

EZ-STATS allows you to copy the DB2 catalog statistics ***used for access path selection*** between DB2 systems, allowing you to Explain SQL in a test environment and generate the same access path as used in production.

EZ-STATS has the following features:-

- **Copies** Statistics from one subsystem to another for one or more databases at a time.
- **Extracts** Statistics from any DB2 subsystem to the **Statistics Warehouse** for one or more databases at a time
- **Display** and **Edit** Statistics in the Statistics Warehouse for a particular system and database
- **Updates** statistics from the Statistics Warehouse to a new DB2 System or database
  - DBNAMEs and table owner/creators may be different on Live and Test systems.
  - **EZ-STATS** will automatically translate DBNAME and entity Owner / Creator names as required when transferring Statistics from one system to another.
  - In some environments (e.g. SAP R/3Z), Index names may vary for a given table in different environments. **EZ-STATS** will automatically translate the index names if

necessary, recognizing corresponding indexes and updating the Statistics for the appropriate index.

- You may Load ALL Statistics for a given database, or Select Statistics to Load for particular Table Space(s), Table(s), Indexes etc.

- **Archive** and **Restore** Statistics on a particular system

The following tables contain statistics that are used for access path selection and are processed by EZ-Stats when Copying, Extracting, Archiving or Updating statistics:-

- SYSCOLDIST
- SYSCOLSTATS
- SYSCOLUMNS
- SYSINDEXES
- SYSKEYTARGETS
- SYSKEYTGTDIST
- SYSROUTINES
- SYSTABLES
- SYSTABLESPACE
- SYSTABSTATS

Refer to Statistics used by Optimizer for Access Path Selection section *A.2* for further details.

Select option 5 to display the EZ-Stats menu as shown in the following figure:-

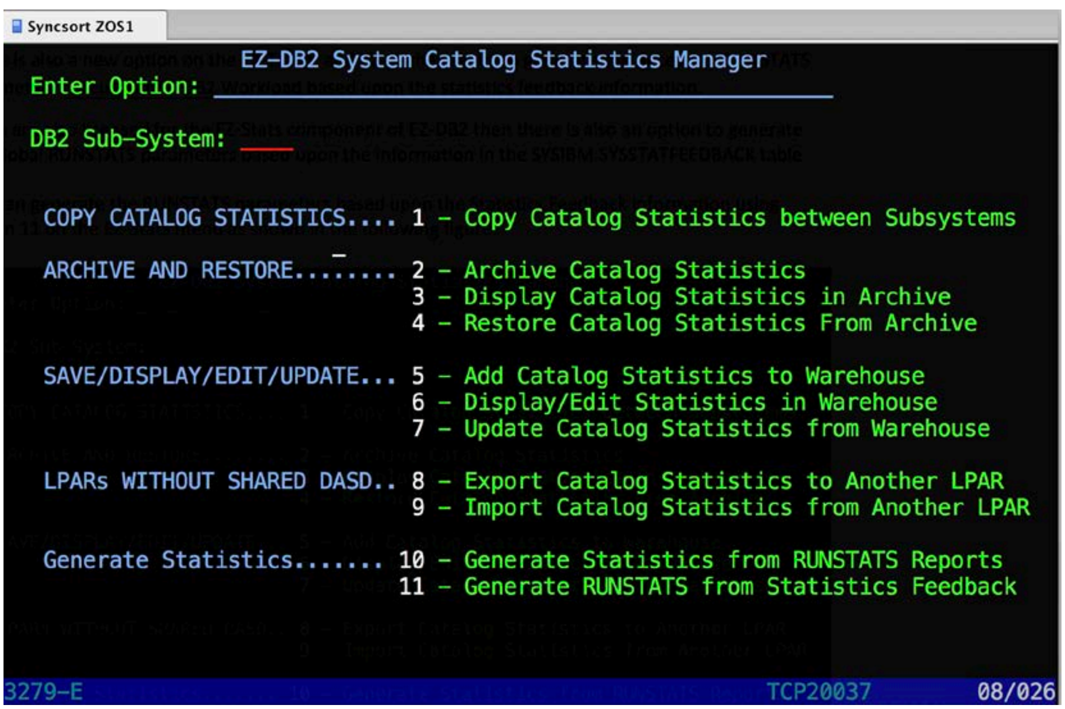

Figure 3. Statistics Manager Menu

These options are described in the following sections.

5.1 Copy Catalog Statistics between Subsystems

This option is used when you wish to copy statistics directly from one DB2 subsystem to another without first saving the statistics in the statistics warehouse.

Select option 5.1 Copy Catalog Statistics between Subsystems to display the following panel:-

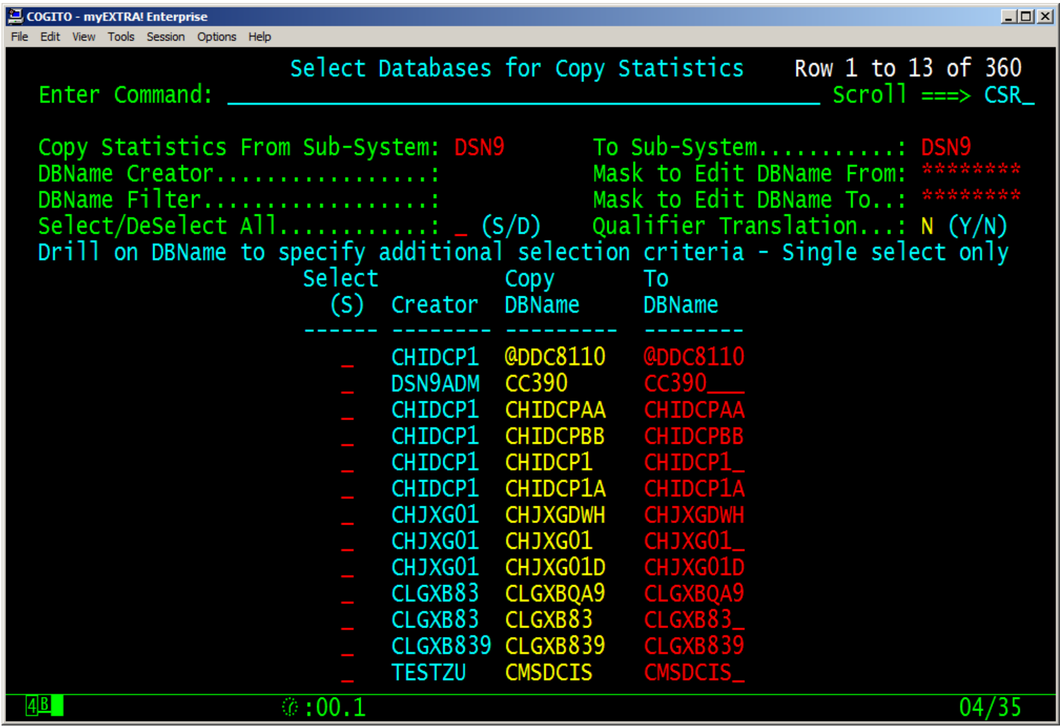

Figure 4 **Select Database for Copy Statistics**

All databases defined on the FROM system are listed unless you have filtered the display using the Dbname Creator and / or DBName filter parameters. See DBName Creator and DBName Filter parameters below:-

The following parameters may be specified on this panel.

Copy Statistics From Sub-System	Specify the DB2 subsystem which you wish to copy the statistics FROM.
DBName Creator	Optionally enter a Creator Name and only Databases created with that Creator will be listed for selection. You may filter the selection using "_"and/or "%" characters. For example, AB_D%
DBName Filter	Optionally enter a filter to reduce the Database name display to those that match the filter. You may filter the selection using "_"and/or "%" characters. For example, AB_D%

Select/Deselect ALL

Enter S or D to Select or Deselect all Database names displayed.

To Sub-System

Specify the DB2 subsystem which you wish to copy the statistics TO.

Mask to Edit DBName From

If the database names on the system to be Copied FROM are different to the database names to be copied TO, you may use this parameter and Mask to Edit DBName To in order to perform a global edit of the FROM and TO DBNames.

For example,

Mask to Edit DBName From is **PR****
Mask to Edit DBName To **QA****

will change all values PR in positions 3 and 4 to QA.

See also To DBName described later.

Mask to Edit DBName To

See Mask to Edit DBName From.

Qualifier Translation

If the Databases that will have their statistics updated were created using different qualifiers to the Databases from which the statistics were extracted then enter **Y** and DRILL <**PF4**> to get a pop-up panel in which you can input the old and new values for object creators as shown in the following figure:-

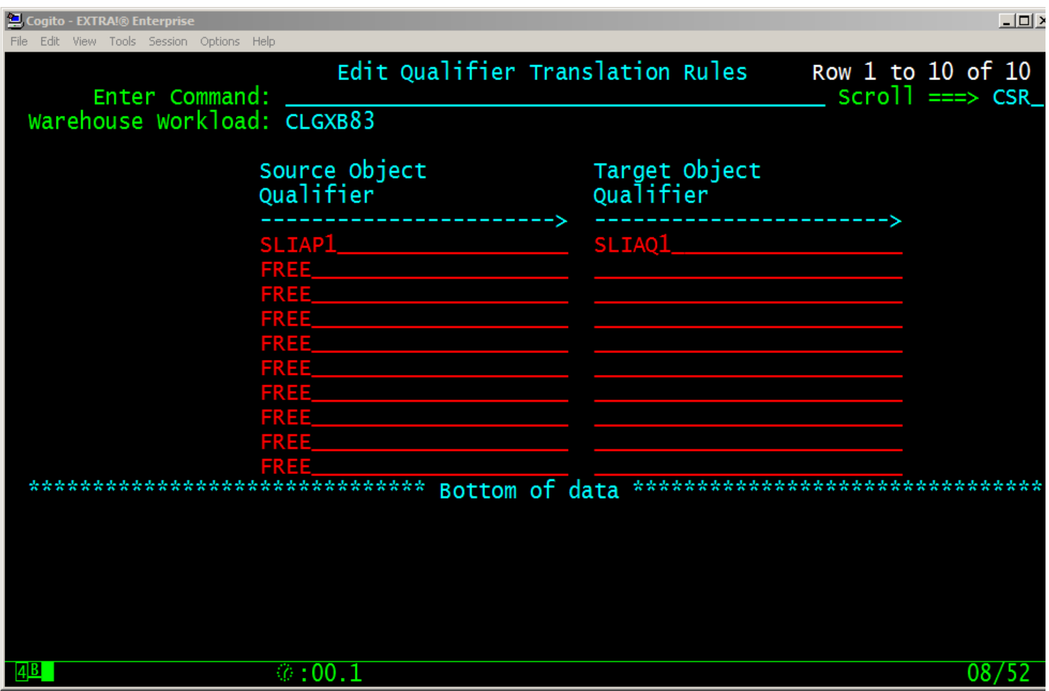

Figure 5 **Qualifier Translation**

The qualifier translation panel will initially display a list of 10 FREE entries into which you can overtype the values for the Source and Target object qualifiers as required.

If you need more than 10 entries, press <**enter**> and an additional 10 free slots will be displayed. Repeat until you have defined all of the qualifier translation requirements for the job.

Select (S)

Enter S to select this database to be copied.

See also Select/DeSelect ALL parameter.

Creator

The database creator name.

Copy DBName

The Name of a database to be copied from the DB2 subsystem specified.

All of the databases on this subsystem will be listed, unless you have filtered by DBName or Creator. See DBName Creator and DBName Filter parameters above.

By default, ALL Tablespaces and Tables will be copied for the selected database(s). However, when selecting a single database, it is possible to DRILL on the database name to further specify the selection by either table space or table name. Multiple table spaces or tables may be specified for the selected database. See Selecting by table space and table.

To DBName

The name of the database to be copied TO. You may overwrite this name if the name is different to the originating database that it is copied FROM.

See also Mask to Edit DBName From/To above.

*Note that you can SORT on DBName or Creator and also use the L(ocate) command for a database name. For example, **L DGRB** to position at the first database that commences with the characters DGRB.*

Selecting by Tablespace and Table

When selecting a single database, it is possible to further refine the selection criteria by DRILLing <PF4> on a database name. The following panel is displayed:-

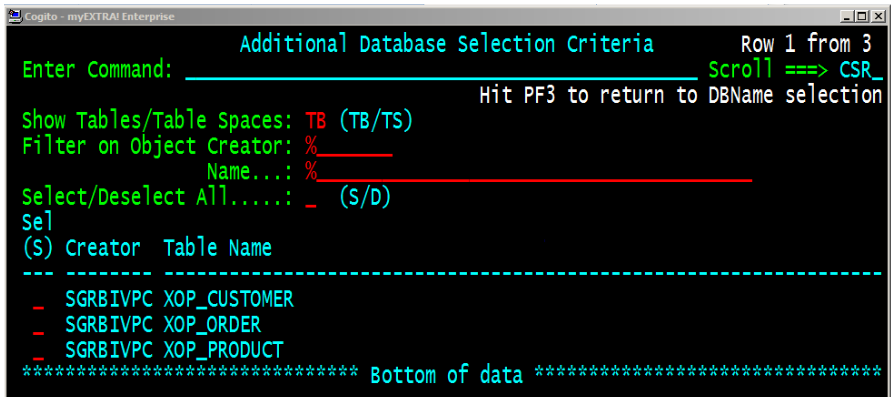

Figure 6 Additional database selection criteria

The panel lists all of the table spaces or all of the tables for the selected database. You may select one or more tables, OR one or more table spaces depending on the Show Tables/Tables Spaces parameter.

The following parameters may be specified on this panel:-

Show Tables/Table spaces (TB/TS)

The panel will display the table names or the table space names for the selected database. The following options may be specified:-

TB List the table names for the selected database.

TS List the table space names for the selected database.

The default value is TB.

All of the table or table space names for the selected database will be listed, unless you have filtered by Object Creator or Name.

See Filter on Object Creator/ Name.

Filter on Object Creator

Optionally enter a Creator Name and only objects created with that Creator will be listed for selection. You may use * or % as a wild card.

For example enter ABC% to list all objects where the creator name begins with ABC.

Filter on Object Name

Optionally enter a Name and only objects with that Name will be listed for selection. You may use * or % as a wild card.

For example enter ABC% to list all objects where the name begins with ABC.

Select/Deselect ALL	Enter S or D to Select or Deselect all Database names displayed.
Select(S) **Creator**	Enter S to Select this object. The Table or Table space Creator will be listed.
Table / Table Space Name	Either the Table name or Table Space Names will be listed depending on whether you have specified TB or TS. See Show Tables/Table spaces. All Tables or Table spaces for the selected database will be shown unless you have filtered by Creator or Name. See Filter on Object Creator and Filter on Object Name.

Unless you select by Table name or Tablespace name, by default EZ-Stats will copy all relevant statistics for the selected database(s). That is,

- SYSCOLDIST
- SYSCOLSTATS
- SYSCOLUMNS
- SYSINDEXES
- SYSKEYTARGETS
- SYSKEYTGTDIST
- SYSTABLES
- SYSTABLESPACE
- SYSTABSTATS

To copy the SYSROUTINES statistics, you must enter a database name of SYSRTNES. SYSROUTINES statistics are not associated with any particular database and are copied separately from database copies. If you select SYSRTNES as a pseudo database name, then ONLY the SYSROUTINES statistics will be copied.

You can override this selection using Select Options for Statistics Update panel shown below.

If you filter by Tablespace for a selected database, then EZ-Stats will copy statistics as above, except only for the selected Tablespace(s).

You can override this selection using Select Options for Statistics Update panel shown below.

If you filter by Table name for a selected database then by default EZ-Stats will copy the following statistics for the selected table(s):-

- SYSCOLDIST
- SYSCOLSTATS
- SYSCOLUMNS
- SYSINDEXES
- SYSKEYTARGETS
- SYSKEYTGTDIST
- SYSTABLES

You can override this selection using Select Options for Statistics Update panel shown below, however, you cannot copy the SYSTABLESPACE or SYSTABSTATS statistics. If you want to copy these statistics, you must select the entire database, or selected Tablespace(s).

Select Options for Statistics Update

When you have identified the Sub-Systems and Databases (and optionally additional database selection criteria), press <**ENTER**> and a further selection panel will be presented to allow filtering of the Statistics to be updated as shown in the following figure:-

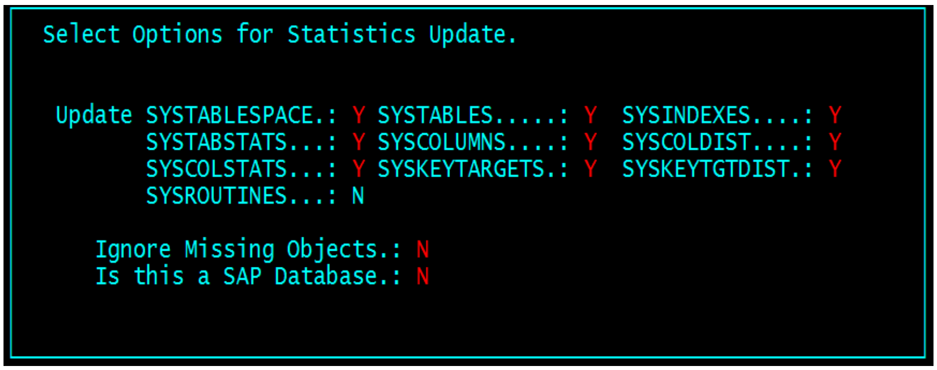

Figure 7 Select Options for Statistics Update

The following fields are displayed:-

Update SYSTABLESPACE Specify Y to update the SYSTABLESPACE Statistics for the selected database(s) or Tablespace(s).

This option may only be selected if you are updating/ copying the entire database or selected Tablespace(s).

It may not be selected if you are filtering by Table names.

Update SYSTABLES Specify Y to update the SYSTABLES Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSINDEXES Specify Y to update the SYSINDEXES Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSTABSTATS Specify Y to update the SYSTABSTATS Statistics for the selected database(s) or Tablespace(s).

This option may only be selected if you are updating/ copying the entire database or selected Tablespace(s).

It may not be selected if you are filtering by Table names.

Update SYSCOLUMNS Specify Y to update the SYSCOLUMNS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSCOLDIST

Specify Y to update the SYSCOLDIST Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSCOLSTATS

Specify Y to update the SYSCOLSTATS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSKEYTARGETS

Specify Y to update the SYSKEYTARGETS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSKEYTGTDIST

Specify Y to update the SYSKEYGTDIST Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSROUTINES

Specify Y to update the SYSROUTINES statistics. This option may only be selected if you have specified a database name of SYSRTNES.

All of the other options will be set to N.

Ignore Missing Objects(Y/N)?

When applying the copied statistics TO the specified sub-system, if the Database(s) against which the statistics are to be applied do not exactly match the databases from which the statistics have been copied FROM, then any Objects missing will be reported.

For example a Tablespace may be defined as Partitioned in the FROM subsystem but not in the TO subsystem, which may be acceptable. However you should use this option with caution as missing Tables may cause SQL to fail to Explain and missing Indexes will cause different access paths to be generated.

Enter Y to ignore missing Objects and continue.
Enter N to terminate with RC 16.

Is this a SAP database (Y/N)?

Specify Yes or No to identify whether the Database to be updated is an **SAP** Database. **EZ-STATS** has some special processing logic to cater for name translation in **SAP** environments, where typically Tablespace names and even Index names vary between systems. The default value is N. Refer to the Section A.5.4 SAP R/37 Processing Considerations for further details.

Specify any required filter parameters and press <**Enter**> to generate the job stream to copy the statistics from one DB2 system to another.

- Type **SUBMIT** in the command line to submit the job for execution.
- Press <**PF3**> to return to the **EZ-Stats** menu.

To save a repeatable job stream that can be run at will, take copies of the files qualified with "userid.IO" to new dataset names of your choice and adjust the saved JCL with the new dataset names.

You may view the output from the Statistics Update job (XOPSTATU) in SDSF.

Refer to section A.5 in the Appendix for further information about XOPSTATU processing.

5.2 Archive Current Catalog Statistics

EZ-Stats provides the ability to save current catalog statistics to an archive file in the Statistics Warehouse. This may be useful if, for example, you wish to replace statistics for a particular database, but then have the option of reverting to the original statistics.

On a test/QA DB2 subsystem you could for example:-

- Archive Current Catalog statistics for a database
- Update Current Catalog Statistics for the database from Statistics Warehouse (for example, using statistics from the Production system)
- Perform your Binds, Explains Testing, Optimization etc. etc. using the production statistics
- Restore Statistics for the database from Archive

With **EZ-Stats**, it is possible to create multiple archive versions through the use of archive PROFILES. Each PROFILE may contain the statistics for one or more databases, which may have been added to the archive at the same or different dates. For example, you may wish to create archives of statistics corresponding to certain events, such as - Before Runstats, or corresponding to the BIND/REBIND of particular applications.

It is possible to view the statistics in a particular archive PROFILE and restore statistics to the DB2 catalog from a particular archive PROFILE.

The default archive PROFILE is called PRIMARY, and this is the PROFILE used by EZ-DB2 Warehouse processes. For example, when you specify that you wish to **Use Live Statistics** and **Preserve the Current Statistics** when EXPLAINing a Workload in the Warehouse environment. EZ-DB2 will:-

- Save the current Statistics to the PRIMARY archive
- Update the Catalog using the LIVE statistics
- EXPLAIN the SQL in the Workload
- Restore the Statistics from the PRIMARY archive

It is recommended that you do not use the PRIMARY archive for your own statistics archives. You should create your own PROFILE as described later.

Refer to the *EZ-DB2 SQL Warehouse User Guide* for further details about the option to **Use Live Statistics** and **Preserve Current Statistics**.

The statistics related to the PRIMARY statistics archive profile are saved in a flat file with a name of *ezdb2hlq.db2ssid*.STATS.ARC.

The statistics related to a *user created* archive PROFILE are saved in a flat file with a name of *ezdb2hlq.db2ssid*.STATS.*profile*

Where

ezdb2hlq	is the High Level qualifier used for the EZ-DB2 Installation datasets
db2ssid	is the applicable DB2 subsystem ID or data sharing group name
profile	is the name of the ARCHIVE PROFILE

To save statistics to an ARCHIVE, Select option 5.2 to Archive Current Catalog Statistics. The following panel is displayed:-

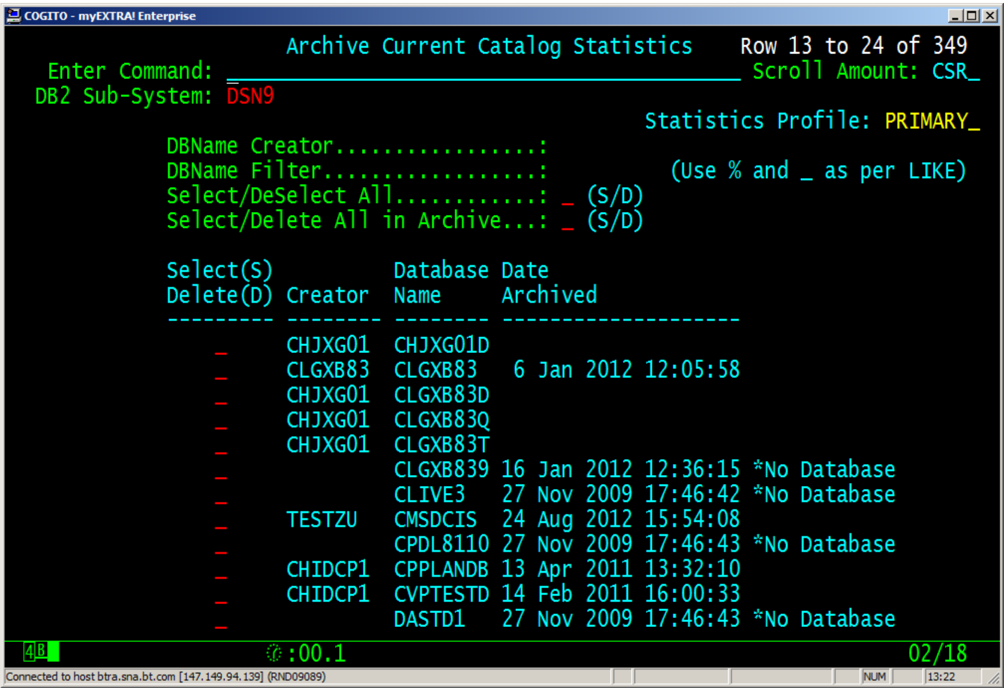

Figure 8 Archive Current Catalog Statistics

The panel lists all of the databases for the selected DB2 subsystem (unless you have used the Creator and / or DBName filters See DBName Creator / DBName Filter parameters below). Databases that have previously been extracted to the applicable archive will show a Date Archived; otherwise the Date Archived will be blank. If the Date Archived is followed by a comment *No Database, this indicates that the statistics for the relevant database are in the archive, but the corresponding database no longer exists. The following parameters may be specified on this panel:-

DB2 Subsystem — Specify the DB2 subsystem from which you wish to add statistics to the Archive.

Statistics Profile — Enter Profile name to be used for this Archive. If you enter a new name, i.e. for a Profile that does not already exist you will see the following panel:-

Figure 9 Create new Archive PROFILE

You may overtype the description (shown in red) to specify your own meaningful description for the new PROFILE. Press <**enter**> to continue to create the new PROFILE, or press <**PF12**> to cancel.

If you don't wish to create a new PROFILE you can specify the name of an existing PROFILE to be used for this archive. Alternatively, you can DRILL <**PF4**> on the PROFILE field to display a list of all of the existing PROFILEs as shown in the following figure:-

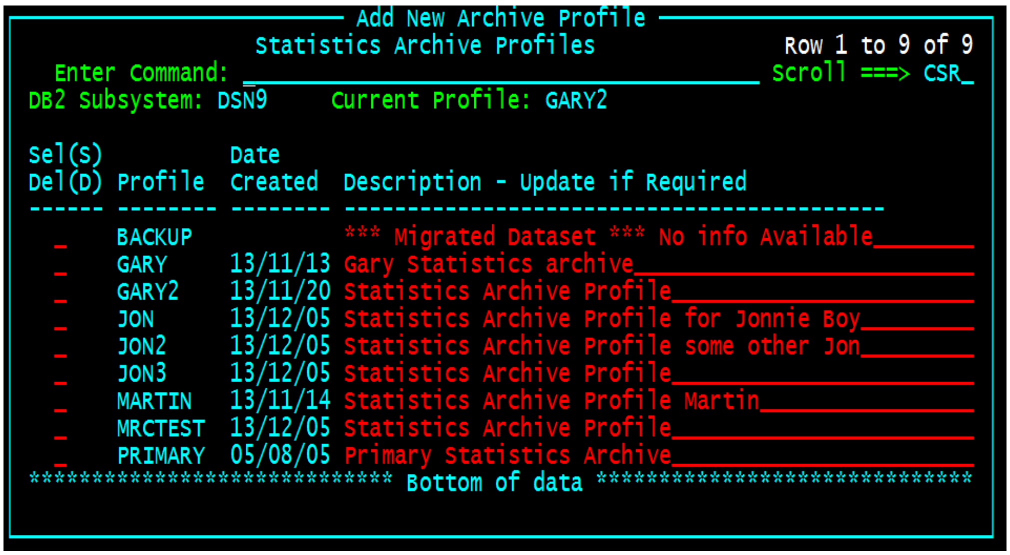

Figure 10 Display Archive PROFILES

Sel(S)
DelD)

Enter S to select the archive PROFILE.

Enter D to delete the applicable PROFILE. The dataset containing the statistics for this profile will be deleted.

Profile

The name of this Profile.

Date Created

The date this PROFILE was created,

Description

The description given to this PROFILE at the time of creation. You can overtype to update the description if required.

Referring back to the remaining parameters on the Archive Catalog Statistics panel (Figure 8)

DBName Creator

Optionally enter a Creator Name and only Databases created with that Creator will be listed for selection.

You may filter the selection using "_"and/or "%" characters. For example, AB_D%

DBName Filter

Optionally enter a filter to reduce the Database name display to those that match the filter.

You may filter the selection using "_"and/or "%" characters. For example, AB_D%

NOTE: To add SYSROUTINES statistics to the Archive for the relevant subsystem, specify a database name of SYSRTNES (SYSRTNES is not actually a database name!)

Select/Deselect ALL (S/D)

Enter S to Select ALL listed databases.

The Select indicator S is set for all listed databases.

Enter D to De-select all selected databases.

The Select indicator is cleared for all of the databases that currently have the Select indicator set.

Select/Delete ALL in Archive.

Enter S to Select ALL databases currently in the Archive PROFILE.

This option may be useful if you wish to refresh the statistics for all of the databases currently held in the archive for this PROFILE.

The Select indicator S is set for all databases currently in the archive profile.

Enter D to DELETE all databases currently in the archive PROFILE. The Delete (D) indicator is set for ALL databases currently in the archive for this PROFILE.

Select (S)
Delete (D)

Enter S to Select a particular database that you wish to add to the archive. If the database is already in the archive, the statistics will be refreshed from the DB2 catalog.

You may select multiple databases.

See also the Select/DeSelect ALL parameter described above,

Enter D to DELETE the statistics for a database from the archive.

You may DELETE multiple databases in one execution.

See also the Select/Delete ALL in Archive parameter described above.

Creator

The database creator name.

Database Name

The Database Name. Note that all databases for the applicable DB2 subsystem will be listed unless you have filtered by Creator or Database name, See DBName Creator and DBName Filter parameters.

*Note that you can SORT on DBName or Creator and also use the L(ocate) command for a database name. For example, **L DGRB** to position at the first database that commences with the characters DGRB.*

Date Archived

The date and time that the statistics for this database were added to the archive.

This is NOT the same as STATSTIME, which indicates when RUNSTATS was last run.

If Date Archived is blank, then the statistics for this database are not currently held in the archive.

If the Date Archived is followed by the comment *No Database", this indicates that there are statistics in the archive for a database that no longer exists. It is recommended that you delete such databases from the archive.

- Press **<Enter>** to generate the job stream to perform the extraction of statistics from the catalog to the archive.
- Type SUBMIT in the command line to submit the job.

After the job has completed, you will see the Date Archived data updated in the display.

You may view the output from the Statistics Extract job (XOPSTATD) in SDSF.

Refer to section A.4 in the Appendix for further information about XOPSTATD processing.

5.3 Display Catalog Statistics in Archive

You may display statistics that have been saved to a statistics archive in the same way as you can display statistics added to the Statistics Warehouse.

Note: When displaying Statistics in an Archive that all statistics fields are protected. You may not edit statistics that have been saved in a Statistics archive.

Select option 5.3 to display the Display Catalog Statistics in Archive menu as shown in the following figure:-

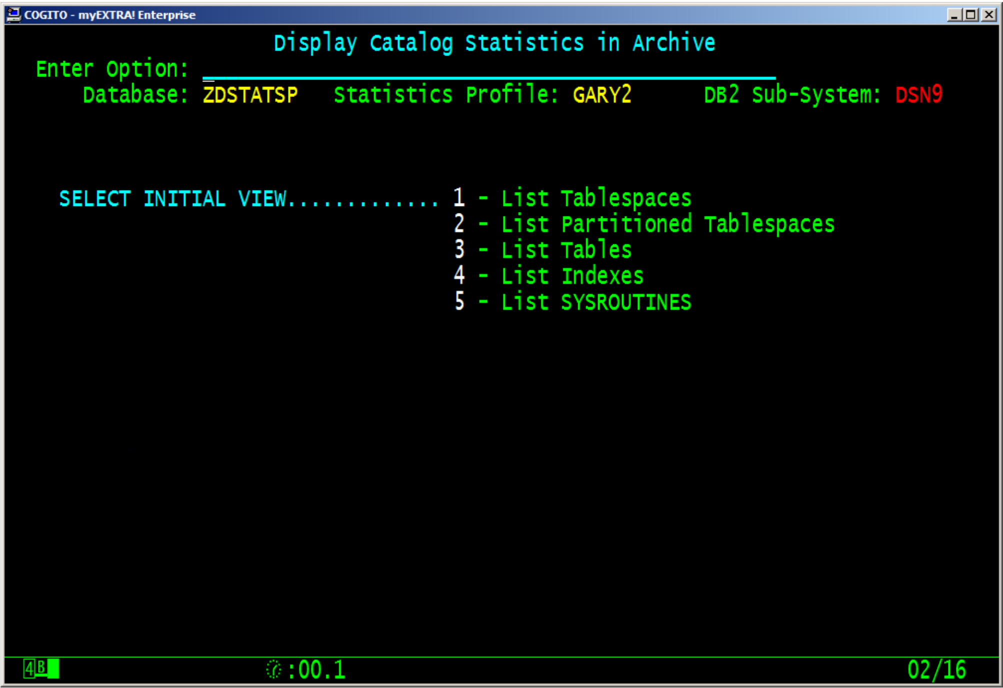

Figure 11 **Display Catalog Statistics in Archive**

Statistics Profile	Enter the name of the Archive PROFILE to be displayed. You can DRILL on this field to display a list of available PROFILEs as described in section 5.2. See Figure 10.

The Display Catalog Statistics in Archive function is essentially the same as the Display/Edit Statistics in Warehouse function (option 5.6), except that you cannot edit the statistics.

Refer to section 5.6 for further information.

5.4 Restore Catalog Statistics from Archive

Select option 5.4 to display the Restore Catalog Statistics from Archive panel as shown in the following figure:-

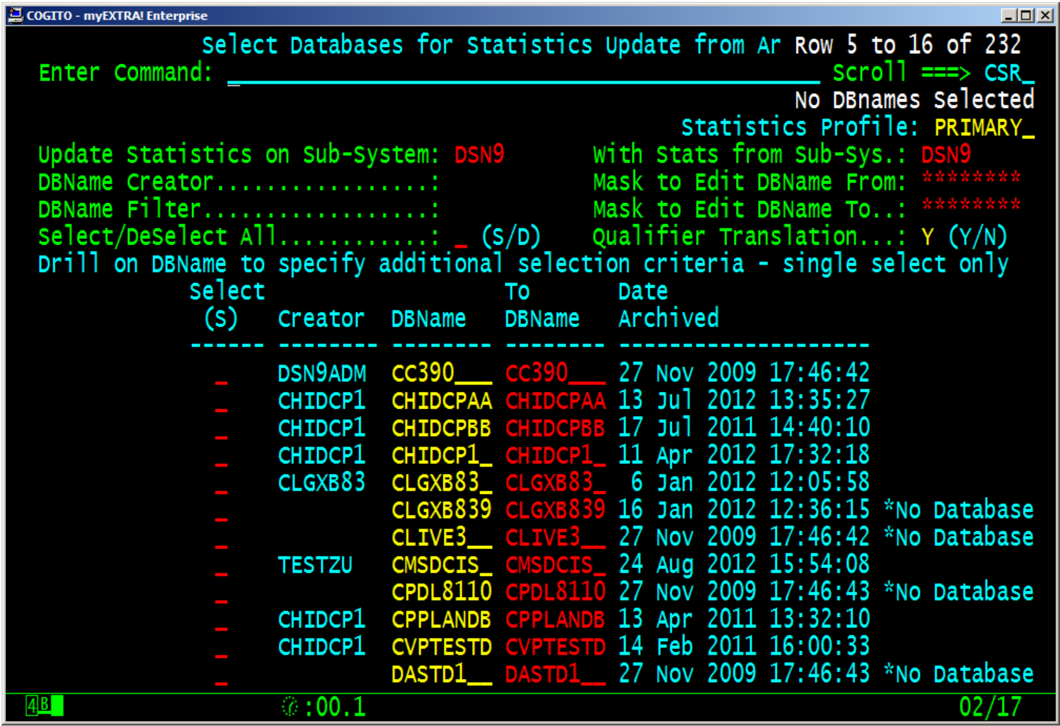

Figure 12 Restore Catalog Statistics from Archive

Statistics Profile

Enter the name of the Archive PROFILE to be restored.

You can DRILL on this field to display a list of available PROFILEs as described in section 5.2. See Figure 10.

The Restore Catalog Statistics from Archive function is essentially the same as the Update Catalog Statistics from Warehouse function, except that the statistics are restored from the specified archive PROFILE.

Refer to section 5.7 for further details.

5.5 Add Catalog Statistics to Warehouse

Select option 5.5 Add Catalog Statistics to Warehouse to select a database to be added to the Statistics Warehouse. The following panel is displayed:-

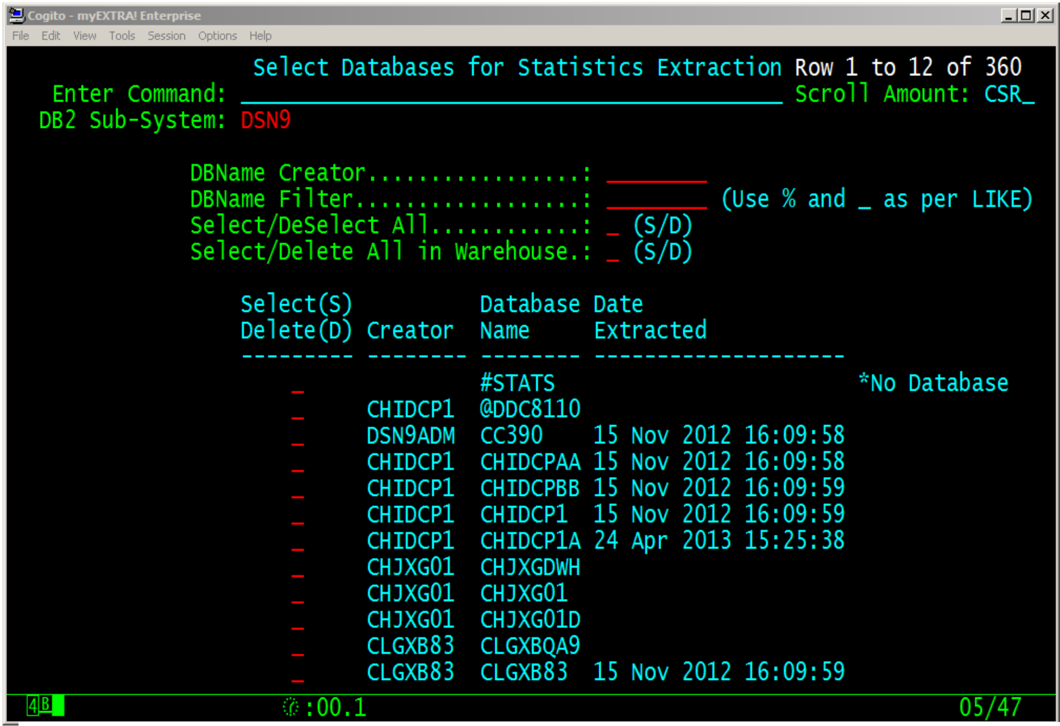

Figure 13 Select database for statistics extraction

The panel lists all of the databases for the selected DB2 subsystem (unless you have used the Creator and / or DBName filters See DBName Creator / DBName Filter parameters below). Databases that have previously been extracted to the statistics warehouse will show a Date Extracted; otherwise the date extracted will be blank. If the Date Extracted is followed by a comment *No Database, this indicates that the statistics for the relevant database are in the statistics warehouse, but the corresponding database no longer exists.

The following parameters may be specified on this panel:-

DB2 Subsystem	Specify the DB2 subsystem from which you wish to add statistics to the Statistics warehouse.
DBName Creator	Optionally enter a Creator Name and only Databases created with that Creator will be listed for selection. You may filter the selection using "_"and/or "%" characters. For example, AB_D%
DBName Filter	Optionally enter a filter to reduce the Database name display to those that match the filter.

You may filter the selection using "_"and/or "%" characters. For example, AB_D%

NOTE: To add SYSROUTINES statistics to the Warehouse for the relevant subsystem, specify a database name of SYSRTNES (SYSRTNES is not actually a database name!)

Select/Deselect ALL (S/D)

Enter S to Select ALL listed databases.

The Select indicator S is set for all listed databases.

Enter D to De-select all selected databases.

The Select indicator is cleared for all of the databases that currently have the Select indicator set.

Select/Delete ALL in Warehouse.

Enter S to Select ALL databases currently in the statistics warehouse.

This option may be useful if you wish to refresh the statistics for all of the databases currently held in the statistics warehouse.

The Select indicator S is set for all databases currently in the warehouse.

Enter D to DELETE all databases currently in the Warehouse. The Delete (D) indicator is set for ALL databases currently in the Warehouse.

Select (S)
Delete (D)

Enter S to Select a particular database that you wish to add to the Statistics Warehouse. If the database is already in the warehouse, the statistics will be refreshed from the DB2 catalog.

You may select multiple databases.

See also the Select/DeSelect ALL parameter described above,

Enter D to DELETE the statistics for a database from the statistics warehouse.

You may DELETE multiple databases in one execution.

See also the Select/Delete ALL in Warehouse parameter described above.

Creator

The database creator name.

Database Name

The Database Name. Note that all databases for the applicable DB2 subsystem will be listed unless you have filtered by Creator or Database name, See DBName Creator and DBName Filter parameters.

Note that you can SORT on DBName or Creator and also use the L(ocate) command for a database name. For example, ***L DGRB*** *to position at the first database that commences with the characters DGRB.*

Date Extracted

The date and time that the statistics for this database were extracted to the statistics warehouse.

This is NOT the same as STATSTIME, which indicates when RUNSTATS was last run.

If Date Extracted is blank, then the statistics for this database are not currently held in the statistics warehouse.

If the Date Extracted is followed by the comment *No Database", this indicates that there are statistics in the warehouse for a database that no longer exists. It is recommended that you delete such databases from the warehouse.

User can now sort on date extracted when displaying database names currently in statistics warehouse.

Enter A or D in the command line and press <enter> with the cursor in the Date Extracted column to sort ascending or descending respectively.

Default sort order is descending.

- Press **<Enter>** to generate the job stream to perform the extraction.
  - Type SUBMIT in the command line to submit the job.

  After the job has completed, you will see the Date Extracted data updated in the display.

You may view the output from the Statistics Extract job (XOPSTATD) in SDSF.

Refer to section A.4 in the Appendix for further information about XOPSTATD processing.

5.6 Display/Edit Statistics in Warehouse

This section describes how to Display and optionally edit statistics in the Statistics Warehouse.

Select option 5.6 to display Display/Edit Statistics in Warehouse menu as shown in the following figure:-

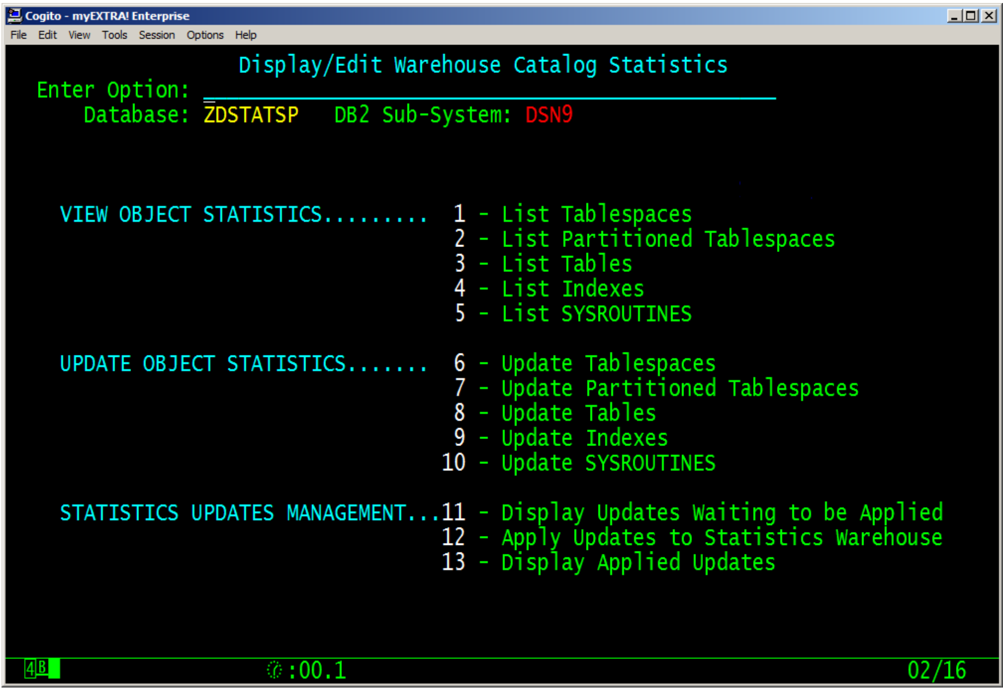

Figure 14. Display/Edit Statistics in Warehouse

From this menu you can select from the following options:-

1. List Tablespaces

Display the Tablespace Statistics (SYSTABLESPACE) for the selected Database. You may DRILL down to view the Tables/Columns/Indexes for the each Tablespace.

2. List Partitioned Tablespaces

Display the Partitioned Tablespace (SYSTABSTATS) statistics for the selected Database. You may DRILL down to view the Tables/Columns/Indexes for the each Partitioned Tablespace.

3. List Tables

Display the Table Statistics (SYSTABLES) for the selected Database. You may DRILL down to view the Columns/Indexes for the each Table.

4. List Indexes

Display the Index Statistics (SYSINDEXES) for the selected Database. You may DRILL down to view the Index Columns and Column distribution statistics for each Index.

5. List SYSROUTINES

Display the SYSROUTINES statistics. Note that this option can only be selected if the Database name is set to SYSRTNES.

6. Update Tablespaces

Edit the Tablespace Statistics (SYSTABLESPACE) for the selected Database. You may DRILL down to edit the Tables/Columns/Indexes for the each Tablespace.

7. Update Partitioned Tablespaces

Edit the Partitioned Tablespace (SYSTABSTATS) statistics for the selected Database. You may DRILL down to edit the Tables/Columns/Indexes for the each Partitioned Tablespace.

8. Update Tables

Edit the Table Statistics (SYSTABLES) for the selected Database. You may DRILL down to edit the Columns/Indexes for the each Table.

9. Update Indexes

Edit the Index Statistics (SYSINDEXES) for the selected Database. You may DRILL down to edit the Index Columns and Column distribution statistics for each Index.

10. Update SYSROUTINES

Edit the SYSROUTINES statistics. Note that this option can only be selected if the database name is set to SYSRTNES.

11. Display Updates Waiting to be

View the Old and New values of Statistics that you have modified and not yet applied to the Statistics Warehouse.

12. Apply Updates to Statistics Warehouse

If required, apply any changes you have made to the copy of the statistics in the Statistics Warehouse.

NOTE: If you exit from this panel without selecting this option, any changes you may have made are DISCARDED.

13. Display Applied Updates

Displays a list of all Statistics that have been updated previously.

Select DB2 Subsystem and database

Before displaying the database objects, you must first specify the DB2 subsystem and database name that you wish to display.

Database

You may explicitly select a database name, or you may display a list of database names by specifying a database name filter.

You may filter the selection using "_"and/or "%" characters. For example, AB_D%.

If you specify a filter, then the Select database panel will be shown as in Figure 15

Note: You may also specify a database name of SYSRTNES. In this case, you will only be able to select the options to List and Update the SYSROUTINES statistics. (SYSRTNES is not actually a database name).

DB2 Subsystem

Specify the DB2 subsystem for the database name that you wish to display or update.

If you have specified a database name filter, the following panel will be display.

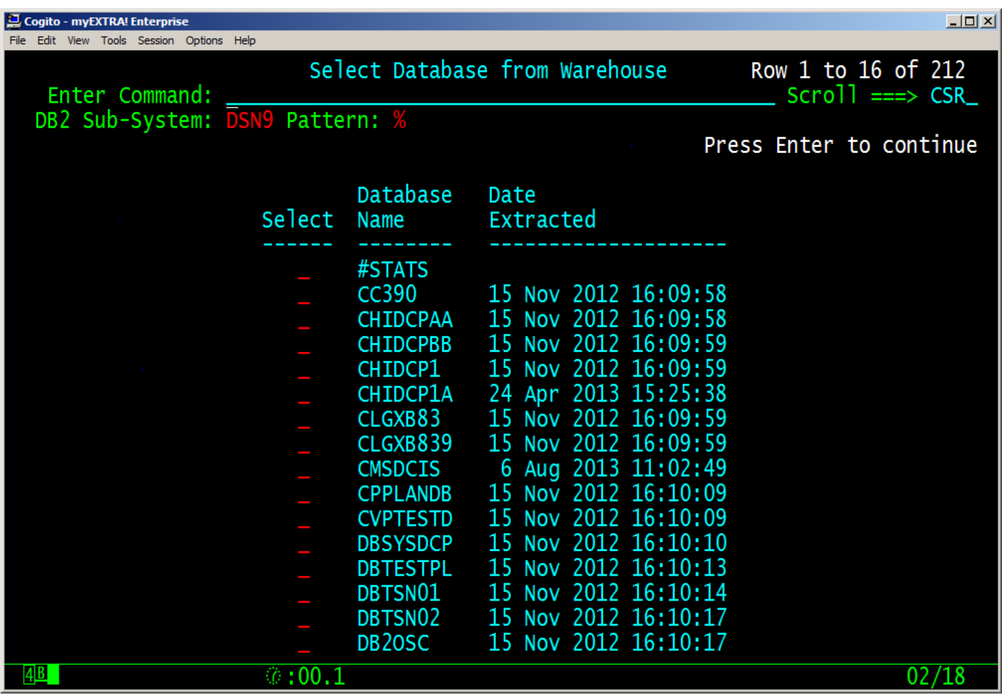

Figure 15 **Select Database from Warehouse**

The panel lists the Database names for the Statistics in the Warehouse for a particular subsystem and the date extracted.

The following fields are displayed:-

DB2 Sub-system
The Current DB2 Subsystem name. You can over type this field.

Pattern
Enter a filter to reduce the number of database names displayed.

You may filter the selection using "_"and/or "%" characters. For example, AB_D%.

Select
Type **S** in this field to select the database name of the statistics you wish to view/edit.

Database name
The name of each database extracted to the Statistics Warehouse for the current DB2 subsystem is listed.

*Note that you can SORT on DBName or Creator and also use the L(ocate) command for a database name. For example, **L DGRB** to position at the first database that commences with the characters DGRB.*

Date Extracted
The date and time when statistics for this database were last extracted.

If the Date Extracted is followed by a comment *No Database, this indicates that the statistics for the relevant database are in the statistics warehouse, but the corresponding database no longer exists.

Select the required database name you wish to View/edit and press **<Enter>** to return to the Display/Edit Statistics in Warehouse menu.

Note: If you wish to view statistics for a database not currently in the Statistics Warehouse, you must first go to option 5.5 to Add Catalog Statistics to Warehouse.

5.6.1 List Tablespaces (SYSTABLESPACE)

The SYSIBM.SYSTABLESPACE table contains one row for each table space.

Select option.5.6.1.to display List Tablespaces panel as shown in the following figure:-

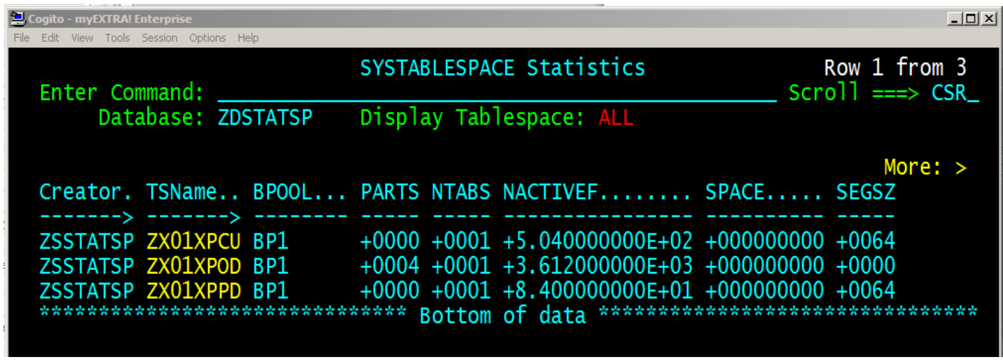

Figure 16 List Tablespaces

The display shows the SYSTABLESPACE statistics for the selected database.

Display Tablespace	**Tablespace Name**	Enter a Tablespace name in this field to display Statistics for a particular tablespace.
	ALL	Enter **ALL** in this field to display statistics for ALL tablespaces for the current database. The default value is **ALL.**

The following information is displayed:-

Creator	The name of the Tablespace creator.
TSName	The Table space name You can DRILL <PF4> on the TSName to see the Tables Detail display for the tables in this tablespace.
BPOOL	Name of the buffer pool used for the table space..
PARTS	Number of partitions of the table space; 0 if the table space is not partitioned.
NTABS	Number of tables defined in the table space.
NACTIVEF	Number of active pages in the table space. A page is termed active if it is formatted for rows,

even if it currently contains none. The value is 0 if statistics have not been gathered.

SPACE	Number of kilobytes of DASD storage allocated to the table space, as determined by the last execution of the STOSPACE utility. The value is 0 if the table space is not related to a storage group, or if STOSPACE has not been run. If the table space is partitioned, the value is the total kilobytes of DASD storage allocated to all partitions that are storage group defined.
SEGSZ	Number of pages in each segment of a segmented table space. The value is 0 if the table space is not segmented.
STATSTIME	If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01-00.00.00.000000'. If the value is '0001-01-02-00.00.00.000000', which indicates that an ALTER TABLE statement was executed to change the length of a VARCHAR column, RUNSTATS should be run to update the statistics before they are used.

Note the following applies to all EZ-Stats statistics displays:-

- Extensive field level **HELP** is available on all fields in the **EZ-Stats** statistics displays. Position the cursor on a field and press <**PF1**> function key to obtain a description of the column.
- You can **SORT** on fields where applicable. Just position the Cursor on a Column and press **<Enter>** or enter the command SORT A or SORT D in the Command Line.
- You can **DRILL** on fields where applicable, for example, DRILL <**PF4**> on a particular Tablespace name to invoke the List Tables display for a particular table space.
- Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.
- Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

5.6.2 List Partitioned Tablespaces (SYSTABSTATS)

The SYSIBM.SYSTABSTATS table contains one row for each partition of a partitioned table space. Rows in this table can be inserted, updated, and deleted.

Select option 5.6.2 to display the List Partitioned Tablespaces panel as shown in the following figure:-

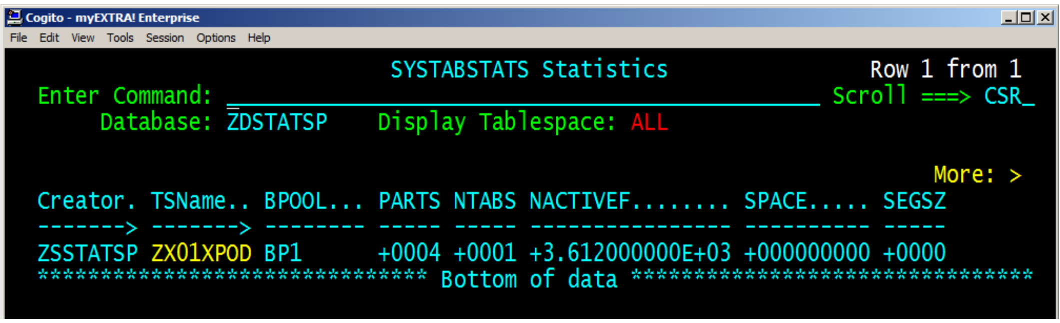

Figure 17 List Partitioned Tablespaces

The display shows the SYSTABSTATS statistics for the selected database (if applicable).

Display Tablespace:	**Tablespace name**	Enter a tablespace name in this field to display Statistics for a particular tablespace.
	ALL	Enter ALL in this field to display statistics for ALL tablespaces for the current database. The default value is **ALL**

The following information is displayed:-

Creator	The name of the Tablespace creator..
TSName	The Tablespace name You can DRILL <PF4> on the TSName to see the Partition detail display for all of the partitions in this tablespace.
BPOOL	Name of the buffer pool used for the table space..
PARTS	Number of partitions of the table space; 0 if the table space is not partitioned.
NTABS	Number of tables defined in the table space.
NACTIVEF	Number of active pages in the table space. A page is termed active if it is formatted for rows, even if it currently contains none. The value is 0 if

statistics have not been gathered.

SPACE Number of kilobytes of DASD storage allocated to the table space, as determined by the last execution of the STOSPACE utility. The value is 0 if the table space is not related to a storage group, or if STOSPACE has not been run. If the table space is partitioned, the value is the total kilobytes of DASD storage allocated to all partitions that are storage group defined.

SEGSZ Number of pages in each segment of a segmented table space. The value is 0 if the table space is not segmented.

STATSTIME If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01-00.00.00.000000'. If the value is '0001-01-02-00.00.00.000000', which indicates that an ALTER TABLE statement was executed to change the length of a VARCHAR column, RUNSTATS should be run to update the statistics before they are used.

Note the following applies to all EZ-Stats statistics displays:-

- Extensive field level **HELP** is available on all fields in the **EZ-Stats** statistics displays. Position the cursor on a field and press <**PF1**> function key to obtain a description of the column.
- You can **SORT** on fields where applicable. Just position the Cursor on a Column and press **<Enter>** or enter the command SORT A or SORT D in the Command Line.
- You can **DRILL** on fields where applicable, for example, DRILL <**PF4**> on a particular Tablespace name to invoke the List Tables display for a particular table space.
- Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.
- Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

Partitioned Tablespace detail display

DRILL on a Partitioned Tablespace to see the Partitioned Tablespace detail display as shown in the following figure:-

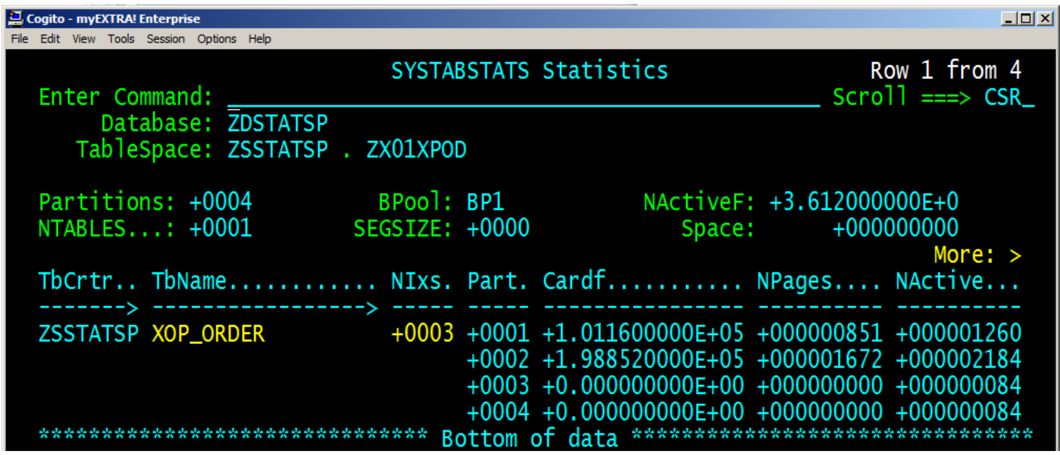

```
Cogito - myEXTRA! Enterprise
File  Edit  View  Tools  Session  Options  Help

                     SYSTABSTATS Statistics                   Row 1 from 4
 Enter Command: _______________________________________________ Scroll ===> CSR_
      Database: ZDSTATSP
    TableSpace: ZSSTATSP . ZX01XPOD

 Partitions: +0004         BPool: BP1          NActiveF: +3.612000000E+0
 NTABLES...: +0001       SEGSIZE: +0000           Space:       +000000000
                                                                  More: >
 TbCrtr.. TbName........... NIxs. Part. Cardf.......... NPages.... NActive...
 -------> -----------------> ----- ----- --------------- ---------- ----------
 ZSSTATSP XOP_ORDER          +0003 +0001 +1.011600000E+05 +000000851 +000001260
                                   +0002 +1.988520000E+05 +000001672 +000002184
                                   +0003 +0.000000000E+00 +000000000 +000000084
                                   +0004 +0.000000000E+00 +000000000 +000000084
 ****************************** Bottom of data ********************************
```

Figure 18 Partitioned Tablespace detail

The following information is displayed:-

Database	Name of the current database.
TableSpace	Name of current table space.
Partitions	Number of partitions of the table space; 0 if the table space is not partitioned.
BPool	Name of the buffer pool used for the table space
NActiveF	Number of active pages in the table space. A page is termed active if it is formatted for rows, even if it currently contains none. The value is -1 if statistics have not been gathered.
NTABLES	Number of tables defined in the table space.
SEGSIZE	Number of pages in each segment of a segmented table space. The value is 0 if the table space is not segmented.
Space	Number of kilobytes of DASD storage allocated to the table space, as determined by the last execution of the STOSPACE utility. The value is 0 if the table space is not related to a storage group, or if STOSPACE has not been run. If the table space is partitioned, the value is the total kilobytes of DASD storage allocated to all partitions that are storage group defined.
TbCrtr	Authorisation ID of the table owner.

TbName

Name of the Table.

You can DRILL <PF4> on this field to see the SYSCOLSTATS statistics display for the table.

NIXs

The number of indexes defined for the table.

You can DRILL <PF4> on this field to see the List Indexes display for the table.

Part

The Partition number.

The following information is repeated for each partition of the partitioned table space.

Cardf

Total number of rows in the table. The value is -1 if statistics have not been gathered.

NPages

Total number of pages on which rows of the table appear. The value is -1 if statistics have not been gathered.

NActive

Number of active pages in the table space. A page is termed active if it is formatted for rows, even if it currently contains none. The value is 0 if statistics have not been gathered.

PCTP

Percentage of total active pages in the partition that contain rows of the table.

PCTRC

Percentage of rows compressed within the total number of active rows in the partition. This includes any row in a table space that is defined with COMPRESS YES.

STATSTIME

If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics.

SYSCOLSTATS statistics

The SYSIBM.SYSCOLSTATS table contains partition statistics for selected columns. For each column, a row exists for each partition in the table. Rows are inserted when RUNSTATS collects either indexed column statistics or non-indexed column statistics for a partitioned table space. No row is inserted if the table space is non-partitioned. Rows in this table can be inserted, updated, and deleted.

To display the SYSCOLSTATS statistic, DRILL <**PF4**> on a table name in the List Partitioned table space detail display.

The following panel is displayed:-

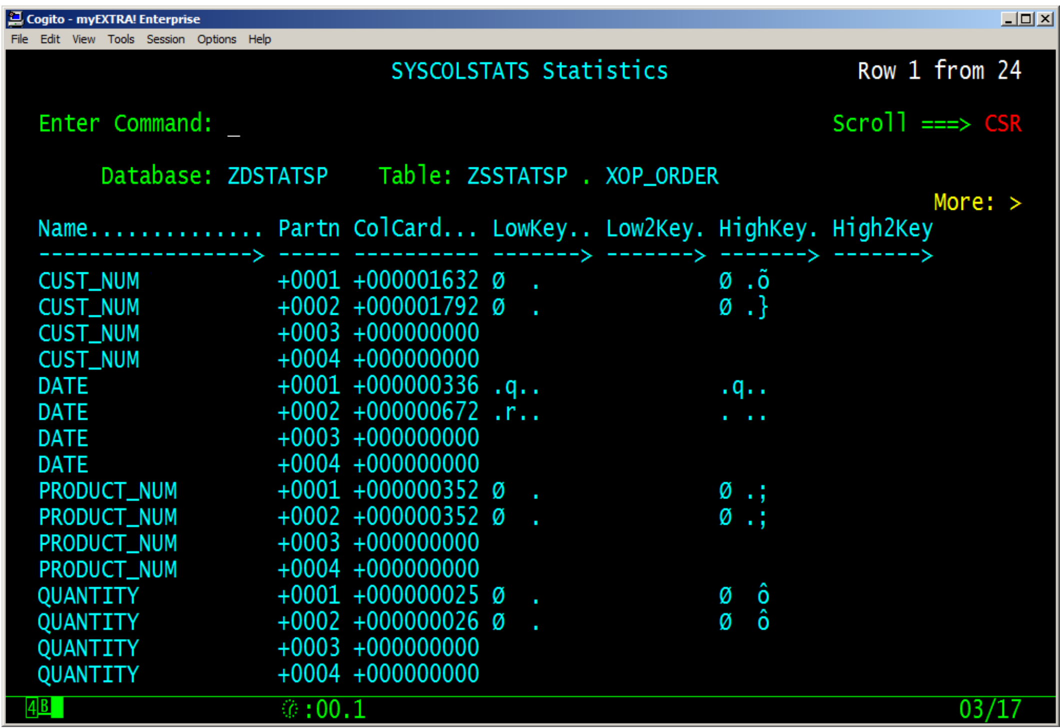

Figure 19 Display SYSCOLSTATS statistics

Database	Name of the current database.
Table	Name of current table.
Name	Name of the column.
Partn	Partition number for the table space that contains the table in which the column is defined.
Colcard	Number of distinct values in the partition.
Lowkey	Lowest value of the column within the partition. Blank if statistics have not been gathered or the column is an indicator column, a node ID column, or a column of an XML table. If the column has a

non-character data type, the data might not be printable. If the partition is empty, the value is a string of length 0.

Low2Key

Second lowest value of the column within the partition. Blank if statistics have not been gathered or the column is an indicator column, a node ID column, or a column of an XML table. If the column has a non-character data type, the data might not be printable. If the partition is empty, the value is a string of length 0.

Highkey

Highest value of the column within the partition. Blank if statistics have not been gathered or the column is an indicator column, a node ID column, or a column of an XML table. If the column has a non-character data type, the data might not be printable. If the partition is empty, the value is a string of length 0.

High2Key

Second highest value of the column within the partition. Blank if statistics have not been gathered or the column is an indicator column, a node ID column, or a column of an XML table. If the column has a non-character data type, the data might not be printable. If the partition is empty, the value is a string of length 0.

StatsTime

If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics.

F(Stats_Format)

The type of statistics gathered:

Blank **S**tatistics have not been collected or varchar column statistical values are padded.

N Varchar column statistical values are not padded.

This is an updatable column.

I(BM Required)

IBM Internal.

5.6.3 List Tables (SYSTABLES)

The SYSIBM.SYSTABLES table contains one row for each table, view, or alias.

Select option 5.6.3 to display the List Tables panel as shown in the following figure:-

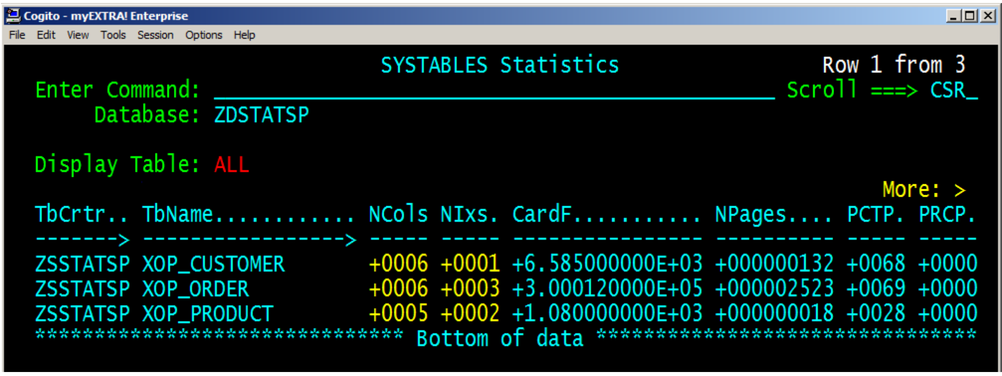

Figure 20 List Tables

The display shows the SYSTABLES statistics for the selected database.

Display Table	Table name	Enter a table name in this field to display Statistics for a particular table.
	ALL	Enter ALL in this field to display statistics for ALL tables for the current database. The default value is ALL.

The following information is displayed:-

TBCrtr	Authorisation ID of the owner of the table.
TbName	Name of the table.
NCols	The number of columns in the table. You can DRILL <PF4> on **NCols** field for a table to display Table Column Statistics which shows statistics for the columns in the selected table.
Nixs	The number of Indexes created for the table. You can DRILL <PF4> on **NIxs** field for a table to display Table Index Statistics which shows statistics for the indexes for the selected table.
CardF	Total number of rows in the table or total number of LOBs in an auxiliary table. The value is -1 if

statistics have not been gathered or the row describes a view, alias, or created temporary table.

NPages Total number of pages on which rows of the table appear. The value is -1 if statistics have not been gathered, or the row describes a view, an alias, a created temporary table, or an auxiliary table.

PCTP Percentage of active table space pages that contain rows of the table. A page is termed active if it is formatted for rows, regardless of whether it contains any. If the table space is segmented, the percentage is based on the number of active pages in the set of segments assigned to the table. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

PRCP Percentage of rows compressed within the total number of active rows in the table. This includes any row in a table space that is defined with COMPRESS YES. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

NPagesF Number of pages used by the table. The value is -1 if statistics have not been gathered or this is an auxiliary table.

SpaceF Kilobytes of DASD storage. The value is -1 if statistics have not been gathered. The value might be non-zero for an auxiliary table in the LOB table space.

KCols Number of columns in the primary key of the table. The value is 0 if the row describes a view, an alias, or a created temporary table.

RecLn For user tables, the maximum length of any record in the table. Length is 8+*N*+*L*, where:

- The number 8 accounts for the header (6 bytes) and the ID map entry (2 bytes).
- *N* is 10 if the table has an edit procedure, or 0 otherwise.
- *L* is the sum of the maximum column lengths. In determining the maximum length of a column, take into account whether the column allows nulls and the data type of the column. If the column can contain nulls and is not a LOB or ROWID column, add 1 byte for a null

indicator. Use 4 bytes for the length of a LOB column and 19 bytes for the length of a ROWID column. If the column has a varying-length data type (for example, VARCHAR, CLOB, or BLOB), add 2 bytes for a length indicator. For more information on column lengths, see “Data types” in the SQL Reference.

The value is 0 if the row describes a view, alias, or auxiliary table. For maximum row and record sizes, see Maximum record size in the SQL Reference.

AvgRowLn	Average length of rows for the tables in the table space. If the table space is compressed, the value is the compressed row length. If the table space is not compressed, the value is the uncompressed row length. The value is -1 if statistics have not been gathered.
STATSTIME	If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'. For a created temporary table, the value of STATSTIME is always the default value.

Note the following applies to all EZ-Stats statistics displays:-

- Extensive field level **HELP** is available on all fields in the **EZ-Stats** statistics displays. Position the cursor on a field and press <**PF1**> function key to obtain a description of the column.
- You can **SORT** on fields where applicable. Just position the Cursor on a Column and press **<Enter>** or enter the command SORT A or SORT D in the Command Line.
- Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.
- Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

Table Column Statistics (SYSCOLUMNS)

The SYSIBM.SYSCOLUMNS table contains one row for every column of each table and view.

To display the Table Column statistics, DRILL <PF4> on the **NCols** field when viewing [List Tables](#) display.

The following panel is displayed:-

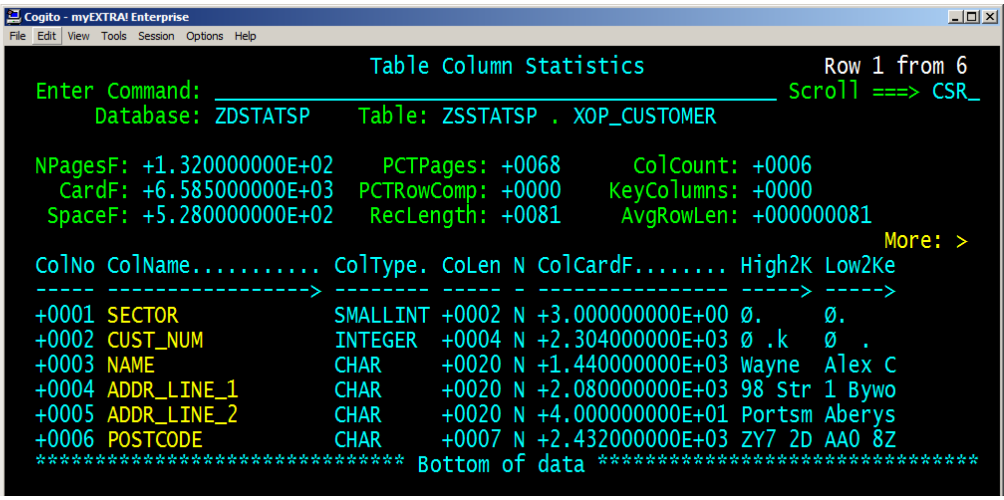

Figure 21 Table Column Statistics

The display includes the SYSTABLES statistics for the selected table and then each applicable table Column with associated SYSCOLUMNS statistics.

The following information is displayed:-

Database	Name of the current database.
Table	Name of current table.
NPagesF	Number of pages used by the table. The value is -1 if statistics have not been gathered or this is an auxiliary table
CardF	Total number of rows in the table or total number of LOBs in an auxiliary table. The value is -1 if statistics have not been gathered or the row describes a view, alias, or created temporary table.
SpaceF	Kilobytes of DASD storage. The value is -1 if statistics have not been gathered. The value might be non-zero for an auxiliary table in the LOB table space.
PCTPages	Percentage of active table space pages that contain rows of the table. A page is termed active

if it is formatted for rows, regardless of whether it contains any. If the table space is segmented, the percentage is based on the number of active pages in the set of segments assigned to the table. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

PCTRowComp

Percentage of rows compressed within the total number of active rows in the table. This includes any row in a table space that is defined with COMPRESS YES. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

RecLength

For user tables, the maximum length of any record in the table. Length is 8+*N*+*L*, where:

- The number 8 accounts for the header (6 bytes) and the ID map entry (2 bytes).
- *N* is 10 if the table has an edit procedure, or 0 otherwise.
- *L* is the sum of the maximum column lengths. In determining the maximum length of a column, take into account whether the column allows nulls and the data type of the column. If the column can contain nulls and is not a LOB or ROWID column, add 1 byte for a null indicator. Use 4 bytes for the length of a LOB column and 19 bytes for the length of a ROWID column. If the column has a varying-length data type (for example, VARCHAR, CLOB, or BLOB), add 2 bytes for a length indicator. For more information on column lengths, see “Data types” in the SQL Reference.

The value is 0 if the row describes a view, alias, or auxiliary table. For maximum row and record sizes, see Maximum record size in the SQL Reference.

ColCount

The number of columns in the table.

KeyColumns

Number of columns in the primary key of the table. The value is 0 if the row describes a view, an alias, or a created temporary table.

AvgRowLen

Average length of rows for the tables in the table space. If the table space is compressed, the value is the compressed row length. If the table space is not compressed, the value is the uncompressed row length. The value is -1 if statistics have not been gathered.

Colno

Numeric place of the column in the table or view; for example 4 (out of 10).

ColName

Name of the column.

You could further DRILL on a particular ColName to display the SYSCOLDIST statistics (Column Distribution Statistics) for the selected column if applicable.

ColType

The type of the column specified in the definition of the column:

INTEGER	Large integer
SMALLINT	Small integer
FLOAT	Floating-point
CHAR	Fixed-length character string
VARCHAR	Varying-length character string
LONGVAR	Varying-length character string (for columns that were added before Version 9)
DECIMAL	Decimal
GRAPHIC	Fixed-length graphic string
VARG	Varying-length graphic string
LONGVARG	Varying-length graphic string (for Columns that were added before Version 9)
DATE	Date
TIME	Time
TIMESTMP	Timestamp
TIMESTZ	Timestamp with time zone
BLOB	Binary large object
CLOB	Character large object
DBCLOB	Double-byte character large object
ROWID	Row ID data type
DISTINCT	Distinct type
XML	XML data type
BIGINT	Big integer
BINARY	Fixed-length binary string
VARBIN	Varying-length binary string
DECFLOAT	Decimal floating point

CoLen

Length attribute of the column or, in the case of a decimal column, its precision. The number does not include the internal prefixes that are used to record the actual length and null state, where applicable

.

N(ulls)

Whether the column can contain null values:

N No

Y Yes

The value can be N for a view column that is derived from an expression or a function. Nevertheless, such a column allows nulls when an outer select list refers to it.

ColCardf

Estimated number of distinct values in the column. For an indicator column, this is the number of LOBs that are not null and have a length greater than zero. The value is -1 if statistics have not been gathered. The value is -2 for the first column of an index of an auxiliary table.

High2Key

Second highest value of the column. Blank if statistics have not been gathered, or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0.

Expand field <**PF6**> to display the full field contents.

Low2Key

Second lowest value of the column. Blank if statistics have not been gathered, or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0. This is an updatable column.

Expand field <**PF6**> to display the full field contents.

STATSTIME

If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'. If the value is '0001-01-02.00.00.00.000000', which indicates that an ALTER TABLE statement was executed to change the length of a VARCHAR column, RUNSTATS should be run to update the statistics before they are used.

Enter <**PF6**> to expand a field, for example if the displayed value is truncated.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

Table Index Statistics (SYSINDEXES)

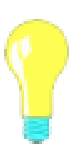

The SYSIBM.SYSINDEXES table contains one row for every index.

To display the Table Index statistics, DRILL <PF4> on the **Nixs** field when viewing List Tables display.

The following panel is displayed:-

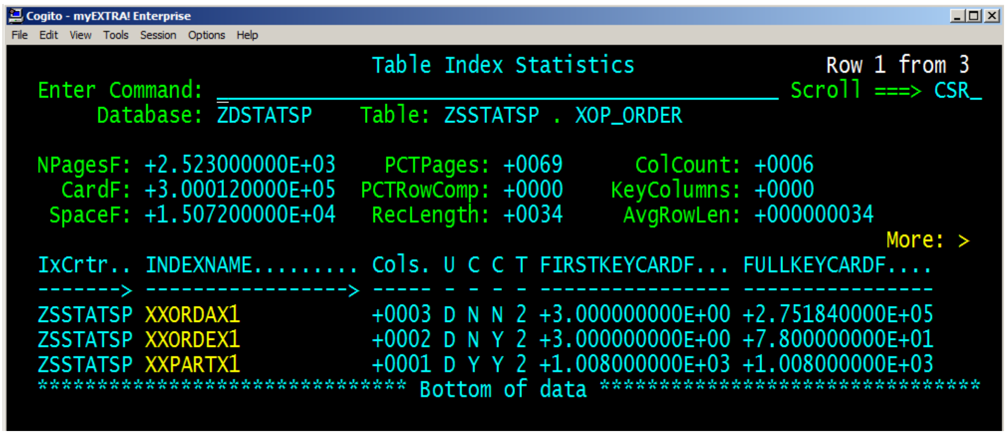

Figure 22 Table Index Statistics

The display includes the SYSTABLES statistics for the selected table and then each applicable index with associated SYSINDEXES statistics.

The following information is displayed:-

Database Name of the current database.

Table Name of current table.

NPagesF Number of pages used by the table. The value is -1 if statistics have not been gathered or this is an auxiliary table

CardF Total number of rows in the table or total number of LOBs in an auxiliary table. The value is -1 if statistics have not been gathered or the row describes a view, alias, or created temporary table.

SpaceF Kilobytes of DASD storage. The value is -1 if statistics have not been gathered. The value might be non-zero for an auxiliary table in the LOB table space.

PCTPages Percentage of active table space pages that contain rows of the table. A page is termed active if it is formatted for rows, regardless of whether it contains any. If the table space is segmented, the percentage is based on the number of active pages in the set of segments assigned to the table. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

PCTRowComp

Percentage of rows compressed within the total number of active rows in the table. This includes any row in a table space that is defined with COMPRESS YES. The value is -1 if statistics have not been gathered, or the row describes a view, alias, created temporary table, or auxiliary table.

RecLength

For user tables, the maximum length of any record in the table. Length is 8+*N*+*L*, where:

- The number 8 accounts for the header (6 bytes) and the ID map entry (2 bytes).
- *N* is 10 if the table has an edit procedure, or 0 otherwise.
- *L* is the sum of the maximum column lengths. In determining the maximum length of a column, take into account whether the column allows nulls and the data type of the column. If the column can contain nulls and is not a LOB or ROWID column, add 1 byte for a null indicator. Use 4 bytes for the length of a LOB column and 19 bytes for the length of a ROWID column. If the column has a varying-length data type (for example, VARCHAR, CLOB, or BLOB), add 2 bytes for a length indicator. For more information on column lengths, see “Data types” in the SQL Reference.

The value is 0 if the row describes a view, alias, or auxiliary table. For maximum row and record sizes, see Maximum record size in the SQL Reference.

ColCount

The number of columns in the table.

KeyColumns

Number of columns in the primary key of the table. The value is 0 if the row describes a view, an alias, or a created temporary table.

AvgRowLen

Average length of rows for the tables in the table space. If the table space is compressed, the value is the compressed row length. If the table space is not compressed, the value is the uncompressed row length. The value is -1 if statistics have not been gathered.

IXCrtr — Authorization ID of the owner of the Index,

INDEXNAME — The Index Name.

You can DRILL <PF4> on the INDEXNAME to display the Index Column Details display for the selected Index.

Cols — The number of columns in the key.

U(nique) — Whether the index is unique:

- **C** Yes, and it is used to enforce the uniqueness of a UNIQUE constraint or hash key columns
- **D** No (duplicates are allowed)
- **U** Yes
- **P** Yes, and it is a primary index (As in prior releases of DB2, a value of P is used for primary keys that are used to enforce a referential constraint.)
- **N** Yes, and it is defined with UNIQUE WHERE NOT NULL
- **R** Yes, and it is an index used to enforce the uniqueness of a non-primary parent key **G** Yes, and it is an index used to enforce the uniqueness of values in a column defined as ROWID GENERATED BY DEFAULT
- **X** Yes, and it is an index used to enforce the uniqueness of values in a column that is used to identify or find XML values associated with a specific row.

C(lustering) — Whether CLUSTER was specified for the index:

- **N** No
- **Y** Yes

C(lustered) — Whether the table is actually clustered by the index:

- **N** A significant number of rows are not in clustering order, or statistics have not been gathered.
- **Y** Most of the rows are in clustering order.
- **Blank** Not applicable.

	This is an updatable column that can also be changed by the RUNSTATS utility
T(ype)	The index type: **2** Type 2 index or a hash overflow index on non-partitioned tables. **Blank** Type 1 index **D** Data-partitioned secondary index **P** An index that is both partitioned and is a partitioning index (index that is on a table that uses table-controlled partitioning).
FIRSTKEYCARDF	Number of distinct values of the first key column. This number is an estimate if updated while collecting statistics on a single partition. The value is -1 if statistics have not been gathered.
FULLKEYCARDF	Number of distinct values of the key. The value is -1 if statistics have not been gathered.
CLUSTERRRATIOF	Percentage of rows that are in clustering order. For a partitioning index, it is the weighted average of all index partitions in terms of the number of rows in the partition. The value is 0 if statistics have not been gathered. The value is -2 if the index is for an auxiliary table.
NLEAF	Number of active leaf pages in the index. The value is -1 if statistics have not been gathered.
NLEVL	Number of levels in the index tree. If the index is partitioned, it is the maximum of the number of levels in the index tree for all the partitions. The value is -1 if statistics have not been gathered.
SPACEF	Kilobytes of DASD storage. The value is -1 if statistics have not been gathered.
IXSPACE	Name of the index space.
BPOOL	Name of the buffer pool used for the index.
PGSIZ	Contains the value 4, 8, 16, or 32 which indicates the size, in KB, of the leaf pages in the index. If the index was created prior to Version 9, the value will be 4096 for a 4 KB page size.
DataRptFactorF	The anticipated number of data pages that will be touched following an index key order. This statistic is only collected when the STATCLUS subsystem parameter is set to ENHANCED. This

	number is -1 if statistics have not been collected. The valid value is -1 or any value that is equal to or greater than 1.
LastUsed	This field is populated from LASTUSED in SYSINDEXSPACESTATS. The date when the index is used for SELECT, FETCH, searched UPDATE, searched DELETE, or used to enforce referential integrity constraints. The default value is NULL..
STATSTIME	If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'.
Cols	Internal Use.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

5.6.4 List Indexes (SYSINDEXES)

The SYSIBM.SYSINDEXES table contains one row for every index.

Select option 5.6.4 to display the List Indexes panel as shown in the following figure:-

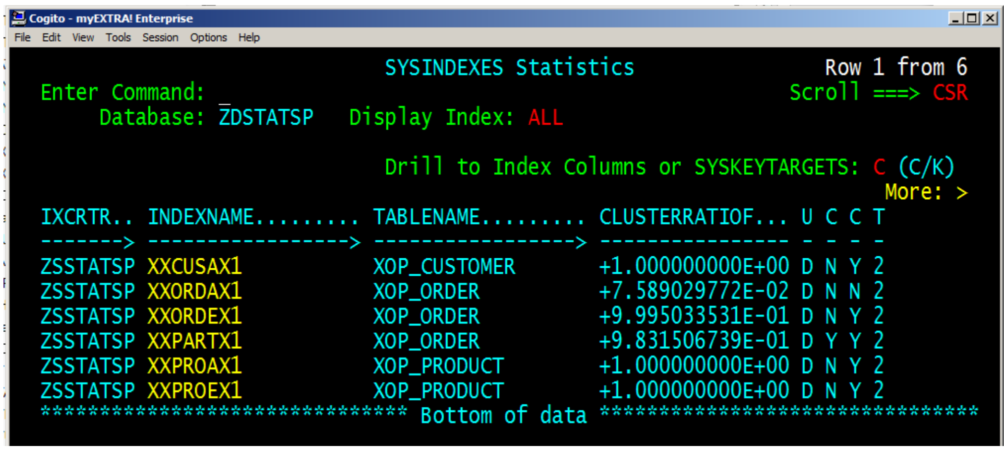

Figure 23 List Indexes

The display shows the SYSINDEXES statistics for all of the indexes for the selected database.

Database		The selected database name.
Display Index	**Index name**	Enter an Index name in this field to display Statistics for a particular index.
	ALL	Enter ALL in this field to display statistics for ALL Indexes for the current database.
		The default value is ALL.
DRILL to Index Columns or SYSKEYTARGETS	**(C/K)**	Enter C to DRILL to Index Columns.
		Enter K to DRILL to SYSKEYTARGETS,

The following information is displayed for each index:-

IXCRTR	The schema of the index.
INDEXNAME	The name of the index.
	You can DRILL <PF4> on the INDEXNAME to display either the Index Columns display (see Index Column Details) or the SYSKEYTARGETS detail display (see SYSKEYTARGETS detail) for

	the selected Index. See the **DRILL to Index Columns or SYSKEYTARGETS** setting.
TABLENAME	Name of the table on which the index is defined.
CLUSTERRRATIOF	Percentage of rows that are in clustering order. For a partitioning index, it is the weighted average of all index partitions in terms of the number of rows in the partition. The value is 0 if statistics have not been gathered. The value is -2 if the index is for an auxiliary table.
U(unique rule)	Whether the index is unique: **C** Yes, and it is used to enforce the uniqueness of a UNIQUE constraint or hash key columns. **D** No (duplicates are allowed) **U** Yes **P** Yes, and it is a primary index (As in prior releases of DB2, a value of P is used for primary keys that are used to enforce a referential constraint.) **N** Yes, and it is defined with UNIQUE WHERE NOT NULL **R** Yes, and it is an index used to enforce the uniqueness of a non-primary parent key **G** Yes, and it is an index used to enforce the uniqueness of values in a column defined as ROWID GENERATED BY DEFAULT **X** Yes, and it is an index used to enforce the uniqueness of values in a column that is used to identify or find XML values associated with a specific row.
C(lustering)	Whether CLUSTER was specified for the index: **N** No **Y** Yes
C(lustered)	Whether the table is actually clustered by the index: **N** A significant number of rows are not in clustering order, or statistics have not been gathered. **Y** Most of the rows are in clustering order. **Blank** Not applicable. This is an updatable column that can also be changed by the RUNSTATS utility

T(ype)	The index type:
	2 Type 2 index or a hash overflow index on non-partitioned tables.
	Blank Type 1 index
	D Data-partitioned secondary index
	P An index that is both partitioned and is a partitioning index (index that is on a table that uses table-controlled partitioning).
FIRSTKEYCARDF	Number of distinct values of the first key column. This number is an estimate if updated while collecting statistics on a single partition. The value is -1 if statistics have not been gathered.
FULLKEYCARDF	Number of distinct values of the key. The value is -1 if statistics have not been gathered.
NLEAF	Number of active leaf pages in the index. The value is -1 if statistics have not been gathered.
NLEVL	Number of levels in the index tree. If the index is partitioned, it is the maximum of the number of levels in the index tree for all the partitions. The value is -1 if statistics have not been gathered.
IXSPACE	Name of the index space.
BPOOL	Name of the buffer pool used for the index.
PGSIZ	Contains the value 4, 8, 16, or 32 which indicates the size, in KB, of the leaf pages in the index. If the index was created prior to Version 9, the value will be 4096 for a 4 KB page size.
SPACEF	Kilobytes of DASD storage. The value is -1 if statistics have not been gathered.
LastUsed	This field is populated from LASTUSED in SYSINDEXSPACESTATS.
	The date when the index was last used for SELECT, FETCH, searched UPDATE, searched DELETE, or used to enforce referential integrity constraints. The default value is NULL.
DataRptFactorF	The anticipated number of data pages that will be touched following an index key order. This statistic is only collected when the STATCLUS subsystem parameter is set to ENHANCED. This number is -1 if statistics have not been collected. The valid value is -1 or any value that is equal to

	or greater than 1.
STATSTIME	If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'.
Cols	Internal Use.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

Index Column Details

To display the Index Column statistics, DRILL <PF4> on the INDEXNAME field when viewing [Table Indexes]() or [List Indexes]() display.

Note. Drill to [Index Columns or SYSKEYTARGETS]() must be to **C**.

The following panel is displayed:-

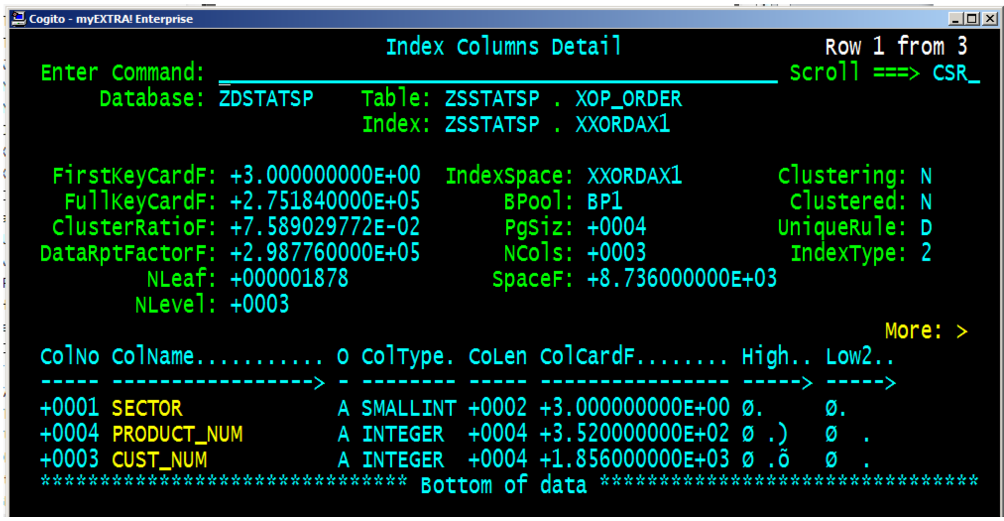

Figure 24 **Index Columns Detail**

The display includes the SYSINDEXES statistics for the selected Index, and then each applicable Column with associated SYSCOLUMNS statistics.

The following information is displayed:-

Database	The selected database name.
Table	The selected table name.
Index	The selected index name
FirstkeycardF	Number of distinct values of the first key column. This number is an estimate if updated while collecting statistics on a single partition. The value is -1 if statistics have not been gathered.
FullkeycardF	Number of distinct values of the key. The value is -1 if statistics have not been gathered.
ClusterRatioF	Percentage of rows that are in clustering order. For a partitioning index, it is the weighted average of all index partitions in terms of the number of rows in the partition. The value is 0 if statistics have not been gathered. The value is -2 if the index is for an auxiliary table.

DataRptFactorF	The anticipated number of data pages that will be touched following an index key order. This statistic is only collected when the STATCLUS subsystem parameter is set to ENHANCED. This number is -1 if statistics have not been collected. The valid value is -1 or any value that is equal to or greater than 1.
NLeaf	Number of active leaf pages in the index. The value is -1 if statistics have not been gathered.
Nlevel	Number of levels in the index tree. If the index is partitioned, it is the maximum of the number of levels in the index tree for all the partitions. The value is -1 if statistics have not been gathered.
IndexSpace	Name of the index space.
BPool	Name of the buffer pool used for the index.
PgSiz	Contains the value 4, 8, 16, or 32 which indicates the size, in KB, of the leaf pages in the index. If the index was created prior to Version 9, the value will be 4096 for a 4 KB page size.
NCols	The number of columns in the key.
SpaceF	Kilobytes of DASD storage. The value is -1 if statistics have not been gathered.
Clustering	Whether CLUSTER was specified for the index: **N** No **Y** Yes
Clustered	Whether the table is actually clustered by the index: **N** A significant number of rows are not in clustering order, or statistics have not been gathered. **Y** Most of the rows are in clustering order. **Blank** Not applicable. This is an updatable column that can also be changed by the RUNSTATS utility
UniqueRule	Whether the index is unique: **C** Yes, and it is used to enforce the uniqueness of a UNIQUE constraint or hash key columns. **D** No (duplicates are allowed) **U** Yes **P** Yes, and it is a primary index (As in prior

releases of DB2, a value of P is used for primary keys that are used to enforce a referential constraint.)

N Yes, and it is defined with UNIQUE WHERE NOT NULL

R Yes, and it is an index used to enforce the uniqueness of a non-primary parent key

G Yes, and it is an index used to enforce the uniqueness of values in a column defined as ROWID GENERATED BY DEFAULT

X Yes, and it is an index used to enforce the uniqueness of values in a column that is used to identify or find XML values associated with a specific row.

IndexType

The index type:

2 Type 2 index or a hash overflow index on non-partitioned tables.

Blank Type 1 index

D Data-partitioned secondary index

P An index that is both partitioned and is a partitioning index (index that is on a table that uses table-controlled partitioning).

Colno

Numeric place of the column in the table or view; for example 4 (out of 10).

ColName

Name of the column.

You could further DRILL on a particular ColName to display the SYSCOLDIST statistics (Column Distribution Statistics) for the selected column if applicable.

O(order)

Order of the column in the key:

Blank Index is an index based on expressions or the column is specified for the index using the INCLUDE clause

A Ascending order

D Descending order

R Random order

ColType

The type of the column specified in the definition of the column:

INTEGER Large integer

SMALLINT Small integer

FLOAT Floating-point

CHAR	Fixed-length character string
VARCHAR	Varying-length character string
LONGVAR	Varying-length character string (for columns that were added before Version 9)
DECIMAL	Decimal
GRAPHIC	Fixed-length graphic string
VARG	Varying-length graphic string
LONGVARG	Varying-length graphic string (for Columns that were added before Version 9)
DATE	Date
TIME	Time
TIMESTMP	Timestamp
TIMESTZ	Timestamp with time zone
BLOB	Binary large object
CLOB	Character large object
DBCLOB	Double-byte character large object
ROWID	Row ID data type
DISTINCT	Distinct type
XML	XML data type
BIGINT	Big integer
BINARY	Fixed-length binary string
VARBIN	Varying-length binary string
DECFLOAT	Decimal floating point

CoLen

Length attribute of the column or, in the case of a decimal column, its precision. The number does not include the internal prefixes that are used to record the actual length and null state, where applicable

ColCardf

Estimated number of distinct values in the column. For an indicator column, this is the number of LOBs that are not null and have a length greater than zero. The value is -1 if statistics have not been gathered. The value is -2 for the first column of an index of an auxiliary table.

High(2key)

Second highest value of the column. Blank if statistics have not been gathered, or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0.

Low2(key)

Second lowest value of the column. Blank if statistics have not been gathered or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0. This is an updatable column.

STATSTIME

If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'.

N(ulls)

Whether the column can contain null values:

N No
Y Yes

The value can be N for a view column that is derived from an expression or a function. Nevertheless, such a column allows nulls when an outer select list refers to it.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

SYSKEYTARGETS details

The SYSIBM.SYSKEYTARGETS table contains one row for each key-target that is participating in an extended index definition. E.g. Index on Expression. . If the selected index is not an extended index definition then the display will be blank.

To display the SYSKEYTARGETS statistics, DRILL <PF4> on the INDEXNAME field from the List Indexes display. Note: Drill to Index Columns or SYSKEYTARGETS must be set to **K**.

The following panel is displayed:-

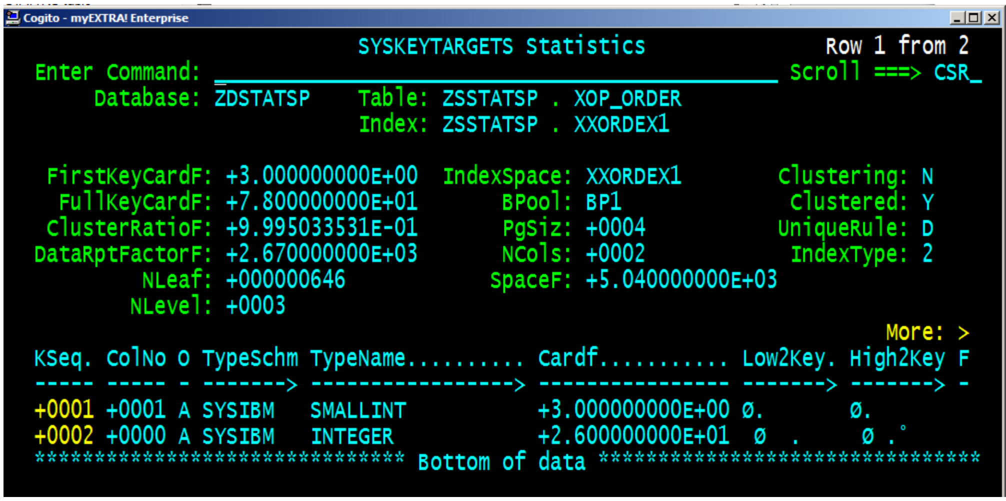

Figure 25 SYSKEYTARGETS Detail

The display includes the SYSINDEXES statistics for the selected Index, and then each applicable Column with associated SYSCOLUMNS statistics
The following information is displayed:-

Database	The selected database name.
Table	The selected table name.
Index	The selected index name
FirstkeycardF	Number of distinct values of the first key column. This number is an estimate if updated while collecting statistics on a single partition. The value is -1 if statistics have not been gathered.
FullkeycardF	Number of distinct values of the key. The value is -1 if statistics have not been gathered.
ClusterRatioF	Percentage of rows that are in clustering order. For a partitioning index, it is the weighted average of all index partitions in terms of the number of rows in the partition. The value is 0 if statistics have not been gathered. The value is -2 if the index is for an auxiliary table.

DataRptFactorF
The anticipated number of data pages that will be touched following an index key order. This statistic is only collected when the STATCLUS subsystem parameter is set to ENHANCED. This number is -1 if statistics have not been collected. The valid value is -1 or any value that is equal to or greater than 1.

NLeaf
Number of active leaf pages in the index. The value is -1 if statistics have not been gathered.

Nlevel
Number of levels in the index tree. If the index is partitioned, it is the maximum of the number of levels in the index tree for all the partitions. The value is -1 if statistics have not been gathered.

IndexSpace
Name of the index space.

BPool
Name of the buffer pool used for the index.

PgSiz
Contains the value 4, 8, 16, or 32 which indicates the size, in KB, of the leaf pages in the index.
If the index was created prior to Version 9, the value will be 4096 for a 4 KB page size.

NCols
The number of columns in the key.

SpaceF
Kilobytes of DASD storage. The value is -1 if statistics have not been gathered.

Clustering
Whether CLUSTER was specified for the index:

N No
Y Yes

Clustered
Whether the table is actually clustered by the index:

N A significant number of rows are not in clustering order, or statistics have not been gathered.

Y Most of the rows are in clustering order.

Blank Not applicable.

This is an updatable column that can also be changed by the RUNSTATS utility

UniqueRule
Whether the index is unique:
C Yes, and it is used to enforce the uniqueness of a UNIQUE constraint or hash key columns.
D No (duplicates are allowed)
U Yes

P Yes, and it is a primary index (As in prior releases of DB2, a value of P is used for primary keys that are used to enforce a referential constraint.)

N Yes, and it is defined with UNIQUE WHERE NOT NULL

R Yes, and it is an index used to enforce the uniqueness of a non-primary parent key

G Yes, and it is an index used to enforce the uniqueness of values in a column defined as ROWID GENERATED BY DEFAULT

X Yes, and it is an index used to enforce the uniqueness of values in a column that is used to identify or find XML values associated with a specific row.

IndexType The index type:

2 Type 2 index or a hash overflow index on non-partitioned tables.

Blank Type 1 index

D Data-partitioned secondary index

P An index that is both partitioned and is a partitioning index (index that is on a table that uses table-controlled partitioning).

KSeq Numeric position of the key-target in the index.

You can DRLL <PF4> on the Kseq value to display the SYSKEYTGTDIST statistics for the Key.

Colno Numeric position of the column in the table if the expression is a single column. Otherwise the value is 0.

For XML indexes, this field is also 0.

O(order) Order of Key.

A Ascending order

TypeSchm Schema of the data type.

TypeName Name of the data type.

Cardf The estimated number of distinct values for the key-target. The value is -2 if the index is a node ID index.

For an XML value index, the statistic is collected for the second key target (the DOCID column).

For all other key targets of the XML value index, a value of -2 is set.

Low2Key — The second lowest key-value. LOW2KEY is an updatable column.

High2Key — The second highest key-value. HIGH2KEY is an updatable column.

F(Stats_Format) — The type of statistics that are gathered:

N	VARCHAR column statistical values are not padded
Blank	Statistics have not been collects or VARCHAR column statistical values are padded.

STATS_FORMAT is an updatable column.

STATSTIME — If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics. The default value is '0001-01-01.00.00.00.000000'.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

Column Distribution Statistics (SYSCOLDIST)

The SYSIBM.SYSCOLDIST table contains one or more rows for the first key column of an index key.

To display the Key Column statistics, DRILL <PF4> on the **Colname** field when viewing Table columns or Index Columns display.

Note that the layout of the display varies depending in whether you select to display the Cardinality, Frequency or Histogram Statistics. The following figure shows an example of the Histogram statistics display.

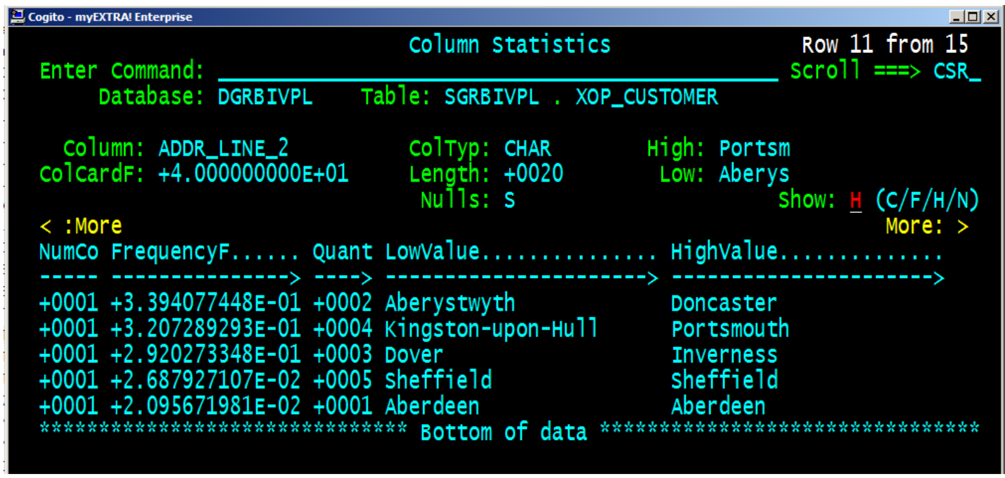

```
Cogito - myEXTRA! Enterprise
                           Column Statistics                    Row 11 from 15
 Enter Command: ______________________________________________  Scroll ===> CSR_
      Database: DGRBIVPL    Table: SGRBIVPL . XOP_CUSTOMER

   Column: ADDR_LINE_2          ColTyp: CHAR        High: Portsm
 ColCardF: +4.000000000E+01     Length: +0020        Low: Aberys
                                 Nulls: S                      Show: H (C/F/H/N)
 < :More                                                                More: >
 NumCo FrequencyF...... Quant LowValue.............. HighValue.............
 ----- ---------------> ----> ---------------------> ---------------------->
 +0001 +3.394077448E-01 +0002 Aberystwyth            Doncaster
 +0001 +3.207289293E-01 +0004 Kingston-upon-Hull     Portsmouth
 +0001 +2.920273348E-01 +0003 Dover                  Inverness
 +0001 +2.687927107E-02 +0005 Sheffield              Sheffield
 +0001 +2.095671981E-02 +0001 Aberdeen               Aberdeen
 ******************************* Bottom of data ********************************
```

Figure 26 Column Distribution Statistics

The display shows the SYSCOLUMNS information for the selected column followed by the column distribution statistics as applicable.

The following information is displayed:-

The following information is displayed:-

Database	The selected database name.
Table	The selected table name.
Column	The selected column name
ColCardf	Estimated number of distinct values in the column. For an indicator column, this is the number of LOBs that are not null and have a length greater than zero. The value is -1 if statistics have not been gathered. The value is -2 for the first column of an index of an auxiliary table.

ColType

The type of the column specified in the definition of the column:

INTEGER	Large integer
SMALLINT	Small integer
FLOAT	Floating-point
CHAR	Fixed-length character string
VARCHAR	Varying-length character string
LONGVAR	Varying-length character string (for columns that were added before Version 9)
DECIMAL	Decimal
GRAPHIC	Fixed-length graphic string
VARG	Varying-length graphic string
LONGVARG	Varying-length graphic string (for Columns that were added before Version 9)
DATE	Date
TIME	Time
TIMESTMP	Timestamp
TIMESTZ	Timestamp with time zone
BLOB	Binary large object
CLOB	Character large object
DBCLOB	Double-byte character large object
ROWID	Row ID data type
DISTINCT	Distinct type
XML	XML data type
BIGINT	Big integer
BINARY	Fixed-length binary string
VARBIN	Varying-length binary string
DECFLOAT	Decimal floating point

Length

Length attribute of the column or, in the case of a decimal column, its precision. The number does not include the internal prefixes that are used to record the actual length and null state, where applicable.

Nulls

Whether the column can contain null values:

N	No
Y	Yes

The value can be N for a view column that is derived from an expression or a function. Nevertheless, such a column allows nulls when an outer select list refers to it.

High(2key)

Second highest value of the column. Blank if statistics have not been gathered, or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0.

Low2(key)	Second lowest value of the column. Blank if statistics have not been gathered or the column is an indicator column or a column of an auxiliary table. If the column has a non-character data type, the data might not be printable. If the table is empty, the value is a string of length 0. This is an updatable column.
Show	C Display the cardinality statistics H Display the Histogram statistics F Display the Frequent value statistics N Display Non-padded frequent value statistics. The remainder of the display varies depending on which statistics are being displayed,
NUmCo(lumns)	Identifies the number of columns associated with the statistics.
ColGroupColNo	Identifies the set of columns associated with the statistics. If the statistics are only associated with a single column, the field contains a zero length. Otherwise, the field is an array of SMALLINT column numbers with a dimension equal to the value in NUMCOLUMNS.
CardF	For TYPE='C', the number of distinct values for the column group. For TYPE='H', the number of distinct values for the column group in a quantile indicated by QUANTILENO.
FrequencyF	Gives the percentage of rows in the table with the value specified in COLVALUE when the number is multiplied by 100. For example, a value of '1' indicates 100%. A value of '.153' indicates 15.3%. When TYPE='H', this is the percentage of rows in table which falls at the quantile indicated by QUANTILENO whose range is limited by [LOWVALUE, HIGHVALUE]. Statistics are not collected for an index on a ROWID column.
ColValue	Contains the data of a frequently occurring value. Statistics are not collected for an index on a ROWID column. If the value has a non-character data type, the data might not be printable.
Quantile(no)	Ordinary sequence number of a quantile in the whole consecutive value range, from low to high. This column is not updatable.

LowValue
For TYPE='H', this is the lower bound for the quantile indicated by QUANTILENO. Not used if TYPE is not 'H'. This column is not updatable

HighValue
For TYPE='H', this is the higher bound for the quantile indicated by QUANTILENO. Not used if TYPE is not 'H'. This column is not updatable

StatsTime
If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

SYSKEYTGTDIST statistics

The SYSIBM.SYSKEYTGTDIST table contains one or more rows for the first key-target of an extended index key.

To display the SYSKEYTGTDST, DRILL <PF4> on the **KSeq** field when viewing the SYSKEYTARGETS details.

Note that the layout of the display varies depending in whether you select to display the Cardinality, Frequency or Histogram Statistics.

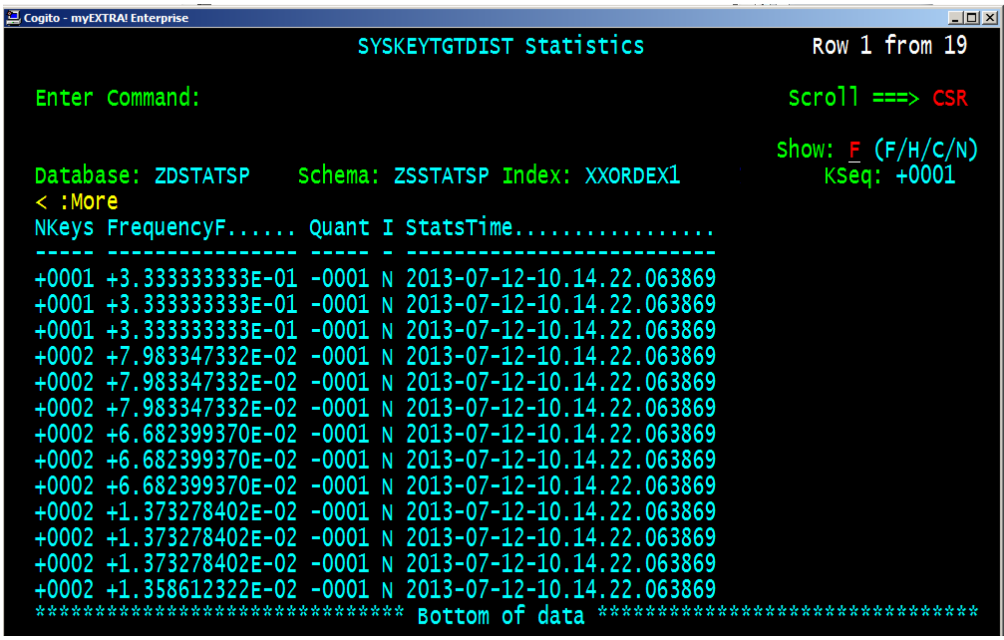

Figure 27 SYSKEYTGTDIST statistics

The following information is displayed:-

Database	The selected database name.
Schema	The qualifier of the index.
Index	The Index name.
Show	C Display the cardinality statistics H Display the Histogram statistics F Display the Frequent value statistics N Display Non-padded frequent value statistics. The remainder of the display varies depending on which statistics are being displayed,
KeyValue	KEYVALUE contains the data of a frequently occurring value. If the value has a non-character

data type, the data might not be printable.

Type

The type of statistics that are gathered:

C Cardinality
F Frequent value
N Non-padded frequent value
H Histogram statistics

CardF

When TYPE='C', CARDF contains the number of distinct values for the key group.

When TYPE='H', CARDF contains the number of distinct values for the key group in a quantile indicated by QUANTILENO.

KeyGroupKeyNo

KEYGROUPKEYNO contains a value that identifies the set of keys that are associated with the statistics.

KEYGROUPKEYNO contains 0 if the statistics are only associated with a single key.
If the statistics are associated with more than a single key,

KEYGROUPKEYNO contains an array of SMALLINT key numbers with a dimension that is equal to the value in NUMKEYS.

NKeys

The number of keys that are associated with the statistics

FrequencyF

When TYPE='F' or 'N', FREQUENCYF contains a value that indicates the percentage of entries in the index that have the value that is contained in the KEYVALUE column.

When TYPE='H', FREQUENCYF contains a value that indicates the percentage of entries in the index that have a value that is in the range of the quantile that is indicated in the QUALTILENO column.

To determine the percentage from the value of FREQUENCYF, multiply the value by 100. For example, a value of '1' indicates 100 percent. A value of '.153' indicates '15.3' percent.

Quant

QUANTILENO contains an ordinary sequence number of a quantile in the whole consecutive value range, from low to high.

LowValue

When TYPE='H', LOWVALUE contains the lower bound for the quantile that is in QUANTILENO. LOWVALUE is not used if TYPE does not equal

	'H'.
HighValue	When TYPE='H', HIGHVALUE contains the upper bound for the quantile that is in QUANTILENO. HIGHVALUE is not used if TYPE does not equal 'H'.
I	IBMREQD.
StatsTime	If RUNSTATS updated the statistics, the date and time when the last invocation of RUNSTATS updated the statistics.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

5.6.5 List SYSROUTINES

The SYSIBM.SYSROUTINES table contains a row for every routine. (A routine can be a user-defined function, cast function, or stored procedure.)

Select option **Error! Reference source not found.** to display the List SYSROUTINES panel as shown in the following figure:-

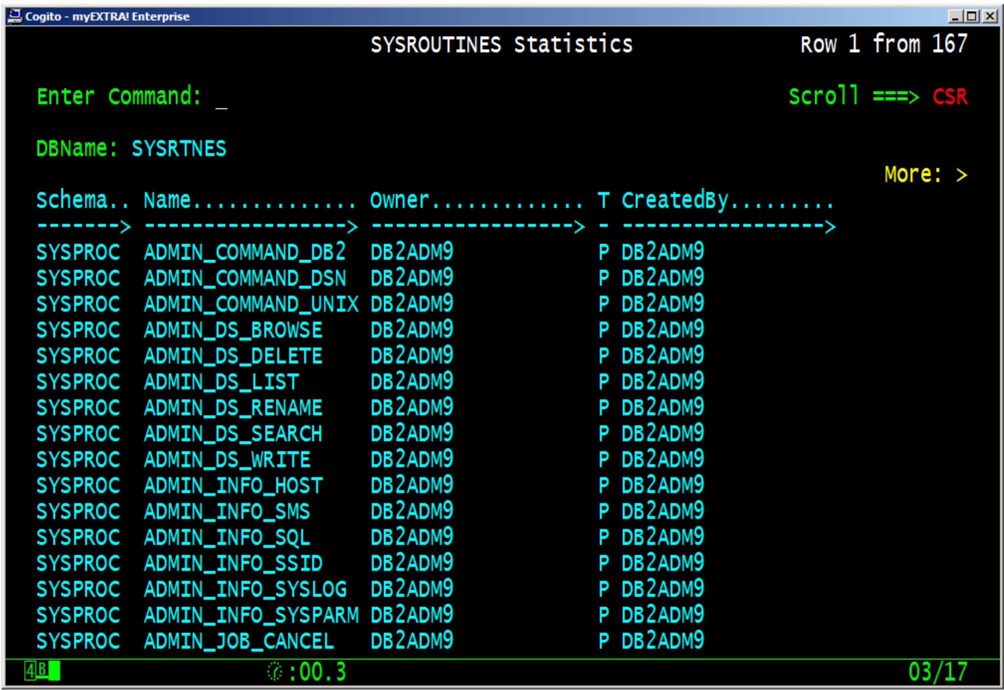

Figure 28 SYSROUTINES statistics

The following information is displayed:-

DBName For SYSROUTINES statistics, the database name must be set to SYSRTNES.

Schema Schema of the routine.

Name Name of the routine.

Owner Owner of the routine.

CreatedBy Primary authorization ID of the user who created the routine.

SpecificName Specific name of the routine.

Cardinality The predicted cardinality of the routine, -1 to trigger the use of the default value (10,000).

InitialInsts INITIAL_INSTS

	Estimated number of machine instructions that are performed the first time or the last time the routine is invoked. The value is -1 if the estimated number is not known.
InitialIOs	INSITIAL_IOS Estimated number of I/O's that are performed the first time or the last time the routine is invoked. The value is -1 if the estimated number is not known.
InstsPerInvoc	INSTS_PER_INVOC Estimated number of I/Os required to execute the routine. The value is -1 if the estimated number is not known.
IOsPerInvoc	IOS_PER_INVOC Estimated number of I/Os required to execute the routine. The value is -1 if the estimated number is not known.
RoutineID	Internal identifier of the routine.
Version	The version identifier for a native SQL procedure (indicated by the value 'N' in the column ORIGIN) or a non-inline SQL scalar function (indicated by the value 'Q' in the column ORIGIN and 'N' in the column INLINE). A zero length string for the rows that are created prior to Version 9 and for the rows that correspond to neither native SQL procedures or non-inline SQL scalar functions.

Enter **<PF11>** or DRILL on the **More: >** field to display further information to the right.

Enter **<PF10>** or DRILL on the **<: More** field to display further information to the left.

5.6.6 Update Tablespaces

Select option 5.6.6 to display the Update Tablespaces panel as shown in the following figure:-

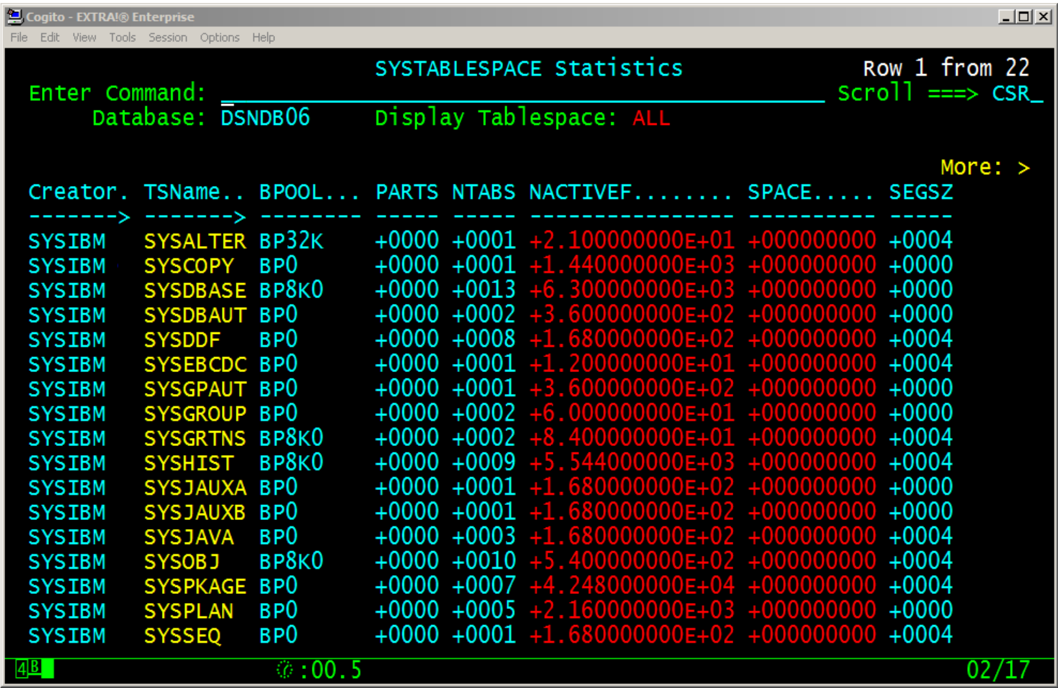

Figure 29 Update Tablespaces

The display shows the SYSTABLESPACE statistics for the selected database and is the same display as the List Tablespace display as described in section 5.6.1 List Tablespaces.

- When using the Update Object options where applicable, some fields are unprotected, allowing you to modify the values as required.
- Note that the modifications are not applied to the Statistics in the Warehouse until you select option 5.6.12 - Apply Updates to Statistics Warehouse.
- You can also use option 5.6.11 Display Updates Waiting to be to display the current updates you have entered for the current database before applying them.
- If you exit from the Display/Edit Statistics in Warehouse menu without applying the updates to the Warehouse, any updates you might have made will be discarded.
- Note that important relationships exist among certain columns in DB2 catalog tables. Consider these relationships if you choose to modify the catalog statistics. Refer to Correlations in Catalog for further information.
- See option 5.6.1 List Tablespaces for a description of the table columns displayed.

Note that you can DRILL from this panel to the Table Columns and Indexes should you wish to view/update the Column and Index statistics including the SYSCOLDIST data etc.

5.6.7 Update Partitioned Tablespaces

Select option 5.6.7 to display the Update Partitioned Tablespaces panel as shown in the following figure:-

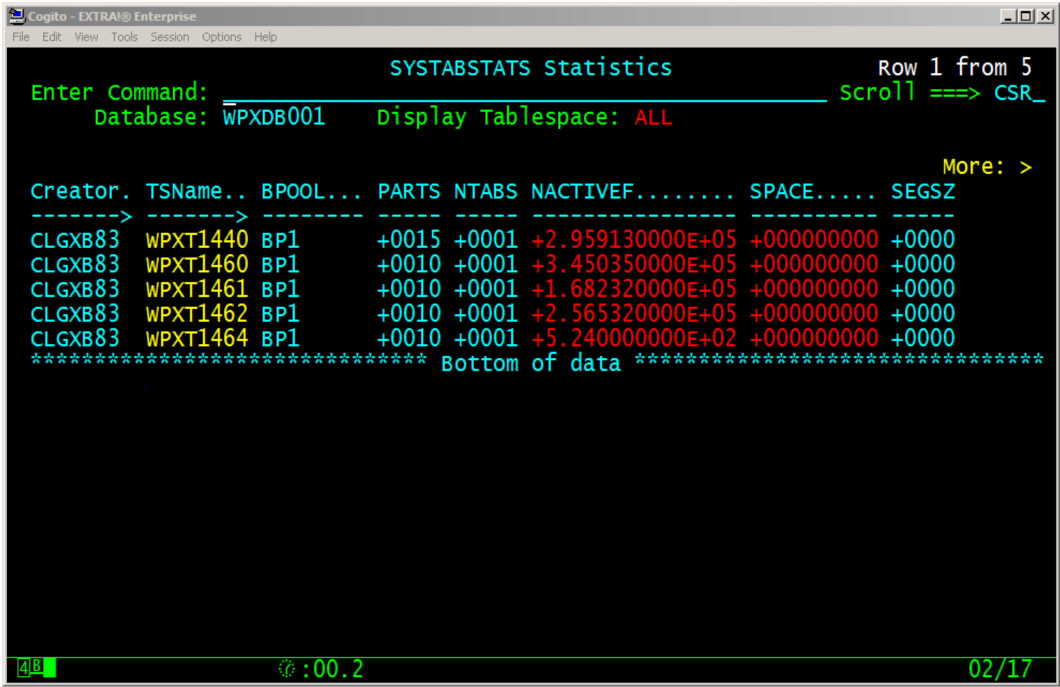

Figure 30 **Update Partitioned Tablespaces**

The display shows the SYSTABSTATS statistics for the selected database and is the same display as the List Partitioned Tablespaces display as described in section 5.6.2

- When using the Update Object options where applicable, some fields are unprotected, allowing you to modify the values as required.
- Note that the modifications are not applied to the Statistics in the Warehouse until you select option 5.6.12 - Apply Updates to Statistics Warehouse.
- You can also use option 5.6.11 Display Updates Waiting to be to display the current updates you have entered for the current database before applying them.
- If you exit from the Display/Edit Statistics in Warehouse menu without applying the updates to the Warehouse, any updates you might have made will be discarded.
- Note that important relationships exist among certain columns in DB2 catalog tables. Consider these relationships if you choose to modify the catalog statistics. Refer to Correlations in Catalog for further information.
- See option 5.6.2 for a description of the table columns displayed.

Note that you can DRILL from this panel to the Table Columns and Indexes should you wish to view/update the Column and Index statistics including the SYSCOLDIST data etc.

5.6.8 Update Tables

Select option 5.6.8 to display the Update Tables panel as shown in the following figure:-

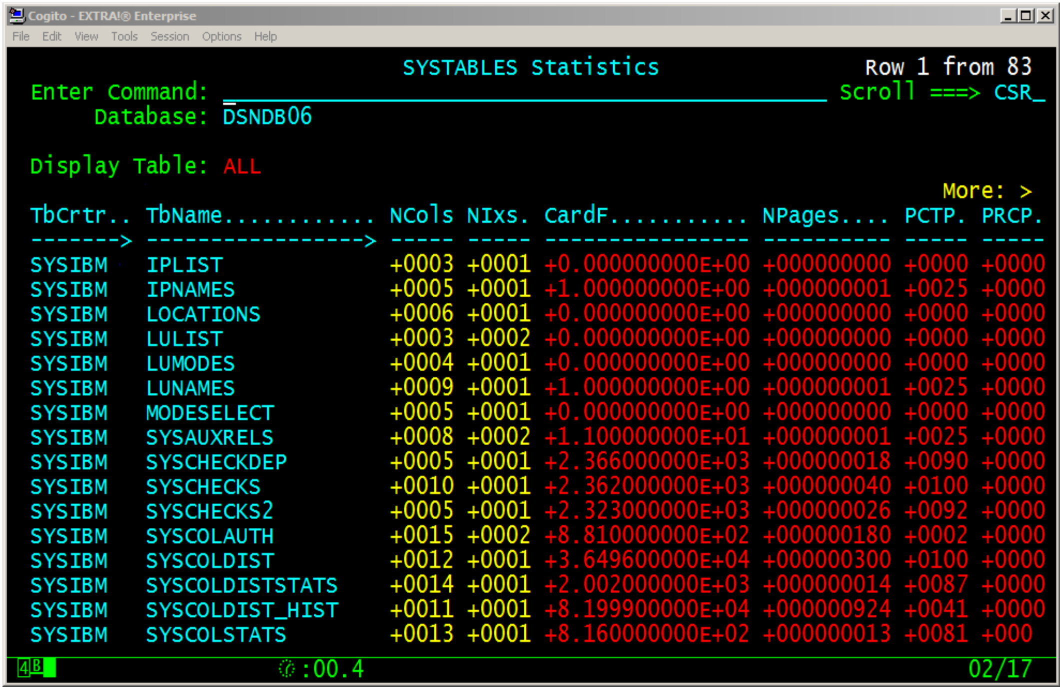
```
Cogito - EXTRA!® Enterprise
File  Edit  View  Tools  Session  Options  Help
                          SYSTABLES Statistics                   Row 1 from 83
Enter Command: ______________________________________________ Scroll ===> CSR_
     Database: DSNDB06

Display Table: ALL
                                                                      More: >
TbCrtr.. TbName........... NCols NIxs. CardF.......... NPages.... PCTP. PRCP.
-------> ----------------> ----- ----- ---------------- ---------- ----- -----
SYSIBM   IPLIST            +0003 +0001 +0.000000000E+00 +000000000 +0000 +0000
SYSIBM   IPNAMES           +0005 +0001 +1.000000000E+00 +000000001 +0025 +0000
SYSIBM   LOCATIONS         +0006 +0001 +0.000000000E+00 +000000000 +0000 +0000
SYSIBM   LULIST            +0003 +0002 +0.000000000E+00 +000000000 +0000 +0000
SYSIBM   LUMODES           +0004 +0001 +0.000000000E+00 +000000000 +0000 +0000
SYSIBM   LUNAMES           +0009 +0001 +1.000000000E+00 +000000001 +0025 +0000
SYSIBM   MODESELECT        +0005 +0001 +0.000000000E+00 +000000000 +0000 +0000
SYSIBM   SYSAUXRELS        +0008 +0002 +1.100000000E+01 +000000001 +0025 +0000
SYSIBM   SYSCHECKDEP       +0005 +0001 +2.366000000E+03 +000000018 +0090 +0000
SYSIBM   SYSCHECKS         +0010 +0001 +2.362000000E+03 +000000040 +0100 +0000
SYSIBM   SYSCHECKS2        +0005 +0001 +2.323000000E+03 +000000026 +0092 +0000
SYSIBM   SYSCOLAUTH        +0015 +0002 +8.810000000E+02 +000000180 +0002 +0000
SYSIBM   SYSCOLDIST        +0012 +0001 +3.649600000E+04 +000000300 +0100 +0000
SYSIBM   SYSCOLDISTSTATS   +0014 +0001 +2.002000000E+03 +000000014 +0087 +0000
SYSIBM   SYSCOLDIST_HIST   +0011 +0001 +8.199900000E+04 +000000924 +0041 +0000
SYSIBM   SYSCOLSTATS       +0013 +0001 +8.160000000E+02 +000000013 +0081 +000
4B                  :00.4                                                02/17
```

Figure 31 Update Tables

The display shows the SYSTABLES statistics for the selected database and is the same display as the List Tables display as described in section 5.6.3

- When using the Update Object options where applicable, some fields are unprotected, allowing you to modify the values as required.
- Note that the modifications are not applied to the Statistics in the Warehouse until you select option 5.6.12 - Apply Updates to Statistics Warehouse.
- You can also use option 5.6.11 Display Updates Waiting to be to display the current updates you have entered for the current database before applying them.
- If you exit from the Display/Edit Statistics in Warehouse menu without applying the updates to the Warehouse, any updates you might have made will be discarded.
- Note that important relationships exist among certain columns in DB2 catalog tables. Consider these relationships if you choose to modify the catalog statistics. Refer to Correlations in Catalog for further information.
- See option 5.6.3 for a description of the table columns displayed.

Note that you can DRILL from this panel to the Table Columns and Indexes should you wish to view/update the Column and Index statistics including the SYSCOLDIST data etc.

5.6.9 Update Indexes

Select option 5.6.9 to display the Update Indexes panel as shown in the following figure:-

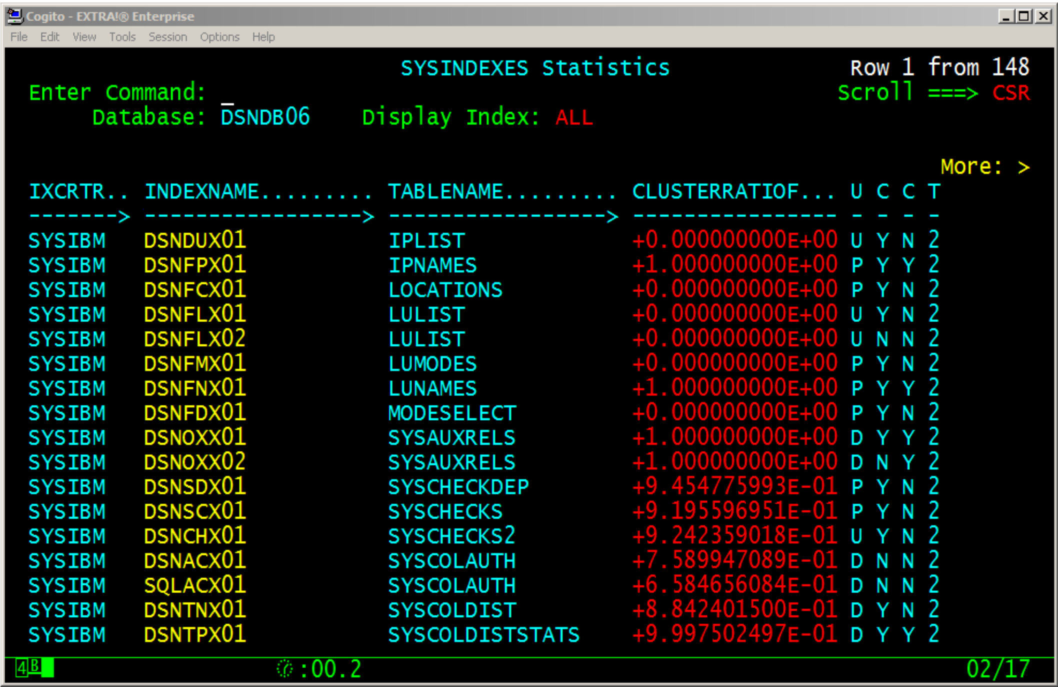

```
Cogito - EXTRA!® Enterprise
File  Edit  View  Tools  Session  Options  Help
                              SYSINDEXES Statistics              Row 1 from 148
Enter Command: _                                                 Scroll ===> CSR
     Database: DSNDB06     Display Index: ALL

                                                                        More: >
IXCRTR.. INDEXNAME......... TABLENAME......... CLUSTERRATIOF... U C C T
-------> -----------------> -----------------> ---------------- - - - -
SYSIBM   DSNDUX01           IPLIST             +0.000000000E+00 U Y N 2
SYSIBM   DSNFPX01           IPNAMES            +1.000000000E+00 P Y Y 2
SYSIBM   DSNFCX01           LOCATIONS          +0.000000000E+00 P Y N 2
SYSIBM   DSNFLX01           LULIST             +0.000000000E+00 U Y N 2
SYSIBM   DSNFLX02           LULIST             +0.000000000E+00 U N N 2
SYSIBM   DSNFMX01           LUMODES            +0.000000000E+00 P Y N 2
SYSIBM   DSNFNX01           LUNAMES            +1.000000000E+00 P Y Y 2
SYSIBM   DSNFDX01           MODESELECT         +0.000000000E+00 P Y N 2
SYSIBM   DSNOXX01           SYSAUXRELS         +1.000000000E+00 D Y Y 2
SYSIBM   DSNOXX02           SYSAUXRELS         +1.000000000E+00 D N Y 2
SYSIBM   DSNSDX01           SYSCHECKDEP        +9.454775993E-01 P Y N 2
SYSIBM   DSNSCX01           SYSCHECKS          +9.195596951E-01 P Y N 2
SYSIBM   DSNCHX01           SYSCHECKS2         +9.242359018E-01 U Y N 2
SYSIBM   DSNACX01           SYSCOLAUTH         +7.589947089E-01 D N N 2
SYSIBM   SQLACX01           SYSCOLAUTH         +6.584656084E-01 D N N 2
SYSIBM   DSNTNX01           SYSCOLDIST         +8.842401500E-01 D Y N 2
SYSIBM   DSNTPX01           SYSCOLDISTSTATS    +9.997502497E-01 D Y Y 2
4B                      :00.2                                           02/17
```

Figure 32 Update Indexes

The display shows the SYSINDEXES statistics for the selected database and is the same display as the List Indexes display as described in section 5.6.4

- When using the Update Object options where applicable, some fields are unprotected, allowing you to modify the values as required.
- Note that the modifications are not applied to the Statistics in the Warehouse until you select option 5.6.12 - Apply Updates to Statistics Warehouse.
- You can also use option 5.6.11 Display Updates Waiting to be to display the current updates you have entered for the current database before applying them.
- If you exit from the Display/Edit Statistics in Warehouse menu without applying the updates to the Warehouse, any updates you might have made will be discarded.
- Note that important relationships exist among certain columns in DB2 catalog tables. Consider these relationships if you choose to modify the catalog statistics. Refer to Correlations in Catalog for further information.
- See option 5.6.4 for a description of the columns displayed.

Note that you can DRILL from this panel to the Index Columns should you wish to view/update the Column and Index statistics including the SYSCOLDIST data etc.

5.6.10 Update SYSROUTINES

Select option 5.6.10 to display the Update SYSROUTINES panel as shown in the following figure:-

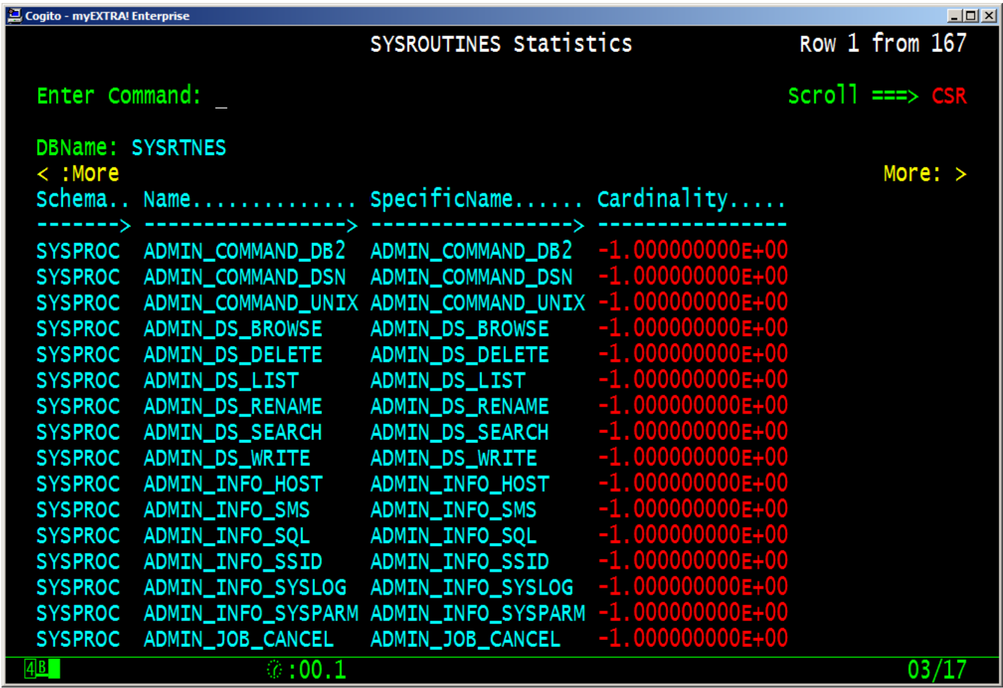

Figure 33 Update SYSROUTINES

The display shows the SYSROUTINES statistics and is the same display as the List SYSROUTINES display as described in section 5.6.5.

Note that to display/Edit the SYSROUTINES statistics, the database name must be set to SYSRTNES.

- When using the Update Object options where applicable, some fields are unprotected, allowing you to modify the values as required.
- Note that the modifications are not applied to the Statistics in the Warehouse until you select option 5.6.12 - Apply Updates to Statistics Warehouse.
- You can also use option 5.6.11 Display Updates Waiting to be to display the current updates you have entered for the current database before applying them.
- If you exit from the Display/Edit Statistics in Warehouse menu without applying the updates to the Warehouse, any updates you might have made will be discarded.
- Note that important relationships exist among certain columns in DB2 catalog tables. Consider these relationships if you choose to modify the catalog statistics. Refer to Correlations in Catalog for further information.

See option 5.6.4 for a description of the columns displayed.

5.6.11 Display Updates Waiting to be applied

You can Display Updates Waiting to be to the Warehouse by selecting option 5.6.11. The following panel is displayed:-

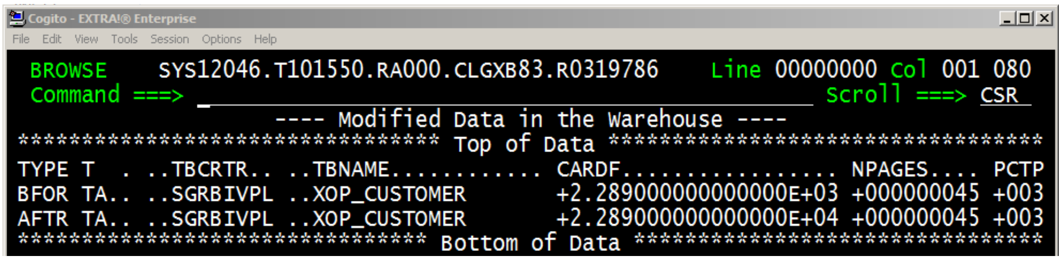

Figure 34 Display Current Updates

A line is displayed for each current update not yet applied to the Warehouse.

Note the following fields are displayed:-

TYPE

BFOR This is the BEFORE image of the record, i.e. This contains the Statistics Values as they are currently stored in the Warehouse.

AFTR This is the AFTER image of the record, i.e., This shows the modified values of the statistics, not yet applied to the Warehouse.

T

Identifies the statistics type as follows:-

TA	SYSTABLES
SA	SYSTABLESPACE
BA	SYSTABSTATS
IA	SYSINDEXES
CA	SYSCOLUMNS
DA	SYSCOLDIST
UA	SYSCOLSTATS
VA	SYSKEYTARGETS
WA	SYSKEYTGTDIST
XA	SYSROUTINES

After reviewing the changes you have made, you may choose to Apply Updates to Statistics Warehouse (option 5.6.12) or exit without saving the changes.

5.6.12 Apply Updates to Statistics Warehouse

Select option 5.6.12 Apply Updates to Statistics Warehouse to apply any changes you have made to the Statistics Warehouse.

You will receive a message -

Updates Applied Successfully -The Current Updates have been applied to the Warehouse.

The Updates you have made have been applied to the copy of the statistics in the Statistics Warehouse. NOTE: They have NOT been applied to the statistics in the system catalog. If you wish to apply the statistics to the system catalog, you must use option 5.7- Update Catalog Statistics from Warehouse.

Or

No Updates to Apply to the Warehouse - There are no current updates to apply to the Warehouse.

*If you use **EZ-Stats** to manually tweak statistics in the Statistics Warehouse, you should keep in mind that there are certain relationships amongst certain columns in the system catalog tables. If you modify values to create inconsistent data, then the DB2 Optimizer may ignore the statistics using default values instead, and give you no warning that it has done so. Refer to A.3 Correlations in the Catalog for further information.*

5.6.13 Display Applied Updates

If you have made edit changes to statistics values, and applied them to the Warehouse, you can view the changes you have made.

Select option 5.6.13 to Display Applied Updates as shown in the following figure:-

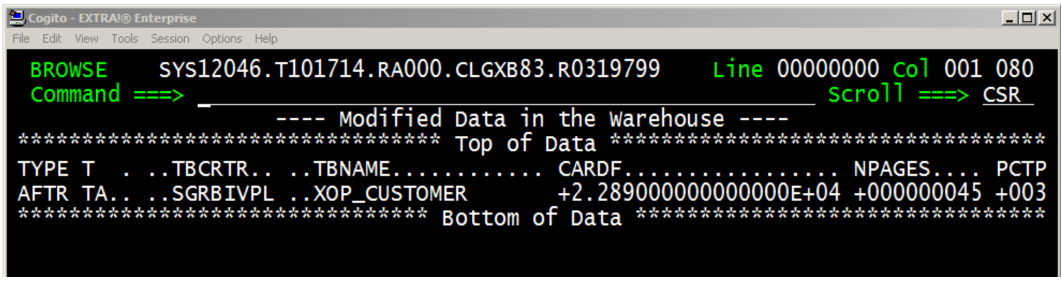

Figure 35 Display Applied Updates

The display shows a line for each statistic in the Warehouse for the current database that has been modified.

Note the following fields displayed:-

TYPE AFTR The record type will always be AFTR identifying that this is an AFTER image (ie. an applied update).

T Identifies the statistics type as follows:-

TA	SYSTABLES
SA	SYSTABLESPACE
BA	SYSTABSTATS
IA	SYSINDEXES
CA	SYSCOLUMNS
DA	SYSCOLDIST
UA	SYSCOLSTATS
VA	SYSKEYTARGETS
WA	SYSKEYTGTDIST
XA	SYSROUTINES

5.7 Update Catalog Statistics from Warehouse

Statistics for a particular database can be updated with statistics that have been added to the Statistics Warehouse. The database to be updated may be the *same or a different database*, on the same or a different DB2 subsystem, from which the statistics were originally extracted.

The database to be updated would be expected to be equivalent to the database from which the statistics were extracted. However, in typical Production/test/QA setups, it is normal for there to be some differences between Production and test/QA versions of the database, and these differences are handled automatically by the **EZ-Stats** update process.

In particular:-

- Database names may be different
- Owner/Creator qualifiers may be different
- Table names and Column Names would be expected to be the same

To cater for **SAP R/37** environments, the following differences are also supported:-

- Table space / Index space names may be different
- Different Column Ordering within a table
- Index Names may be different
- Index Keys may be ORDERED differently, e.g. ASC, DESC

Refer to the **SAP=YES** parameter and Section A.5.4- SAP R/37 Processing Considerations for further information about how **EZ-Stats** handles these **SAP** differences.

Select option 5.7 to display the Update Catalog Statistics from Warehouse panel as shown in the following figure:-

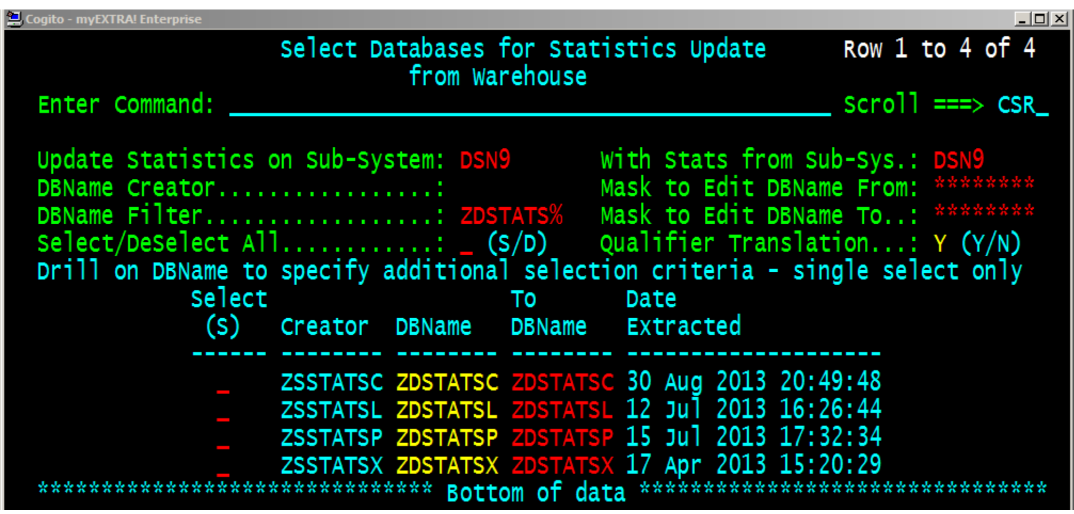

Figure 36 Update Catalog Statistics from Warehouse

The panel lists the Database names for the Statistics in the Warehouse FROM a particular subsystem and the date extracted. The display will list ALL of the databases in the Warehouse, unless you have filtered the display using the DBName Creator or DBName Filter described below. If the Date Extracted is followed by a comment *No Database, this indicates that the statistics for the relevant database are in the statistics warehouse, but the corresponding database no longer exists.

The following parameters may be specified on this panel.

Update Statistics On Sub-System	Specify the DB2 subsystem on which you wish to UPDATE the statistics from the Statistics warehouse.
DBName Creator	Optionally enter a Creator Name and only Databases created with that Creator will be listed for selection. You may filter the selection using "_"and/or "%" characters. For example, AB_D%.
DBName Filter	Enter a filter to reduce the Database name display to those that match the filter. You may filter the selection using "_"and/or "%" characters. For example, AB_D%. Enter % to display all database names that have been extracted to the statistics warehouse for the applicable FROM DB2 subsystem.
Select/DeSelect ALL	Enter **S** or **D** to Select or Deselect all Database names displayed.
With Stats from Subsys	Specify the DB2 subsystem FROM which the catalog statistics were added to the statistics warehouse.
Mask to Edit DBName From	If the database names on the system to be updated are different to the database names in the warehouse, you may use this parameter and Mask to Edit DBName To in order to perform a global edit of the FROM and TO DBNames. For example, Mask to Edit DBName From is **PR**** Mask to Edit DBName To is **QA**** will change all values PR in positions 3 and 4 to QA. See also To DBName described later.
Mask to Edit DBName To	See Mask to Edit DBName From.
Qualifier Translation	If the Databases that will have their statistics updated were created using different qualifiers to the Databases from which the statistics were extracted then enter Y and DRILL <**PF4**> to get a pop-up panel in which you can input the old and new values for object creators as shown in the following figure:-

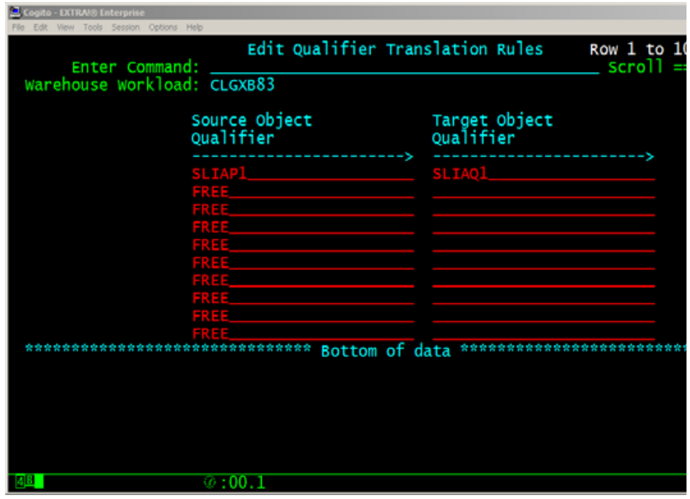

Figure 37 **Qualifier translation**

The qualifier translation panel will initially display a list of 10 FREE entries into which you can overtype the values for the Source and Target object qualifiers as required.

If you need more than 10 entries, press **<Enter>** and an additional 10 free slots will be displayed. Repeat until you have defined all of the qualifier translation requirements for the job.

Select (S)

Enter S to select this database to be updated.

See also Select/DeSelect ALL parameter.

Creator

The database Creator name.

DBName

The name of a database to be Updated with the statistics in the statistics warehouse.

All of the databases currently in the statistics warehouse will be listed, unless you have filtered by DBName or Creator. See DBName Creator and DBName Filter parameters above.

By default, ALL Tablespaces and Tables will be updated for the selected database(s). However, when selecting a single database, it is possible to DRILL on the database name to further specify the selection by either table space or table name. Multiple table spaces or tables may be specified for the selected database. See Selecting by table space and table.

To DBName

The name of the database to be UPDATED. You may overwrite this name if the name is different to the originating database that it is saved in the Warehouse.

See also Mask to Edit DBName from/To above.

*Note that you can SORT on DBName or Creator and also use the L(ocate) command for a database name. For example, **L DGRB** to position at the first database that commences with the characters DGRB.*

Date Extracted

The date/time when the statistics for this database were added to the statistics warehouse.

If the Date Extracted is followed by a comment *No Database, this indicates that the statistics for the relevant database are in the statistics warehouse, but the corresponding database no longer exists.

When you have identified the Sub-Systems and Databases, press **<Enter>** and a further selection panel will be presented to allow filtering of the Statistics to be updated as shown in the following figure.

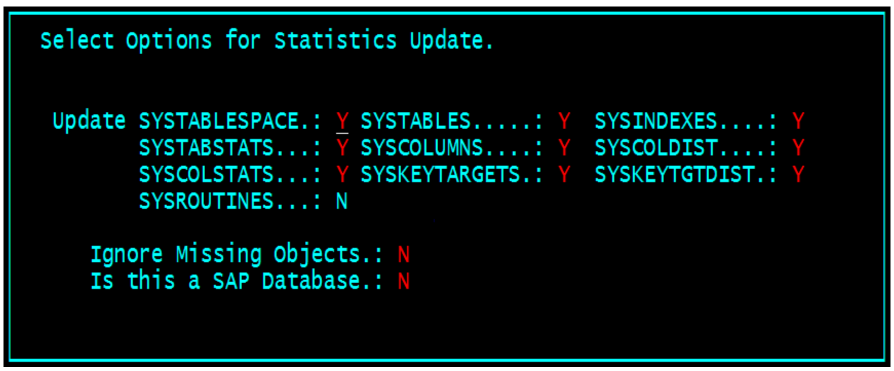

Figure 38 Select Options for Statistics Update

The following fields are displayed:-

Update SYSTABLESPACE

Specify Y to update the SYSTABLESPACE Statistics for the selected database(s) or Tablespace(s).

This option may only be selected if you are updating/ copying the entire database or selected Tablespace(s).

It may not be selected if you are filtering by Table names.

Update SYSTABLES

Specify Y to update the SYSTABLES Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSINDEXES

Specify Y to update the SYSINDEXES Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSTABSTATS

Specify Y to update the SYSTABSTATS Statistics for the selected database(s) or Tablespace(s).

This option may only be selected if you are updating/ copying the entire database or selected Tablespace(s).

It may not be selected if you are filtering by Table names.

Update SYSCOLUMNS

Specify Y to update the SYSCOLUMNS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

Update SYSCOLDIST

Specify Y to update the SYSCOLDIST Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSCOLSTATS

Specify Y to update the SYSCOLSTATS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSKEYTARGETS

Specify Y to update the SYSKEYTARGETS Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSKEYTGTDIST

Specify Y to update the SYSKEYGTDIST Statistics for the selected database(s), Tablespace(s) or Table(s).

This option always defaults to Y.

SYSROUTINES

Specify Y to update the SYSROUTINES statistics. This option may only be selected if you have specified a database name of SYSRTNES.

All of the other options will be set to N.

Ignore Missing Objects(Y/N)?

When applying the copied statistics TO the specified sub-system, if the Database(s) against which the statistics are to be applied do not exactly match the databases from which the statistics have been copied FROM, then any Objects missing will be reported.

For example a Tablespace may be defined as Partitioned in the FROM subsystem but not in the TO subsystem, which may be acceptable. However you should use this option with caution as missing Tables may cause SQL to fail to Explain and missing Indexes will cause different access paths to be generated.

Enter Y to ignore missing Objects and continue.
Enter N to terminate with RC 16.

Is this a SAP database (Y/N)?

Specify Yes or No to identify whether the Database to be updated is an **SAP** Database. **EZ-STATS** has some special processing logic to cater for name translation in **SAP** environments, where typically Tablespace names and even Index names vary between systems. The default value is N. Refer to the Section A.5.4 SAP R/37 Processing Considerations for further details.

To save a repeatable job stream that can be run at will, take copies of the files qualified with "userid.IO" to new dataset names of your choice and adjust the saved jcl with the new dataset names.

You may view the output from the Statistics Update job (XOPSTATU) in SDSF.

Refer to section A.5 in the Appendix for further information about XOPSTATU processing.

5.8 Export Catalog Statistics to another LPAR

If you wish to update statistics on a DB2 system that is located on another LPAR and there is no shared DASD between the LPARS, you will need to transfer the statistics from the Statistics Warehouse to that LPAR.

This transfer is not required if the DB2 system to be updated is on the same LPAR, or the Statistics Warehouse is allocated on shared DASD (accessible to the other LPAR).

Select option 5.8 - Export Catalog Statistics to another LPAR.

The file transfer may be performed using XMIT or FTP syntax, depending on the **EZ-DB2** Environmentals option selected for the current subsystem. Refer to the **EZ-DB2** Installation Guide for further details.

The following panel is displayed:-

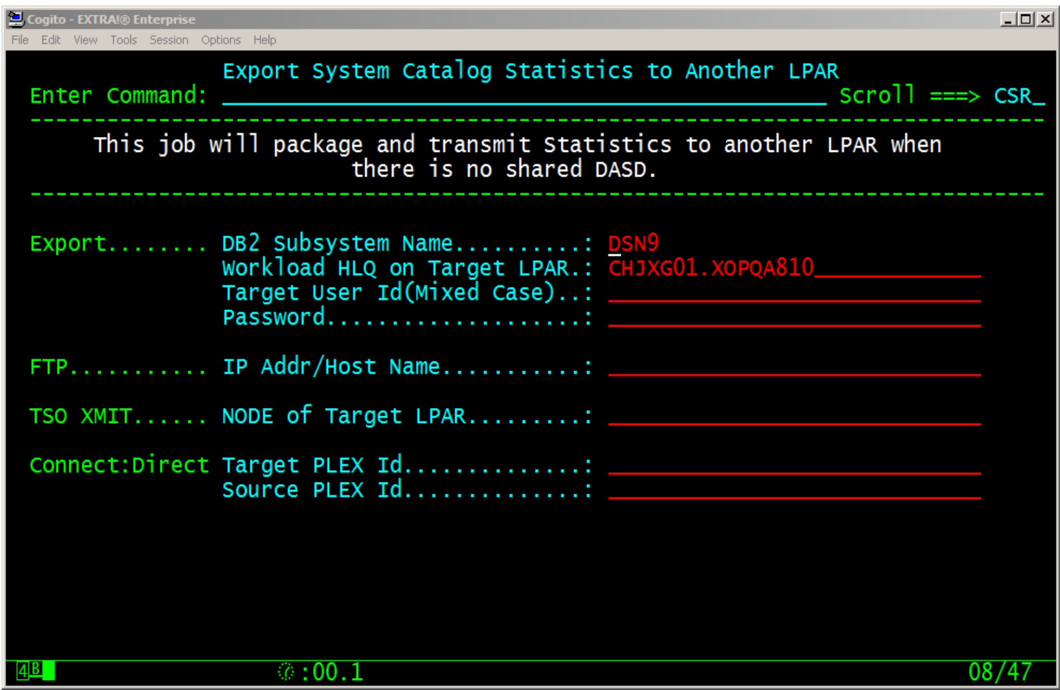

Figure 39 **Export Catalog Statistics to another LPAR**

The following fields must be specified on this panel.

Export

DB2 Subsystem Name	This is the name of the Subsystem whose Statistics you want to Export/Import to/from another LPAR.
Workload HLQ on Target LPAR	Set this field to the high level qualifier for the EZ-DB2 Workload files on the remote LPAR. This is used for FTP to 'PUT' the packaged Statistics file into 'tgtprfx.EXPORT.STATsubsys' or to specify the &FILEOUT1 parameter for Connect:Direct.

	Target Userid(Mixed case)	Set this field to the username to sign onto the remote FTP server, or to the Target User Id for the TSO XMIT. For mainframe FTP, this is normally the remote TSO userid.
	Password	Set this field to the password to use for signing onto the remote FTP server. For mainframe FTP, this is normally the TSO Userid's password. Note: This password is NOT stored within EZ-DB2.
FTP		
	IP Addr/Host Name	Set this field to the IP address or Host Name of the remote LPAR. This field is only required when using FTP for the transfer.
TSO XMIT		
	Node of Target LPAR	Set this field to the JES Node of the remote LPAR, if you wish to use XMIT to send the packaged data to the remote LPAR.
Connect:Direct		
	Target PLEX Id	Set this to the Id of the Target PLEX if you are using Connect:Direct to send the packaged statistics file. This value is substituted for the literal **dctgt** in the USRXMCSU CNTL member**.**
	Source PLEX Id	Set this to the Id of the Source PLEX if you are using Connect:Direct to send the packaged statistics file. This value is substituted for the literal **dctgt** in the USRXMCSU CNTL member.

Press the **<Enter>** key to generate the job stream to transmit the Statistics Warehouse dataset to the required LPAR.

X Type SUBMIT in the command line to submit the job for execution
X Press <**PF3**> function key to return to the **EZ-Stats** main menu.

After the job completes successfully, switch to the other LPAR and use Option 5.9 to receive the files on the other LPAR.

5.9 Import Catalog Statistics from another LPAR

Use this option to RECEIVE the files on another LPAR after you have exported them using option 5.8. Export Catalog Statistics to another LPAR

Select Option 5.9 Import Catalog Statistics from another LPAR to receive the files on the other LPAR. The following panel is displayed:-

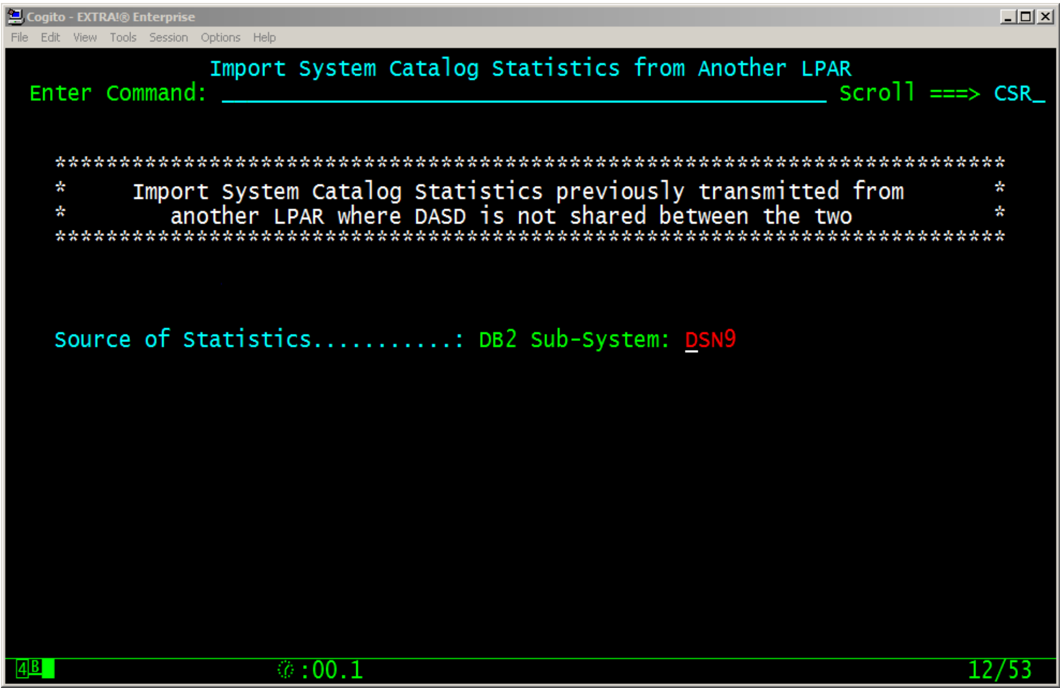

Figure 40 **Import Catalog Statistics from another LPAR**

The following fields must be specified on this panel.

DB2 Sub-system	The name of the DB2 Sub-System from which the Statistics were obtained.

Press the **<Enter>** key to generate the job stream to receive the Statistics Warehouse dataset to the required LPAR.

- X Type **SUBMIT** in the command line to submit the job for execution
- X Press <**PF3**> function key to return to the **EZ-Stats** main menu.

After the job completes successfully, you can use the normal **EZ-Stats** options, for example, to Copy Statistics from the Warehouse to DB2 subsystems on this LPAR.

5.10 Generate Statistics from RUNSTATS Reports ONLY

There are situations where a user is reluctant to run RUNSTATS in a production system. For example, EZ-XOP requires collection of ALL column statistics, not just index columns. This may not be the standard practice at some sites and the DBA may be unable or unwilling to change the RUNSTATS parameters in production to generate the additional statistics required by EZ-XOP. There is now a resolution to this problem.

An option has been added to EZ-STATS to populate an EZ-DB2 stats repository using the RUNSTATS REPORT ONLY option. Using this option, statistics are not updated on the DB2 catalog, thereby resolving any concern a DBA may have about updating production system catalog statistics.

A post process analyses the report data and creates a stats member in the statistics warehouse that looks exactly as if it had been created using the EZ-STATS option 5.5 Add catalog statistics to Warehouse.

This member may be used to update statistics in another subsystem. For example, by explicitly copying using EZ-STATS, or when creating a workload in the EZ-DB2 warehouse environment for processing by EZ-XOP.

This feature is dedicated to the memory of Gordon Froment of DST who was the inspiration for this idea.

Select Option 5.10 to display the Generate Statistics from RUNSTATS REPORTS ONLY menu.

Figure 41 Generate Statistics from RUNSTATS reports

5.10.1 Select Databases for RUNSTATS Report

Select option 5.10.1 to display the Select Databases for RUNSTATS Report menu as shown in the following figure:-

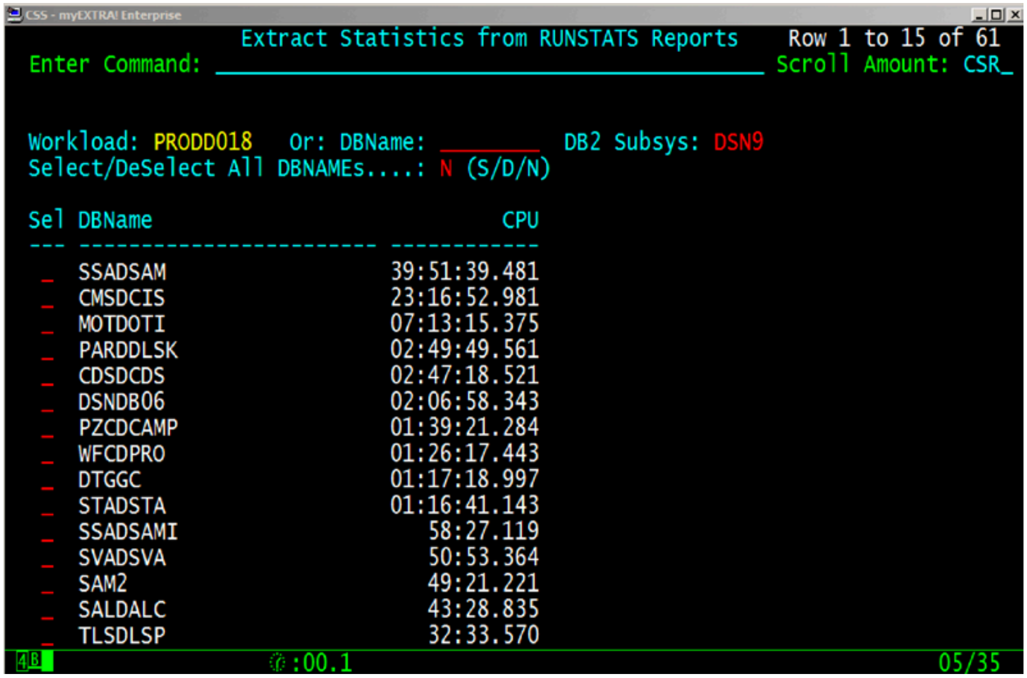

Figure 42 Select databases for RUNSTATS reports

Enter either a DBNAME and a DB2 subsystem, or a Workload Name.

If a Workload name is entered the panel will display a selection list of all of the databases associated with the selected workload.

Workload

Optionally, specify a Workload name. If Workload name is specified, the panel will generate a list of all the DBNames associated with the selected workload.

You may DRILL <**PF4**> on this field for a list of EZ-DB2 workloads.

DBName

If Workload name is not specified, specify a DBNAME and DB2 subsystem.

DB2 Subsys

If DBName is specified you must also specify the DB2 subsystem name,

Sel

If you have specified a workload name, a list of DBNames associated with the workload will be generate.

S — Specify S in this columns to Select the name of the database(s) required. Multiple databases may be selected.

DBName If you have specified a workload name, a list of databases will be generated.

CPU The cumulative CPU cost for SQL in the workload for the applicable database

Pressing enter will generate a job that allocates required datasets and generates the RUNSTATS input files.

Type <submit> in the command line to execute the job.

5.10.2 Edit RUNSTATS Control Parameters

Select option 5.10.2 to EDIT RUNSTATS Control Parameters. The following panel is displayed.

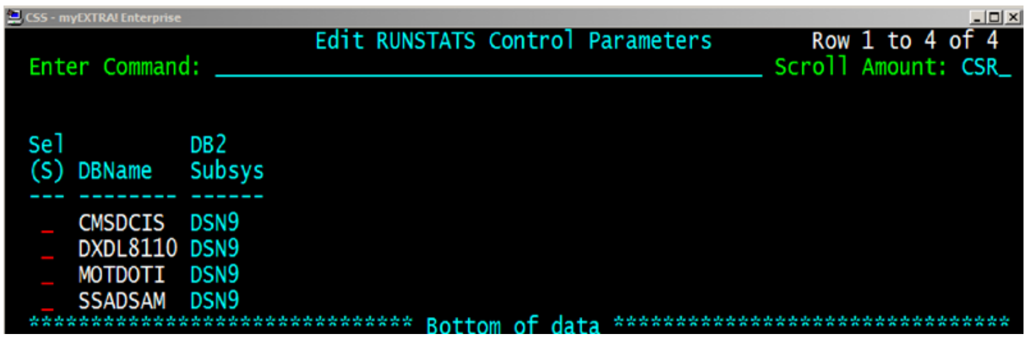

Figure 43 Edit RUNSTATS control parameters

The panel displays a list of previously selected databases.

Sel S Specify S in this columns to Select the name of the database(s) you wish to edit.

DBName The database name.

DB2 Subsys The DB2 subsystem name.

A TSO EDIT panel is displayed showing the RUNSTATS parameters for the selected database(s). You may edit these parameters as required. After reviewing the parameters and making any required changes, hit <**PF3**> to return to the EZ-Stats menu.

Note that the parameters default to REPORT YES UPDATE NONE. These values should NOT be changed.

5.10.3 Submit RUNSTATS

Select option 5.10.3 to generate the RUNSTATS job. The following panel is displayed:-

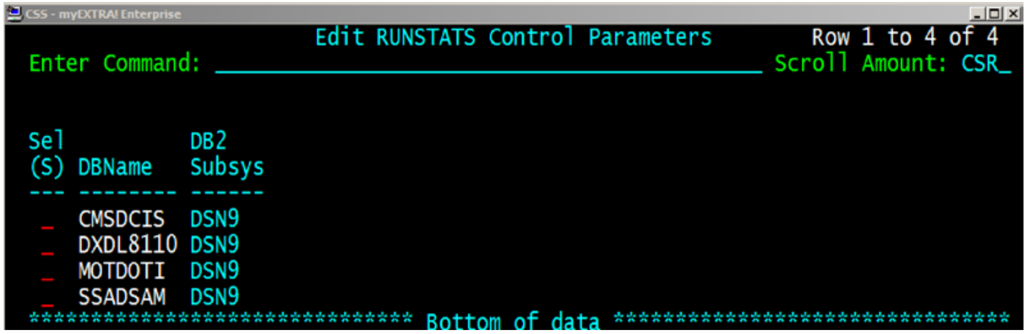

Figure 44 Submit RUNSTATS

The panel displays a list of previously selected databases.

Sel	S	Specify S in this columns to Select the name of the database(s) you wish to Submit.
DBName	The database name.	
DB2 Subsys	The DB2 subsystem name.	

The RUNSTATS JCL is displayed for the selected database(s). Enter <Submit> in the command line to execute the generated job.

The job stream executes RUNSTATS and then executes a post process routine that reads the RUNSTATS reports data and builds an EZ-Stats STATS member for each selected database in the EZ-Stats repository.

After execution, review the job output in SDSF.

5.10.4 Update Catalog Statistics from Reports

Select option 5.10.4 to update the DB2 catalog from the report data.

The panel and functionality is identical to option 5.7 – Update Catalog Statistics from Warehouse.

Refer to option 5.7 for further information.

5.11 Generate RUNSTATS from Statistics Feedback

Effective with DB2 11, the DB2 optimizer provides information about missing or conflicting statistics. Refer to the IBM DB2 11 Managing Performance Guide for further information.

This information is displayed by EZ-Tracer and EZ-Cache. For example, when viewing SQL detail, you can use the new **SF** command to display the Statistics feedback information for an SQL.

There is also a new option on the EZ-Tracer and EZ-Cache menus to generate the necessary RUNSTATS parameters for current EZ-DB2 Workload based upon the statistics feedback information.

If you are also licensed for the EZ-Stats component of EZ-DB2 then there is also an option to generate the global RUNSTATS parameters based upon the information in the SYSIBM.SYSSTATFEEDBACK table.

Select option 5.11 to generate RUNSTATS from statistics feedback.

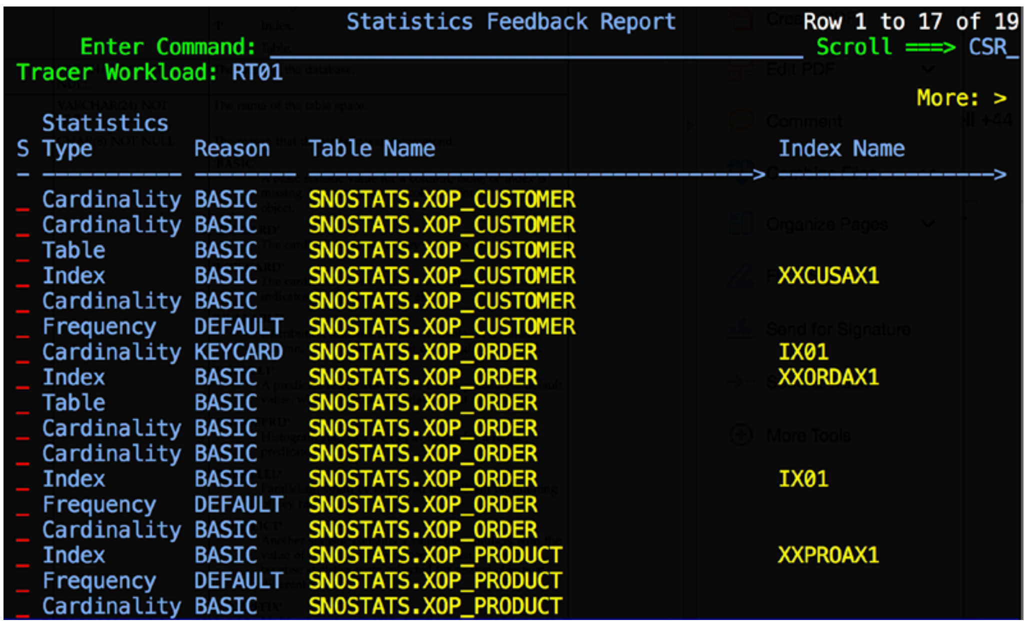

Figure 45 Statistics feedback report

The report shows the statistics feedback report based upon current information held in SYSIBM.SYSSTATFEEDBACK. Note that SYSIBM. SYSSTATFEEDBACK is updated whenever a BIND or Explain is performed when the optimizer finds missing or inconsistent statistics. It is also updated by RUNSTATS when the applicable rows are removed from SYSIBN.SYSSTATFEEDBACK.

From this display the user can generate the corresponding RUNSTATS parameters.

The Select (S) Column is used to generate the sample RUNSTATS parameters as follows:

S	The following values may be entered	
	S	Select a single line
	T	Select all lines for the applicable Table.

A Select ALL lines in the report

After selecting the required lines in the report hit <**enter**> and the following panel is displayed:

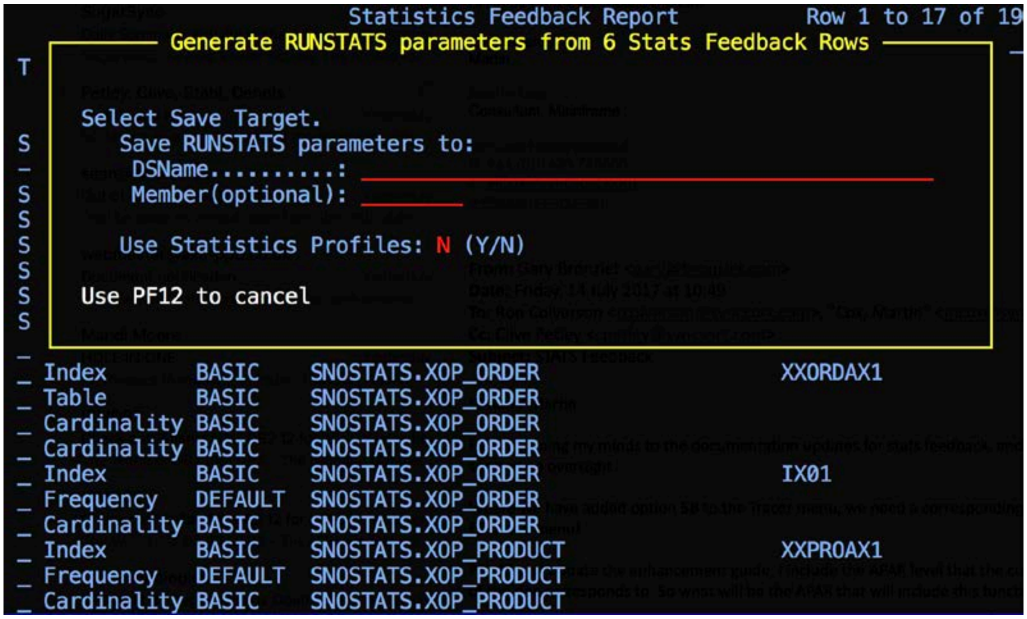

Figure 46 Generate RUNSTATS parameters from SYSIBM.SYSSTATFEEDBACK

Specify the destination for the RUNSTATS parameters. The following parameters may be specified:

DSName — Specify the dataset name where the RUNSTATS parameters should be saved.

Member(Optional) — Optionally specify a member name if DSName is a PDS.

Use Statistics Profiles — If a statistics profile exists for the selected table(s), you can specify whether the applicable statistics profile should be used when generating the RUNSTATS parameters. Refer to the IBM Managing Performance Guide for further information.

Y Use Statistics profile definition if it exists for the applicable table(s). The RUNSTATS parameters will include the syntax USE PROFILE for each table as applicable.

N Do not use Statistics Profile definition.

The generated RUNSTATS parameters will be saved in DSName (optional member name).

Appendix

A. EZ-Stats Considerations

A.1 Statistics used by Optimizer for Access Path Selection

DB2 uses certain values in the various catalog tables for access path selection. The definitive list of which tables and columns are used for access path selection is documented in the IBM manual *Managing Performance* (SC19-2978).

For your information the tables and columns used by the DB2 Optimizer for access path selection is described below:-

SYSCOLDIST

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
CARDF	Yes	Yes	Yes	Yes
COLGROUPCOLNO	Yes	Yes	Yes	Yes
COLVALUE	Yes	Yes	Yes	Yes
FREQUENCYF	Yes	Yes	Yes	Yes
HIGHVALUE	Yes	No	Yes	Yes
LOWVALUE	Yes	No	Yes	Yes
NUMCOLUMNS	Yes	Yes	Yes	Yes
TYPE	Yes	Yes	Yes	Yes
QUANTILENO	Yes	No	Yes	Yes

The existing SYSCOLDIST data is DELETED, and the new data INSERTED.

SYSCOLSTATS

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
COLCARD	Yes	Yes	Yes	Yes
COLCARDDATA	Yes	Yes	Yes	Yes
HIGHKEY	Yes	Yes	Yes	Yes
HIGH2KEY	Yes	Yes	No	Yes
LOWKEY	Yes	Yes	Yes	Yes
LOW2KEY	Yes	Yes	No	Yes
PARTITION	Yes	No	Yes	Yes

The existing SYSCOLSTATS data is DELETED and the new data INSERTED.

SYSCOLUMNS

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
COLCARDF	Yes	Yes	Yes	Yes
HIGH2KEY	Yes	Yes	Yes	Yes
LOW2KEY	Yes	Yes	Yes	Yes

The existing SYSCOLUMNS data is UPDATED for the above columns only.

SYSINDEXES

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
AVGKEYLEN	Yes	No	No	No
CLUSTERED	Yes	Yes	No	Yes
CLUSTERING	No	No	Yes	No
CLUSTERRATIOF	Yes	Yes	Yes	Yes
FIRSTKEYCARDF	Yes	Yes	Yes	Yes
FULLKEYCARDF	Yes	Yes	Yes	Yes
NLEAF	Yes	Yes	Yes	Yes
NLEVELS	Yes	Yes	Yes	Yes
SPACEF	Yes	Yes	No	Yes
DATAREPEATFACTORF	Yes	Yes	Yes	Yes

The existing SYSINDEXES data is UPDATED for the above columns only.

SYSKEYTARGETS

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
CARDF	Yes	No	Yes	Yes
HIGH2KEY	Yes	Yes	Yes	Yes
LOW2KEY	Yes	Yes	Yes	Yes
STATS_FORMAT	Yes	Yes	Yes	Yes

The existing SYSKEYTARGETS data is UPDATED for the above columns only.

SYSKEYTGTDIST

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
CARDF	Yes	Yes	Yes	Yes
KEYGROUPKEYNO	Yes	Yes	Yes	Yes
KEYVALUE	Yes	Yes	Yes	Yes
FREQUENCYF	Yes	Yes	Yes	Yes
HIGHVALUE	Yes	Yes	Yes	Yes
LOWVALUE	Yes	Yes	Yes	Yes
NUMKEYS	Yes	Yes	Yes	Yes
TYPE	Yes	Yes	Yes	Yes
QUANTILENO	Yes	Yes	Yes	Yes

The existing SYSKEYTGTDIST data is DELETED, and the new data INSERTED.

SYSROUTINES

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
CARDINALITY	No	Yes	Yes	Yes
INITIAL_INSTS	No	Yes	Yes	Yes
INITIAL_IOS	No	Yes	Yes	Yes
INSTS_PER_INVOC	No	Yes	Yes	Yes

IOS_PER_INVOC	No	Yes	Yes	Yes

The existing SYSROUTINES data is UPDATED for the above columns only.

SYSTABLES

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
AVGROWLEN	Yes	Yes	No	Yes
CARDF	Yes	Yes	Yes	Yes
EDPROC	No	No	Yes	No
NPAGES	Yes	Yes	Yes	Yes
NPAGESF	Yes	Yes	Yes	Yes
PCTPAGES	Yes	Yes	No	Yes
PCTROWCOMP	Yes	Yes	Yes	Yes
SPACEF	Yes	Yes	No	Yes

The existing SYSTABLES data is UPDATED for the above columns only.

SYSTABLESPACE

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
AVGROWLEN	Yes	No	No	No
NACTIVEF	Yes	Yes	Yes	Yes
SPACE	Yes	No	No	No
SPACEF	Yes	Yes	No	No

The existing SYSTABLESPACE data is UPDATED for the above columns only.

SYSTABSTATS

Column Name	Set by RUNSTATS	User Can Update	Used for Access Paths	Updated by EZ-STATS
CARDF	Yes	Yes	Yes	Yes
NACTIVE	Yes	Yes	No	Yes
NPAGES	Yes	Yes	Yes	Yes
PCTPAGES	Yes	Yes	No	Yes
PCTROWCOMP	Yes	Yes	No	Yes

The existing SYSTABSTATS data is DELETED, and the new data INSERTED.

A.2 Modelling your Production System

It should be noted that Binding or Explaining an application in a test or sand-box system using production statistics does not guarantee that identical access paths will be chosen by the DB2 Optimizer to the access path chosen when the same application is Explained or Bound in Production.

The access path should be similar, but there are other factors that affect the choice of access path, for example:-

- Different processor models.

- Different number of processors. (Differences in the number of processors can affect the degree of parallelism that is obtained.)

- Different buffer pool sizes.

- Different RID pool sizes.

- Mismatched data in SYSIBM.SYSCOLDIST.

  (This should not occur when using **EZ-Stats** to copy the statistics from production to the test system).

- Different service levels.

- Different values of optimization subsystem parameters, such as STARJOIN, NPGTHRSH, and PARAMDEG (MAX DEGREE on installation panel DSNTIP4).

Where possible, you should therefore attempt to ensure that your **EZ-DB2** analysis system matches the live system being analyzed as closely as possible to achieve the best possible results.

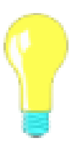

It is now possible to simulate these other factors that affect access path selection using the new DB2 functionality to enable environment virtualization. Using PROFILE definitions you may simulate BUFFER POOL, RID POOL and SORT POOL sizes. And new DSNZPARM parameters allow you to simulate CPU speed and CPU Count. An easy user interface to take advantage of this functionality is implemented in EZ-DB2 Warehouse. Refer to the ***EZ-DB2 SQL Warehouse User Guide*** *section* ***Environment Virtualization*** *for further information.*

Refer to the *DB2 Managing Performance Guide* (SC19-2078) section *Modeling a Production Environment on a test subsystem* for further information about factors affecting choice of access paths.

A.3 Correlations in the Catalog

If you use **EZ-Stats** to manually tweak statistics in the Statistics Warehouse, you should keep in mind that there are certain relationships amongst certain columns in the system catalog tables. If you modify values to create inconsistent data, then the DB2 Optimizer may ignore the statistics using default values instead, and give you no warning that it has done so.

The following relationships exist among certain columns of the catalog tables:

- Columns within table SYSCOLUMNS
- Columns in the tables SYSCOLUMNS and SYSINDEXES
- Columns in the tables SYSCOLUMNS and SYSCOLDIST
- Columns in the tables SYSCOLUMNS, SYSCOLDIST, and SYSINDEXES
- Columns with table space statistics and columns for partition-level statistics, as described in "Statistics for partitioned table spaces"

If you plan to update some values, keep in mind the following correlations:

COLCARDF and FIRSTKEYCARDF
For a column that is the first column of an index, those two values are equal. If the index has only that one column, the two values are also equal to the value of FULLKEYCARDF.

COLCARDF, LOW2KEY and HIGH2KEY
If the COLCARDF value is not '-1' or'-2', DB2 assumes that statistics exist for the column. In particular, it uses the values of LOW2KEY and HIGH2KEY in calculating filter factors. If COLDARDF = 1 or if COLCARDF = 2, DB2 uses HIGH2KEY and LOW2KEY as domain statistics, and generates frequencies on HIGH2KEY and LOW2KEY.

CARDF in SYSCOLDIST
CARDF is related to COLCARDF in SYSIBM.SYSCOLUMNS and to FIRSTKEYCARDF and FULLKEYCARDF in SYSIBM.SYSINDEXES. CARDF must be the minimum of the following values:

- A value between FIRSTKEYCARDF and FULLKEYCARDF if the index contains the same set of columns
- A value between MAX(COLCARDF of each column in the column group) and the product of multiplying together the COLCARDF of each column in the column group.

CARDF in SYSTABLES
CARDF must be equal to or larger than any of the other cardinalities,
Such as, COLCARDF, FIRSTKEYCARDF, FULLKEYCARDF, and CARDF in SYSIBM.SYSCOLDIST.

FREQUENCYF and COLCARDF or CARDF
The number of frequencies collected must be less than or equal to COLCARDF for the column or CARDF for the column group.

FREQUENCYF
The sum of frequencies collected for a column or column group must be less than or equal to 1.

Refer to the *DB2 Managing Performance Guide* section *Correlations in the Catalog* for further information.

A.4 XOPSTATD processing

XOPSTATD accesses the DB2 system catalog, reading all of the statistics used for access path selection for the selected database(s) and writing them to a member in the Statistics Warehouse or the Statistics Warehouse Archive depending on the option selected.

XOPSTATD works at the DATABASE level. This means that all of the optimizer data for a particular DATABASE will be written to a member in the Statistics Warehouse partitioned dataset (PDS) where the member name is the database name.

The program processes Catalog tables in the following order:

SYSIBM.SYSTABLESPACE
SYSIBM.SYSTABLES
SYSIBM.SYSINDEXES
SYSIBM.SYSCOLUMNS
SYSIBM.SYSCOLDIST
SYSIBM.SYSTABSTATS
SYSIBM.SYSCOLSTATS
SYSIBM.SYSKEYTARGETS
SYSIBM.SYSKEYTGTDIST
SYSIBM.SYSROUTINES

The SYSROUTINES statistics are only extracted if the database name is set to SYSRTNES. In which case, the SYSIBM.SYSROUTINES statistics will be written to a member in the statistics PDS called SYSRTNES. No other statistics will be extracted when the database name is SYSRTNES.

For information only purposes, XOPSTATD also dumps the following tables:-

SYSIBM.SYSTABLEPART
SYSIBM.SYSINDEXPART
SYSIBM.SYSINDEXSPACESTATS

However, no statistics from these last three tables is updated by XOPSTATU.

The following figure shows an example of the XOPSTATD processing LOGFILE written to SDSF. The report contains statistics for the number of rows read for each statistics type, for each database processed.

```
XOPPARM  11.35.21 : EZ-DB2 Release 9.1 APAR Level: CX910225
XOPSTATD 11.35.21 : R9.10 Starting Call Attach to DB2 Subsystem DSN9
XOPSTATD 11.35.21 : DB2 Version 10.1.0 CM9
XOPSTATD 11.35.32 : R9.10 XOPSTATD - Dump Catalog Statistics Starting -------
XOPSTATD 11.35.32 : Unloading database name ZDSTATSP
XOPSTATD 11.35.37 : XOPSTATD  :    225 Output Records Written
XOPSTATD 11.35.37 : XOPSTATD  :     03 Rows Read From SYSIBM.SYSTABLESPACE
XOPSTATD 11.35.37 : XOPSTATD  :     03 Rows Read From SYSIBM.SYSTABLES
XOPSTATD 11.35.37 : XOPSTATD  :     04 Rows Read From SYSIBM.SYSTABSTATS
XOPSTATD 11.35.37 : XOPSTATD  :     06 Rows Read From SYSIBM.SYSINDEXES
XOPSTATD 11.35.37 : XOPSTATD  :     17 Rows Read From SYSIBM.SYSCOLUMNS
XOPSTATD 11.35.37 : XOPSTATD  :     82 Rows Read From SYSIBM.SYSCOLDIST
XOPSTATD 11.35.37 : XOPSTATD  :     06 Rows Read From SYSIBM.SYSTABLEPART
XOPSTATD 11.35.37 : XOPSTATD  :     09 Rows Read From SYSIBM.SYSINDEXPART
XOPSTATD 11.35.37 : XOPSTATD  :     09 Rows Read From SYSIBM.SYSINDEXSPACESTATS
XOPSTATD 11.35.37 : XOPSTATD  :     24 Rows Read From SYSIBM.SYSCOLSTATS
XOPSTATD 11.35.37 : XOPSTATD  :     04 Rows Read From SYSIBM.SYSKEYTARGETS
XOPSTATD 11.35.37 : XOPSTATD  :     38 Rows Read From SYSIBM.SYSKEYTGTDIST
XOPSTATD 11.35.37 : XOPSTATD  :     00 Rows Read From SYSIBM.SYSROUTINES
XOPSTATD 11.35.37 : XOPSTATD  :    196 Rows Read in Total
XOPSTATD 11.35.37 : R9.10 Stopping Call Attach
XOPSTATD 11.35.37 : R9.10 XOPSTATD - Dump Catalog Statistics Completed ------
XOPSTATD 11.35.37 : 13/09/04 Ending            RETURN CODE - 0000
```

Figure 47 XOPSTATD Logfile

XOPSTATD return codes

RC	Description
0	Processing completed successfully.
4	No Databases Selected
8	One or more selected databases not found.
12	Invalid XOPPARMS parameter. See LOGFILE for details.
16	SQL error occurred. See LOGFILE for SQL diagnostics.
20	IO Error on file - check joblog for details
24	Dynamic allocation error. Check LOGFILE for Dynalloc error and information codes. The meaning of the Dynalloc error and information codes can be found in the **IBM** manual Authorized Assembler Services Guide. **Note**: Corresponding IKJ messages with further diagnostic information are written to JESMSGLG
28	CAF error. Check LOGFILE for details.
32	Unsupported DB2 version
36	Product authorization code invalid or expired.
99	Unexpected error.

A.5 XOPSTATU processing

XOPSTATU reads Catalog statistics saved in the Statistics Warehouse for one or more databases, and uses them to update the statistics on specified databases on specified DB2 sub-systems. The names of the databases and DB2 subsystem updated (the *target* database(s)) may be different from the original databases and DB2 subsystems from where the catalog statistics were extracted to the Warehouse (the *Source* databases).

XOPSTATU calls an internal routine XOPSQUAL to automatically translate the qualifiers as saved in the Statistics Warehouse as applicable for the database to be updated.

In the context of the following discussion, the database from which the statistics were extracted is referred to as the **A**Live database**@** and the database to be updated is referred to as the **A**Explain**@** database.

If XOPSQUAL qualifier translation is successful, XOPSTATU will proceed to update the catalog statistics on the database selected for update as appropriate. You will see messages similar to that shown in Figure 46.

The report displays the selected database names in the following format:-

```
XOPSQUAL 16.30.30 : Selected DBNAMEs:
XOPSQUAL 16.30.30 :
XOPSQUAL 16.30.30 : Live DBN Explain DBN
XOPSQUAL 16.30.30 : -------- ---------
XOPSQUAL 16.30.30 : DFSTP3NL DFXOP3NL
```

Where Live DBN and Explain DBN are the Statistics Warehouse database name and corresponding database to be updated as selected and specified in option 5.7- Update Catalog Statistics from Warehouse.

The report displays the qualifier translation for each applicable qualifier in the following format:-

```
XOPSQUAL 16.30.31 : Live      Explain     Live
XOPSQUAL 16.30.31 : Qualifier Qualifier DBName
XOPSQUAL 16.30.31 : --------- --------- --------
XOPSQUAL 16.30.31 : FSTP3     FXOP3     DFSTP3NL
XOPSQUAL 16.30.31 : FSTP3D    FXOP3D    DFSTP3NL
XOPSQUAL 16.30.31 : FSTP3RO   FXOP3RO   DFSTP3NL
XOPSQUAL 16.30.31 : FSTP3U1   FXOP3U1   DFSTP3NL
```

XOPSTATU will then update the catalog statistics as applicable on the database(s) selected for update.

The number of Catalog statistics updated is included in the report.

The job should normally terminate with return code zero.

```
XOPPARM  11.51.50 : EZ-DB2 Release 9.1 APAR Level: CX910225
XOPSTATU 11.51.50 : R9.10 Starting Call Attach to DB2 Subsystem DSN9
XOPSTATU 11.51.50 : DB2 Version 10.1.0 CM9
XOPSTATU 11.51.50 : R9.10 ---------- Update Catalog Statistics Starting -----
XOPSQUAL 11.51.50 : *******************************************************
XOPSQUAL 11.51.50 :  Starting Live to Image Qualifier translation
XOPSQUAL 11.51.50 : *******************************************************
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Selected DBNAMEs:
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Live DBN                 Image DBN
XOPSQUAL 11.51.50 : ------------------------ ------------------------
XOPSQUAL 11.51.50 : ZDSTATSP                 ZDSTATSC
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Existing Qualifiers from QUALS file:
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Live               Image              Live
XOPSQUAL 11.51.50 : Qualifier          Qualifier          DBName
XOPSQUAL 11.51.50 : ------------------ ------------------ ----------------------
XOPSQUAL 11.51.50 : ZSSTATSP           ZSSTATSC           ZDSTATSP
XOPSQUAL 11.51.50 : +++++++++++++++++++++++++++++++++++++++++++++++++++++++
XOPSQUAL 11.51.50 : Reading XOPCATS for DBNAME  ZDSTATSP
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Updated QUALS file output:
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : Live               Image              Live
XOPSQUAL 11.51.50 : Qualifier          Qualifier          DBName
XOPSQUAL 11.51.50 : ------------------ ------------------ ----------------------
XOPSQUAL 11.51.50 : ZSSTATSP           ZSSTATSC           ZDSTATSP
XOPSQUAL 11.51.50 :
XOPSQUAL 11.51.50 : *******************************************************
XOPSQUAL 11.51.50 :  End of Live to Image Qualifier translation
XOPSQUAL 11.51.50 : *******************************************************
XOPSQUAL 11.51.50 :
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : Tablespace Name Translation
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : Translating Old TSName: ZX01XPCU                 to New TSNa
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : Selected Live DBNAME: ZDSTATSP                 Image DBNAME:
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : Translating Old Qual: ZSSTATSP                 to New Qual :
XOPSTATU 11.51.50 : Selected Live DBNAME: ZDSTATSP                 Image DBNAME:
XOPSTATU 11.51.50 : ZDSTATSC STATS Updates Applied
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 :
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : Index Name Translation
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : ordered/Named: ZSSTATSP XXCUSAX1           To ZSSTATSC YYCUS
XOPSTATU 11.51.50 : ordered/Named: ZSSTATSP XXORDAX1           To ZSSTATSC YYORD
XOPSTATU 11.51.50 : Index Renamed: ZSSTATSP XXORDEX1           To ZSSTATSC YYORD
XOPSTATU 11.51.50 : Index Renamed: ZSSTATSP XXPARTX1           To ZSSTATSC YYPAR
XOPSTATU 11.51.50 : ordered/Named: ZSSTATSP XXPROAX1           To ZSSTATSC YYPRO
XOPSTATU 11.51.50 : Index Renamed: ZSSTATSP XXPROEX1           To ZSSTATSC YYPRO
XOPSTATU 11.51.50 : ------------------------------------------------------------
XOPSTATU 11.51.50 : XOPSTATU  :    226 Input Records Read
XOPSTATU 11.51.50 : XOPSTATU  :     44 Input Comment Records Read
XOPSTATU 11.51.50 : XOPSTATU  :     03 Rows Updated in  SYSIBM.SYSTABLESPACE
```

```
XOPSTATU 11.51.50 : XOPSTATU  :     03 Rows Updated in  SYSIBM.SYSTABLES
XOPSTATU 11.51.50 : XOPSTATU  :     04 Rows Deleted in  SYSIBM.SYSTABSTATS
XOPSTATU 11.51.50 : XOPSTATU  :     04 Rows Inserted in SYSIBM.SYSTABSTATS
XOPSTATU 11.51.50 : XOPSTATU  :     06 Rows Updated in  SYSIBM.SYSINDEXES
XOPSTATU 11.51.50 : XOPSTATU  :     17 Rows Updated in  SYSIBM.SYSCOLUMNS
XOPSTATU 11.51.50 : XOPSTATU  :     82 Rows Deleted in  SYSIBM.SYSCOLDIST
XOPSTATU 11.51.50 : XOPSTATU  :     82 Rows Inserted in SYSIBM.SYSCOLDIST
XOPSTATU 11.51.50 : XOPSTATU  :     24 Rows Deleted in  SYSIBM.SYSCOLSTATS
XOPSTATU 11.51.50 : XOPSTATU  :     24 Rows Inserted in SYSIBM.SYSCOLSTATS
XOPSTATU 11.51.50 : XOPSTATU  :     04 Rows Updated in  SYSIBM.SYSKEYTARGETS
XOPSTATU 11.51.50 : XOPSTATU  :     38 Rows Deleted in  SYSIBM.SYSKEYTGTDIST
XOPSTATU 11.51.50 : XOPSTATU  :     38 Rows Inserted in SYSIBM.SYSKEYTGTDIST
XOPSTATU 11.51.50 : XOPSTATU  :     00 Rows Updated in  SYSIBM.SYSROUTINES
XOPSTATU 11.51.50 : XOPSTATU  :     00 Indexes with no stats
XOPSTATU 11.51.50 : XOPSTATU  :     03 Total Indexes Renamed
XOPSTATU 11.51.50 : XOPSTATU  :     03 Total Indexes Reordered
XOPSTATU 11.51.50 : XOPSTATU  :    177 Rows Updated in Total
XOPSTATU 11.51.50 : XOPSTATU  :     00 Rows Excluded by Qualifier
XOPSTATU 11.51.50 : XOPSTATU  :     00 Total Missing Objects
XOPSTATU 11.51.50 : --------------------------------------------------
XOPSTATU 11.51.50 : R9.10 Stopping Call Attach
XOPSTATU 11.51.50 : R9.10 ---------- Update Catalog Statistics Completed ----
XOPSTATU 11.51.50 : 13/09/04 Ending           RETURN CODE - 0004
```

Figure 48 **XOPSTATU LOGFILE**

XOPSTATU Return Codes

RC	Description
00	Data Loaded successfully
04	No DBNAME selected or Some SYSTABSTATS missing or Some SYSTABLESPACE missing or Some Index Renamed/Reordered or Updates suppressed by XOPPARMS option.
08	No rows found for Database name used.
12	Invalid parameter string passed.
16	SQL Error (or Missing Objects and IGN_MISSING_OBJECTS set to NO). See Load Catalog Statistics – Missing Objects.
20	XOPCATS I/O error
24	Allocation error
28	CAF error
99	Unexpected RC.

A.5.1 Load Catalog Statistics – Missing Objects

It is intended that the database being updated is identical in definition to the corresponding database from which statistics were extracted to the Statistics warehouse. If that is not the case, the XOPSTATU LOGFILE may include a MISSING OBJECTS report as shown in the following figure, and the job will terminate with RC=16.

```
---------------------------------------------------------------------------------
EZX20009 16.30.33 :Objects Missing-Statistics cannot be loaded-
EZX20009 16.30.33 :Program will terminate with RC 16
EZX20009 16.30.33 :
---------------------------------------------------------------------------------
EZX20009 16.30.33 : Missing Object Report
EZX20009 16.30.33 : ---------------------
EZX20009 16.30.33 :
EZX20009 16.30.33 : Database Object       Object   Object              Owning   Table
EZX20009 16.30.33 : Name     Type         Creator  Name                Creator  Name
EZX20009 16.30.33 :
---------------------------------------------------------------------------------
EZX20009 16.30.33 : DFXOP3NL INDEX        FXOP3    ICPJ02                       TCPJBDET
EZX20009 16.30.33 : DFXOP3NL INDEX        FXOP3    ICPJ03                       TCPJBDET
EZX20009 16.30.33 : DFXOP3NL INDEX        FXOP3    ICPJ04                       TCPJBDET
EZX20009 16.30.35 :
---------------------------------------------------------------------------------
```

Figure 49 Load Catalog Statistics -Missing Objects Report

The report identifies the missing object(s) and object type for each database processed as applicable.

In the event that you receive missing object reports, you must correct the database definition of the database being updated and re-run option 5.7 Update Catalog Statistics from Warehouse .

Note: see also IGNORE Missing objects (Y/N) parameter. If you specify IGNORE Missing Objects = Y, then XOPSTATU will still produce the missing objects report but will end with RC=4.

If you specify Ignore Missing Objects, for missing Indexes, XOPSTATU will attempt to find equivalent indexes with different names, or with different column sequence as described in section A.5.6 XOPSTATU – SAP Processing.

A.5.2 Load Catalog Statistics Qualifier Translation Error

In some cases, it may not be possible for XOPSQUAL to unambiguously perform the qualifier translation.

In this case, the XOPSTATU LOGFILE will include messages such as the following:-

```
XOPSQUAL 15.25.45 : +++++++++++++++++++++++++++++++++++++++++++++++++++++++
XOPSQUAL 15.25.45 : Unable to uniquely identify qualifier for Tables:
XOPSQUAL 15.25.45 : +++++++++++++++++++++++++++++++++++++++++++++++++++++++
XOPSQUAL 15.25.45 : FSTP3U1 .PLAN_TABLE   has multiple qualifiers in Explain DB- DFXOP3NL
XOPSQUAL 15.25.45 : FSTP3RO .PLAN_TABLE   has multiple qualifiers in Explain DB- DFXOP3NL
XOPSQUAL 15.25.45 : +++++++++++++++++++++++++++++++++++++++++++++++++++++++
XOPSQUAL 15.25.45 :
XOPSQUAL 15.25.45 : Updated QUALS file output:
XOPSQUAL 15.25.45 :
XOPSQUAL 15.25.45 : Live      Explain   Live
XOPSQUAL 15.25.45 : Qualifier Qualifier DBName
XOPSQUAL 15.25.45 : --------- --------- --------
XOPSQUAL 15.25.45 : FSTP3     FXOP3     DFSTP3NL
XOPSQUAL 15.25.45 : FSTP3D    FXOP3D    DFSTP3NL
XOPSQUAL 15.25.45 : FSTP3RO   ????????  DFSTP3NL
XOPSQUAL 15.25.45 : FSTP3U1   ????????  DFSTP3NL
XOPSQUAL 15.25.45 :
XOPSQUAL 15.25.45 : *******************************************************
XOPSQUAL 15.25.45 :  End of Live to Explain Qualifier translation
```

Figure 50 XOPSQUAL - Qualifier Translation Failure

Typically, errors such as this would occur for example if

- The database definition of the database being updated is not in synch with the source database definition, e.g., there are missing tables and / or indexes.
- A table exists in the Target database with multiple qualifiers. XOPSTATU would not know which qualifier to match.

When such errors occur, the Job step will terminate with a Return Code of 99.

In this case, you must use the option Qualifier Translation (option 5.7) to explicitly set the Target Qualifier names as applicable before re-executing the update current statistics step.

A.5.3 Load Catalog Statistics - SAP Databases

When loading the Catalog Statistics for an **SAP** database (Refer to Update Current Statistics option **SAP** = YES), the XOPSTATU process includes special provision for handling unavoidable inconsistencies between the **SAP** source database and the target database being updated.

Refer to SAP R/3**7** Processing Considerations- section A.5.4 for further information.

A.5.4 SAP R/3® Processing Considerations

EZ-DB2 includes special **SAP** specific processing to deal with the fact that it is impossible to create an **SAP** Explain database in the **EZ-DB2** analysis system that is a true clone of the equivalent Live **SAP** database. **SAP** users will be aware that it is not possible to guarantee that a table will be in the same Tablespace in the Explain environment as it is in the Live. Neither is it possible to guarantee that Index names will be the same in the Explain as in the Live environment.

This would normally cause a problem for **EZ-DB2** when migrating the catalog statistics from the Live to the Analysis environment. To cope with this, a special **SAP** WORKLOAD parameter has been introduced in the Workload Trace Parameters. When **SAP** WORKLOAD is set to YES, some special **SAP** specific processing is performed when migrating the catalog statistics as described in the following section.

A.5.5 XOPSTATD – SAP Processing

XOPSTATD (Dump Catalog Statistics) executes as part of the **EZ-Stats** product and also as part of the Workload Trace process in EZ-Tracer. XOPSTATD writes the catalog statistics for each applicable database to a CATSTATS file member.

When **SAP** WORKLOAD = YES, the CATSTATS file will include, for each index, some specially encoded information allowing the indexes to be recognised by the column names and their ordering. This will allow XOPSTATU to identify the equivalent indexes in the analysis system even when the index names do not match.

A.5.6 XOPSTATU – SAP Processing

XOPSTATU (Update Catalog Statistics) executes as part of **EZ-Stats** and also the Import/Update SQL Workloads process of the EZ-SQL Warehouse manager.

XOPSTATU reads the CATSTATS member for each applicable database and attempts to update the catalog statistics as applicable for the Explain database. When processing **SAP** databases, XOPSTATU will handle NOT FOUND conditions for various Catalog entities as follows:-

SYSTABLESPACE.

When **SAP** WORKLOAD=YES, if the update returns ANot Found@, then an alternative SQL statement will be executed to update SYSTABLESPACE on CREATOR and NAME instead of DBNAME and NAME resulting in a Tablespace scan.

If the Tablespace name does not exist, then the missing Tablespace name will be reported to the log file. For missing Table space the job will terminate with RC=4.

SYSTABSTATS.

When **SAP** WORKLOAD =YES, XOPSTATU will not include DBNAME in the SQL predicates, so it will update the appropriate entry if it is exists. If the update fails it will be reported in the normal way.

Missing TABSTATS will terminate with RC=4.

SYSINDEXES

When processing SYSINDEXES, if a particular index is not found on the Explain database, XOPSTATU will attempt to identify an equivalent index (same column names and ordering) using the encoded information stored by XOPSTATD, and if found, will report the index name change and apply the statistics. If a match is not found, it will attempt to identify an index with the same column names in a different sequence, and if successful, will report the re-named and re-ordered index and apply the statistics.

If the index is not found, XOPSTATU will report the missing index.

Renamed Indexes, the job will terminate with RC=0. Re-ordered indexes the Job will terminate with RC=4. Missing Indexes, the job will terminate with RC=16.

New Indexes on Explain database

For each table in CATSTATS, XOPSTATU will identify all indexes for the table with missing statistics and report them to the log. If any indexes are found with missing statistics the job will terminate with RC=16.

The following figure shows an example of the possible LOGFILE messages written by XOPSTATU when processing an **SAP** database:-

```
The following errors/warnings are due to running with SAP=YES

----------------------------------------------------------------------------------------------
*Severity *    Error/Warning Type: Creator  Name
----------------------------------------------------------------------------------------------
*Warning  *    Tablespace Missing: SXQC4114 SX01XPCU
*Warning  *    Tablespace Missing: SXQC4114 SX01XPOD
*Warning  *    Tablespace Missing: SXQC4114 SX01XPPD
*Warning  *         Index Missing: SXQC4114 XXCUSAX1
*Warning  * Index Reordered/Named: SXQC4114 XXCUSAX2          TO SXQC4114 XXCUSAY2
*Error    *    Index Has No Stats: SXQC4114 XXCUSAX3
*Warning  *         Index Missing: SXQC4114 XXORDAX1
*Warning  *         Index Missing: SXQC4114 XXPROAX1
*Warning  * Index Reordered/Named: SXQC4114 XXPROAX2          TO SXQC4114 XXPROAY2
----------------------------------------------------------------------------------------------
XOPSTATU  :     01 Indexes with no stats
XOPSTATU  :     03 Total Indexes Renamed
XOPSTATU  :     02 Total Indexes Reordered
```

Figure 51 XOPSTATU messages possible when running with SAP = YES

The following describes the messages included in the XOPSTATU LOGFILE for **SAP**

Tablespace Missing: Signifies that XOPSTATU was unable to update SYSTABLESPACE due to missing TABLE SPACES in the Explain database. This is a WARNING message. The Job will terminate with RC=4.

Index Missing: Signifies that XOPSTATU was not able to find a corresponding index on the Explain database. This is an ERROR message. The job will terminate with RC=16.

Index Reordered/Renamed:
Index Renamed indicates XOPSTATU found a match on the **A**ordered**@** hash key passed from XOPSTATD, so we assume this is really the same index, and update the new named index instead. This is a WARNING message. The job will terminate with RC=0.

Index Reordered indicates XOPSTATU found a match on the **A**unordered**@** hash key passed from XOPSTATD. This means we have found a similar index ie. Same Columns but different ordering. In this case we update the new named index. This is a WARNING message. The job will terminate with RC=4.

Index has no Stats: Signifies that XOPSTATU found indexes on the Explain environment with no equivalent on the Live environment, and was therefore unable to load any statistics. This is an ERROR message. The Job will terminate with RC=16.

If there are any missing indexes in the Explain database, or indexes in the Explain for which we have no statistics, then the user will have to decide how to proceed.

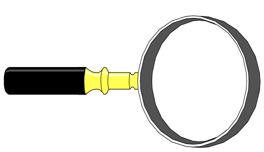

Index Missing: *Indexes missing in Explain database can be created manually and XOPSTATU re-run to apply the statistics.*

Index has no Stats: *If there is an Explain database Index with no statistics then it could be dropped if not Primary, Clustering, Foreign or the only index on the table. If Foreign, then the dependency could be removed.*

If it can't be dropped, you will have to generate statistics for it through RUNSTATS after the table is populated with sample data. (This assumes the table gets selected for Index Optimization), and live with the fact the statistics are inappropriate. Otherwise let the DB2 *Optimizer work on default statistics for Original Costs and Impact Analyzer.*