precisely

Syncsort Optimize DB2
(EZ-DB2)

SQL Warehouse User Guide
Release 9.10

Version 6.12

July 2014

Syncsort DB2 SQL Warehouse User Guide
(formerly EZ-DB2 SQL Warehouse User Guide)

NOTICE

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.

TABLE OF CONTENTS

FIGURES

New Features

Multiple Warehouse Databases

It is now permitted to install multiple EZ-DB2 Warehouses. Previously, only one EZ-DB2 Warehouse could be installed in the DB2 subsystem designated as the EZ-DB2 analysis environment.

You may still only have one EZ-DB2 warehouse per subsystem, but you can now have different warehouses in different environments.

This means for example, that you can now install a warehouse database in a production environment, and perform index analysis directly in the production environment.

When creating an EZ-DB2 Warehouse, the user must specify whether the environment is a production or development environment, as certain functions are prohibited in production.

For example,

- EZ-XOP (Index Optimizer) may not be executed in a production environment.
- When importing SQL into a warehouse database in a production environment, the option to “use live statistics” is suppressed.

When activating EZ-DB2 for a particular sub system, the user now specifies whether the subsystem is Production, QA or test.

Refer to the **EZ-DB2 Installation Reference Guide** for further information.

Importing Warehouse Workloads into Warehouse Workloads

It is now possible to choose from Tracer, Impact Analyzer and Warehouse workloads when importing SQL data into a Warehouse workload.

For example, you might want to use filters to create a subset of an existing warehouse workload.

Refer to the Impact Analyzer/Tracer/Warehouse (I/T/W) parameter in Option 6.1 – Select and Import SQL Traces for further information.

Enhanced Filtering when populating EZ-DB2 Warehouse Workload

When loading a tracer workload into an EZ-DB2 Warehouse workload, the user may now filter to include only selected Plans and/or Packages. All of the filtering has been moved to a new screen.

Refer to the **Filter** parameter in Workload Analysis and SQL and Statistics Controls for further information.

The ability to filter at this level has a number of uses. For example, if using EZ-Impact Analyzer against a Warehouse workload, you may wish to analyze selected plans and packages, while also taking into account execution frequencies from live workloads.

SQL Text and Access Path after Bind

There is a new option to display SQL detail information after the BIND of an application.

Refer to option 6.7 - SQL Text and Access Path after Bind for further information.

Adjust SQL Frequency of Execution and Exclude SQL

The Adjust SQL Frequency panel has been enhanced to include additional columns to the right, specifically - Total Clock, Total CPU, Total GetPages and Total Timerons. Page Right <PF11> to see the additional fields.

This allows the user to sort by these additional columns when selecting which SQL to modify and / or exclude from further analysis.

In addition, the user can now enter D or X in the delete column. For example, enter

D To delete a particular SQL from the workload. As before.

X New! To delete all SQL to the end of the display.

This makes it easier for a user to remove the lowest cost SQL from the analysis. Rather than having to enter D against each SQL that is no longer required the user can sort by Total CPU (for example) and then enter X to exclude all SQL from a particular SQL onwards.

Refer to option 6.5 for further information.

Environment Virtualization

EZ-DB2 now exploits DB2 capability to virtualize a DB2 production environment in test.

For example, for access path testing in a test or sand-box environment, it is common practise to copy production statistics to a *test* or *sand-box* environment. However, there are other factors that influence access path selection. These are:-

- CPU Speed
- Number of CPUS
- Buffer Pool Sizes
- RID Pool Size
- Sort Pool Size

and various DSNZPARM settings, i.e.

- NPGTHRSH
- PARAMDEG
- STARJOIN

Through the use of new DSNZPARM options and PROFILEs, it is now possible to simulate all of these values.

Environment Virtualization is implemented in EZ-DB2 as part of the EZ-SQL Warehouse Manager.

Refer to option 6.8 for further information.

Importing CA Detector® data into Tracer and Warehouse Workloads

It is now possible to create an EZ-DB2 Tracer and EZ-SQL Warehouse workload from SQL captured by the CA Detector® product.

Therefore, it is now possible for users of the CA Detector® product to take advantage of the Index Analyzer and Index Optimizer features of EZ-DB2.

In addition, when viewing the CA Detector data in EZ-Tracer and EZ-SQL Warehouse, CA Detector® users can benefit from the unique SQL consolidation functionality of EZ-DB2.

Refer to CA Detector® documentation for information about externalising Detector® data.

Refer to the Section Importing CA Detector data for information about importing the externalized CA Detector® data into a EZ-DB2 workloads.

CA Detector® is a program product of CA Technologies Inc.

Toggle between date created/date last updated in workload displays

The workload selection panels in all products have been updated to allow toggle between date created and date last updated.

```
EZ-Tracer Workload Manager                                Row 167 to 177 of 331
  Enter Command: ____________________________________________ Scroll ===> CSR_
Trace  Workload: IVP                          List Archived Workloads: N (Y/N)

Sel(S) Trace    Last                    Show Date Created/Last Updated: L (C/L)
Del(D) Workload Updated  Description - Update if Required
------ -------- -------- ------------------------------------------------------
  __    IVP      14/07/14 IVP Performance Trace_____________________________
  __    IVPTEST  12/07/16 Cache Trace_______________________________________
  __    JCAPAR38 14/04/07 Cache Trace_______________________________________
  __    JCCOLLID 14/04/14 Cache Trace_______________________________________
  __    JCDST    14/01/15 Cache Trace_______________________________________
  __    JCSETID  13/02/17 Cache Trace_______________________________________
  __    JCWAIT   13/02/04 Cache Trace_______________________________________
  __    JC01     14/02/18 Cache Trace_______________________________________
  __    JC02     14/01/27 Cache Trace_______________________________________
  __    JTAPAR38 14/03/05 Performance Trace_________________________________
  __    JTBGWAIT 14/03/01 Performance Trace_________________________________
```

Figure 1 **Workload Manager panel**

A new option has been added to the panel whenever a workload selection panel is displayed. The user can switch the List Mode to display Date Created or Date Last Updated.

Show Date Created/Last Updated	L	Show date workload last Updated
	C	Show date workload created

The setting will be remembered as the default for the current user.

After selecting the List Mode, the User can SORT on the date field, by positioning the cursor in the field and pressing the <**ENTER**> key.

Filtering of Workload Lists

The workload selection panels in all products have been updated to support the specification of wild-cards to filter the list of workload names displayed. This is particularly useful for sites that have large numbers of workloads.
You may enter % to match remaining characters and underscore to match any character in that position. For example, enter Z_A% in workload name field to display following list of workloads.

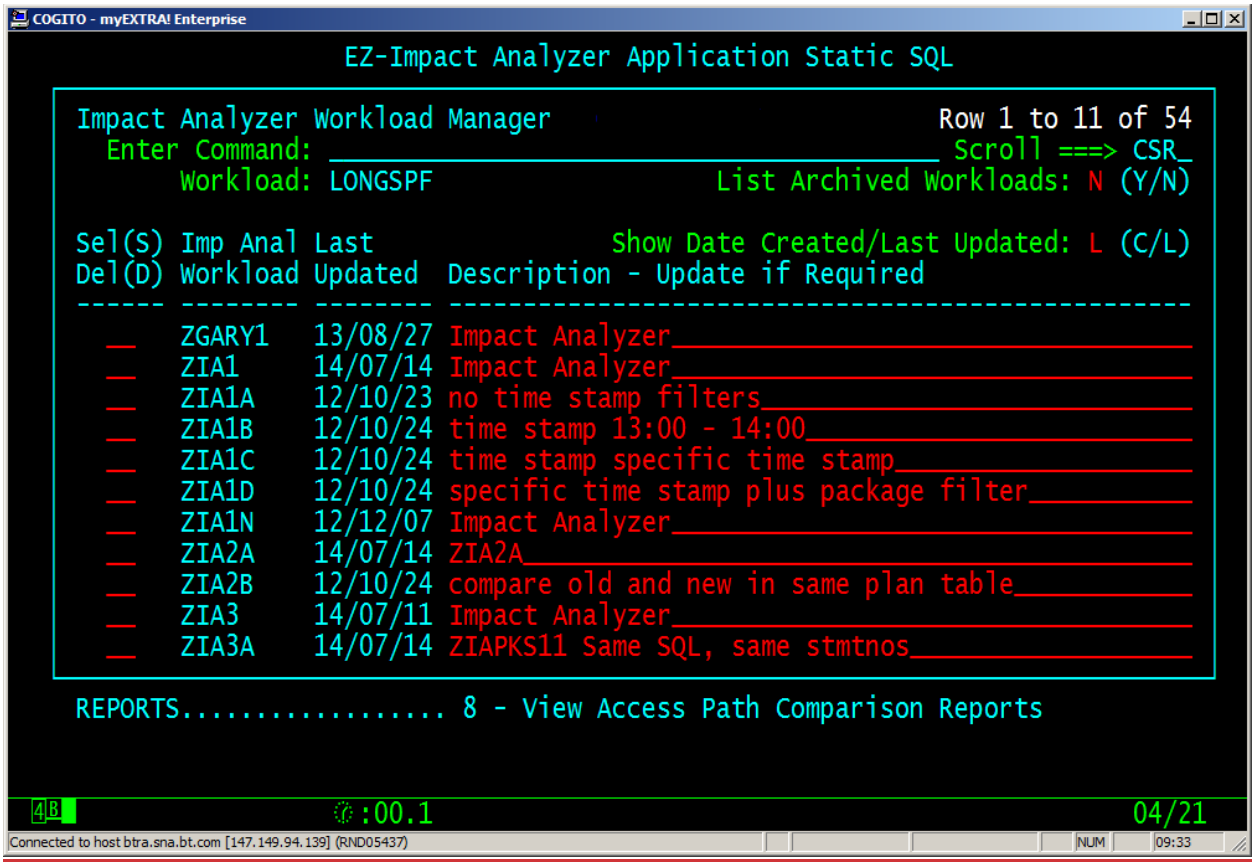

Figure 2 **Filter List of Workload names**

Archiving Workloads

It is possible to remove workload names from the list of workloads displayed.

A new option has been added to the panel whenever a workload selection panel is displayed. The user can switch the List Mode to display Active or Archived Workloads.

List Archived Workloads Specify whether you wish to display the active or the archived workloads.

Y List the archived workloads. If you have specified a MASK, only the workloads matching the MASK will be displayed.

N List the active workloads. If you have specified a MASK, only the workloads matching the MASK will be displayed.

A new command has been added to specify that workloads are to be moved to the archive list.

Sel(S)
Del(D)

A - Archive the Workload. The Workload data is retained but the name is removed from the active workload list. To view the workload names that have been archived, change List Archived Workloads to Y.

R - Recall Workload. If a workload has not been referenced for a while, it's possible that the applicable datasets have been archived by DFSMSHSM. Specify R to send a Recall request to DFSMSHSM for all of the workload datasets,

If the workload has been archived it will also be restored from the list of archived workloads.

Note: You can use block commands, **AA**, **DD** and **RR** to archive, delete and restore multiple workloads.

6 EZ-SQL Warehouse Manager

The purpose of the SQL Warehouse manager is to allow you to create and manipulate SQL Warehouse *Analysis Workloads* for processing by EZ-DB2 components, such as EZ-Index Analyzer and EZ-XOP Index Optimizer.

You may create an SQL *Warehouse Workload*

- By extracting SQL from the system catalog for one or more applications
- You may also import SQL Trace Files from one or more *Trace Workloads* created by the EZ-Tracer or EZ-Cache components.

You may enhance an SQL Workload by

- Loading additional *Free-form* SQL text

  This option allows you to perform EZ-DB2 analysis against applications that might not yet exist. For example, before the application is developed you could enter the proposed SQL in Free-form into an SQL Warehouse Workload, and manually set the anticipated/projected SQL frequency

- Refining SQL Workload frequencies

  You may manually override the SQL frequency in the *Warehouse Workload*, thereby allowing you to explicitly set the frequency of SQL loaded from the System Catalog or Free-Form SQL , and also to modify the frequency of traced applications, thereby allowing you to model the effect of anticipated Workload growth.

Note that the names of *Trace Workloads* and SQL *Warehouse Workloads* MUST be unique. This allows a *many- to-many* relationship, enabling you to load SQL Trace Files from many *Trace Workloads* into a single SQL *Warehouse Workload*, or SQL Trace Files from a *Trace Workload* into many SQL *Warehouse Workloads*.

Throughout this document, the term *Warehouse Workload* is synonymous with *Analysis Warehouse Workload* or *SQL Warehouse Workload.* These names all imply Workloads created in the Warehouse (Analysis) environment as opposed to *Trace Workloads* that are created by the EZ-Tracer or EZ-Cache components.

Select Option 6 to display the SQL Warehouse Manager menu as shown in the following figure:-

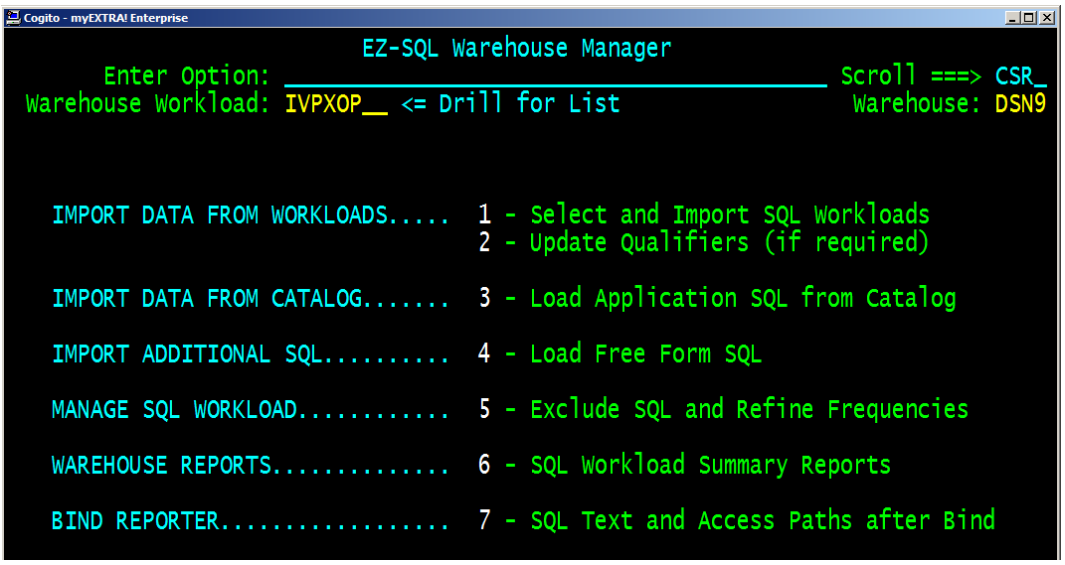

Figure 3 SQL Warehouse Manager

Note that in a multi-warehouse environment, the current Warehouse environment DB2 subsystem is shown on this panel. In the example above (Figure 3), the current Warehouse is installed on the DB2 subsystem DSN9.

Refer to the EZ-DB2 Installation Reference Guide for information about installing multiple warehouse databases.

You may DRILL <PF4> on the subsystem name to display a pop-up window showing all of the DB2 subsystems where an EZ-DB2 Warehouse is installed. This is shown in the following figure:-

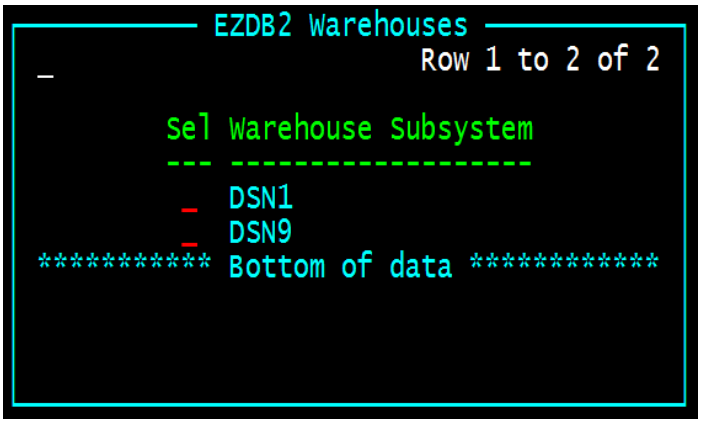

Figure 4 Multiple EZ-DB2 Warehouses

To change the current EZ-DB2 Warehouse system, enter S in the Sel column for the applicable Warehouse subsystem.

The creation of the *SQL Warehouse workload* is now described in the following sections.

Create/Select Warehouse Workload

To select a Warehouse Workload to work with, enter the name in the Warehouse Workload name field. If you enter a name that doesn't currently exist, the following pop-up panel will be displayed:-

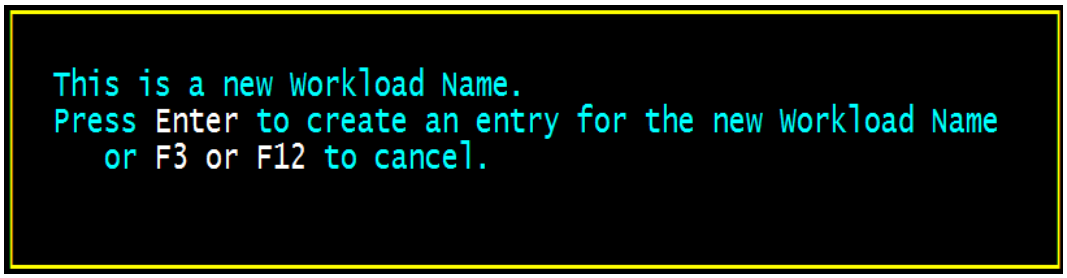

Figure 5 Create New Warehouse Workload

Press <**enter**> to create the new Workload, or press <**PF3**> or <**PF12**> to cancel.

Note that the names of Trace Workloads and SQL Warehouse Workloads MUST be unique. This allows a many- to-many relationship enabling you to load SQL Trace Files from many Trace Workloads into a single SQL Warehouse Workload, or SQL Trace Files from a Trace Workload into many SQL Warehouse Workloads.

The newly created Workload is set as the Current Workload.

To display a list of existing workloads, you can DRILL <PF4> on the Warehouse Workload field.

You can filter the list of Workloads displayed by entering a MASK in the Warehouse Workload field before pressing the <ENTER> or DRILL <PF4> key.

To specify a MASK use % to match remaining characters. Use _ (underscore) to match any character in that position.

The following figure shows the Warehouse Workload Manager panel after the user has specified GX% in the Warehouse Workload field and pressed <ENTER>

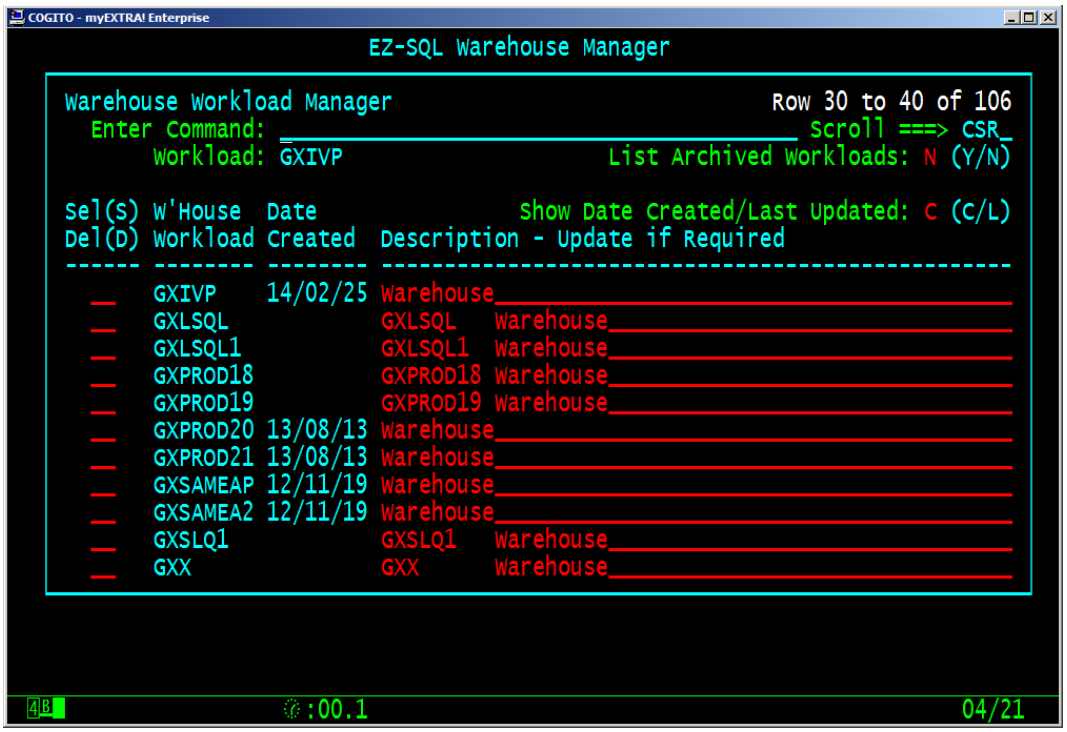

Figure 6 Warehouse Workload Manager

From this panel, you can *set* the current SQL Warehouse Workload, *delete* an existing SQL Warehouse Workloads or enter a workload description.

The panel fields are described below.

Workload

The name of the current Warehouse Workload is displayed.

List Archived Workloads

Specify whether you wish to display the active or the archived workloads.

Y List the archived workloads. If you have specified a MASK, only the workloads matching the MASK will be displayed.

N List the active workloads. If you have specified a MASK, only the workloads matching the MASK will be displayed.

Show Date Created/Last Updated

You may toggle the display to show either the Date Created or the Date Last Updated for each workload.

C Show Date Created.

L Show Date Last Updated

You may SORT the workloads in the list by positioning the cursor in the Date Created/ate Last Updated column and pressing the <ENTER> key.

Sel(S)
Del(D)

You can specify a select value against an existing Warehouse Workload as follows:-

S - Set this Workload as the current Workload.

Note: You can also set the current warehouse Workload by over typing the Warehouse Workload name on any of the EZ-DB2 panels where the Warehouse Workload name appears.

D - Delete this Workload. A Job will be generated to delete the Workload datasets.

A - Archive the Workload. The Workload data is retained but the name is removed from the active workload list. To view the workload names that have been archived, change List Archived Workloads to Y.

R - Recall Workload. If a workload has not been referenced for a while, it's possible that the applicable datasets have been archived by DFSMSHSM. Specify R to send a Recall request to DFSMSHSM for all of the workload datasets,

If the workload has been archived it will also be restored from the list of archived workloads.

Note: You can use block commands, **AA**, **DD** and **RR** to archive, delete and restore multiple workloads.

W'House Workload

The Warehouse Workload name.

This field is sortable. Position cursor in this field and press the <enter> key.

Date Created / Date Last Updated

Either the Date Created or the Date Last Updated is displayed, depending on the setting of the Date Created/Last Updated toggle.

This field is sortable. Position cursor in this field and press the <enter> key.

Description

You may over type this field to enter some descriptive text that describes this Warehouse Workload.

6.1 Select and Import SQL Traces

Selecting and importing SQL Traces into the SQL Warehouse Workload may be an iterative process. After loading the SQL trace data into the SQL Warehouse Workload, you can use the various online displays to review the components of the workload and further refine the workload by running subsequent Traces which may be selected and imported into this SQL Warehouse workload. You may also refine the SQL Warehouse Workload by adding SQL from the Catalog, adding Freeform SQL and also by modifying the SQL trace frequencies as described later. This process may continue until you are satisfied that the SQL Warehouse Workload is sufficiently representative of the Workload you wish to Analyze and/or Optimize.

Select option 6.1 to display the Select and Import SQL Traces panel as shown in the following figure:-

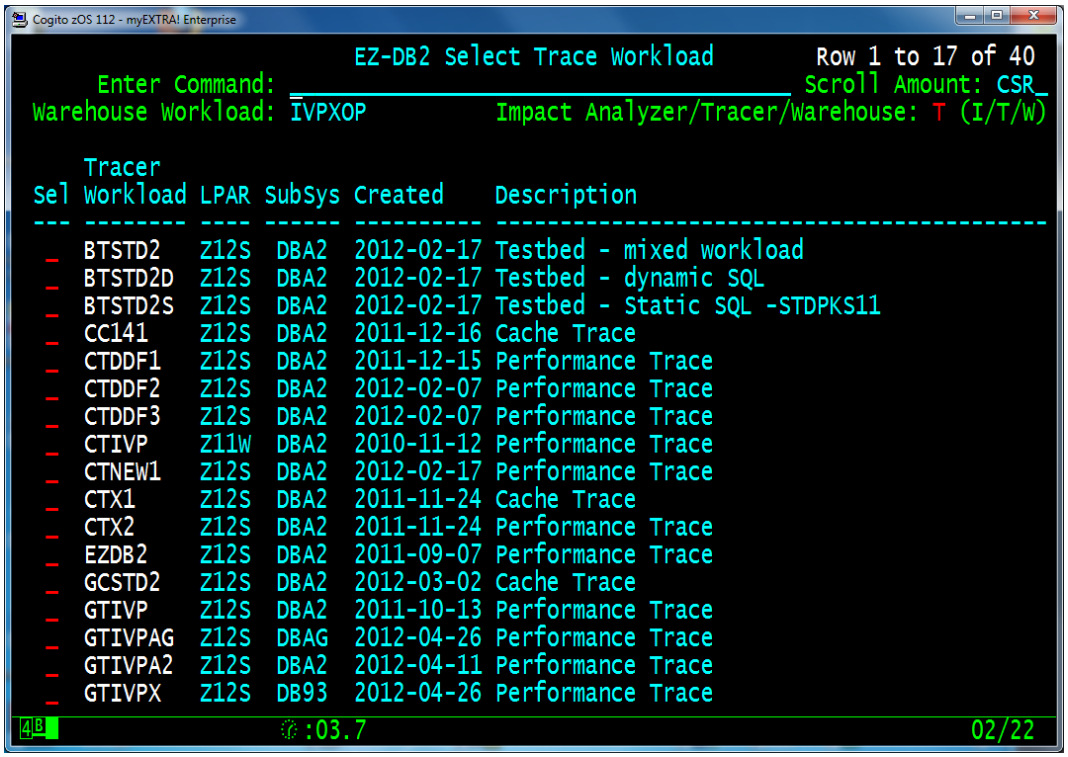

Figure 7 Select and Import SQL Traces

This option presents you with a panel listing all available Workloads. You are required to Select the Trace Workload from which you wish to load SQL Trace Files.

The following information is displayed on this Panel:-

Warehouse Workload	The current Warehouse workload name that will be populated with the selected data.

Impact Analyzer/Tracer/Warehouse (I/T/W)

You may populate Warehouse workloads with data from Tracer, Impact Analyzer or indeed from other Warehouse workloads. To select which type of Workload should be displayed, specify:-

I – All available Impact Analyzer Workloads will be displayed.

T – All available Tracer Workloads will be displayed.

W – All Available Warehouse Workloads will be displayed.

Sel

Specify S in the Sel column against the Workload from which you wish to load Trace data.

Tracer Workload

The name of the Workload as created by the EZ-Tracer or EZ-Cache product.

Note; if you have specified that you wish to import data from Impact Analyzer or Warehouse workloads, the column header will change as applicable.

DB2 LPAR

The LPAR on which the Trace Workload was created.

Note: If the LPAR is a different LPAR to where the EZ-DB2 Analysis LPAR, and the datasets are not allocated on shared DASD, then you must use option 9.4 - Export Workload and option 9.5- Import Workload to Transfer Traces to the EZ-DB2 analysis LPAR before you can import them into the SQL Warehouse. Refer to EZ-DB2 Housekeeping and Controls Guide for further information.

SubSys

The DB2 Subsystem where the Trace Workload was created.

Date Created

Date the Trace Workload was created.

Description

Some descriptive text describing this Trace Workload if this was entered when creating the Workload.

Select the required Trace Workload and press <**enter**>. A panel is displayed showing the available SQL Trace Files for the selected Tracer Workload as shown in the following figure:-

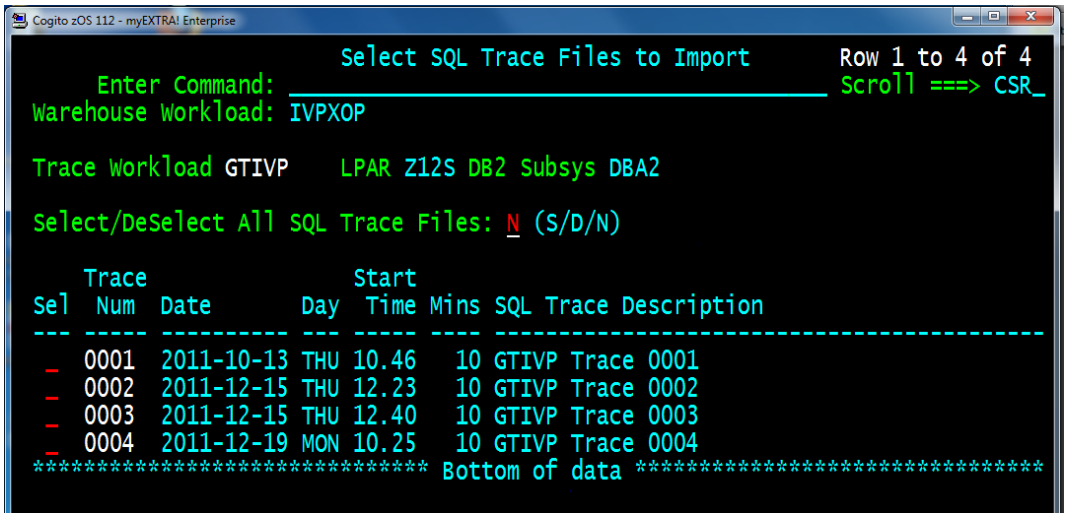

Figure 8 Select SQL Traceiles to Import

The following information is displayed on this Panel:-

Warehouse Workload	The current Warehouse Workload name.
Trace Workload	The Trace workload name selected
LPAR	The LPAR on which the trace was created
Subsys	The DB2 Subsystem where the trace Workload was created
Select/DeSelect All SQL Trace Files (S/D/N)	Enter S to Select all, D to reset all selection indicators or N to manually select specific files.
Sel	Specify S in the Select Column to select a particular Trace file. See also Select/DeSelect All SQL Inputs described previously.
Trace Num	The number of the SQL Trace File
Date	The date when the SQL Trace File was created
Day	The day of the week that the SQLTrace File was created.
Start Time	The time that the SQL Trace was started
Mins	The duration of the trace.
SQL Trace Description	The description of the SQL Trace File.

Select all of the applicable SQL Trace Files for this Trace Workload that you wish to include in the current SQL Warehouse Workload. When you have selected all of the applicable SQL Trace Files, press **<enter>** to save the selection and exit to the Workload Analysis SQL and Statistics Controls panel as described in the next section.

You may if you wish repeat this process, selecting SQL Trace Files from other Trace Workloads into the current SQL Warehouse Workload until all of the relevant SQL Trace Files have been selected.

Workload Analysis SQL and Statistics Controls

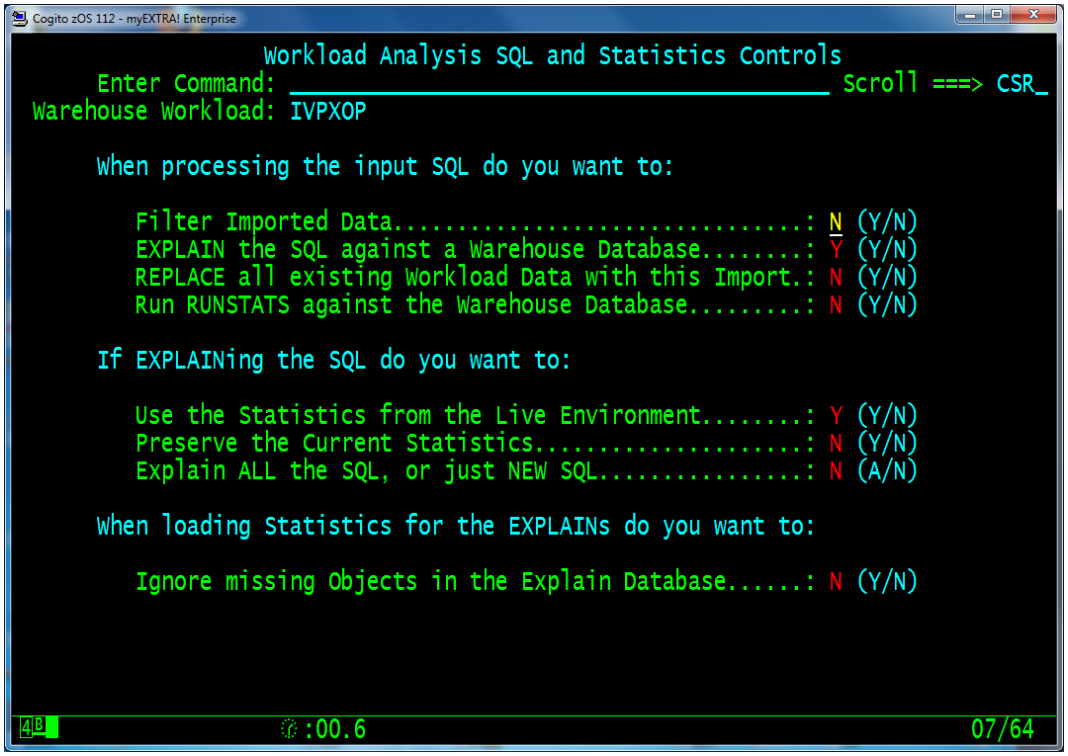

Figure 9 **Workload Analysis SQL and Statistics Controls**

When importing *Trace Workloads* into the SQL Warehouse, you can decide whether you wish to Explain the SQL in the SQL Warehouse analysis environment at the same time. To Explain the SQL, you must have a database in the analysis environment suitable for explaining against. See Warehouse DBName in Select Databases for Processing panel in the next section.

It is not required for you to Explain the SQL if you only plan to use EZ-Alerts and/or EZ-Impact Analyzer. However, you must have an *Explain database* and Explain the SQL if you are intending to use EZ-Index Analyzer or EZ-Index Optimizer.

The following options may be specified:-

When Processing the Input SQL do you want to:-

Filter Imported Data (Y/N) You may choose to import ALL data from the selected Tracer/Warehouse/Impact Analyzer Workloads, or you may choose to filter the imported data based upon applicable selection criteria.

To filter the data, enter Y or DRILL <PF4> on the highlighted field. The following panel will be displayed:-

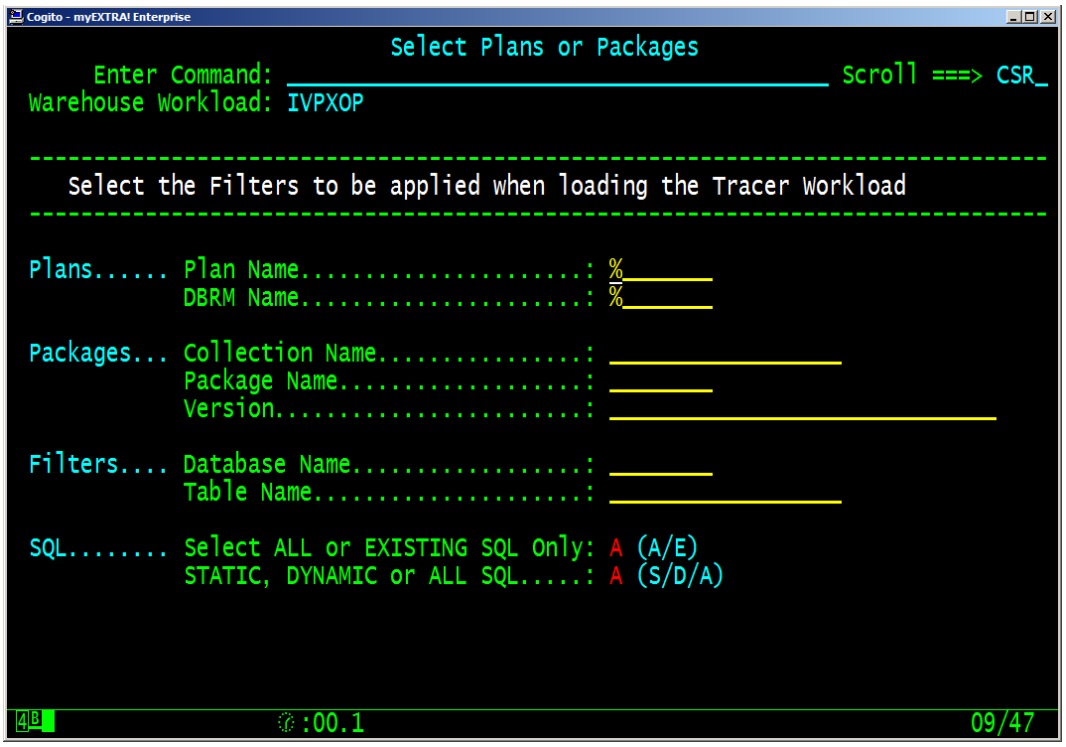

Figure 10 Workload Import Filter parameters

Plans

Plan Name

Specify the Plan Name and DBRM for the program that you want to include in the workload.

NOTE: You can use the % sign in the same manner as a LIKE condition to load SQL from groups of PLANS or PACKAGES.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

NOTE: You can Drill (PF4) on the Plan Name or DBRM Name fields to add up to 8 Plan/DBRM name pairs

DBRM Name

Specify the Plan Name and DBRM for the program that you want to that you want to include in the workload.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

Packages

Collection

Specify the Collid, Package name and optionally the Version for the program that you want to want to include in the workload. If Version is omitted, then Warehouse Manager will use the latest version of the Package in the catalog.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

NOTE: You can use the % sign in the same manner as a LIKE condition to load SQL from groups of PLANS or PACKAGES

NOTE: You can Drill(PF4) on the Collection, Package or Version fields to add up to 8 Collection/Package/Version sets

Package

See Collection.

Version

See Collection

Filters

Database name

You may provide Database Filters so the only SQL selected from the nominated Plans or Packages are those that reference the specified Database Name(s).

This field will contain the database name when only a single database name has been selected. When more than one is selected the field will display the word **MULTIPLE**.

You may select up to 8 database names. However, if you also select Table Names only one database name may be selected.

Use <PF4> to drill on the field to obtain a list of available Database names. Use a prefix and % to limit the displayed list to all databases with that prefix.
E.g. ABC% for all database names starting ABC.

EZ-DB2 will display a list of the database names available from the selected subsystem and will indicate any which have been selected.

If EZ-DB2 is unable to access the subsystem then a default panel will display showing any database names currently selected. You can manually enter database names using this panel.

Table Name

You may also provide Table Name Filters so the only SQL selected from the nominated Plans or Packages are those that reference the specified Table Name(s).

This field will contain the Table name when only a single table name has been selected. When more than one is selected the field will display the word **MULTIPLE**.

You may select up to 50 table names. However, when you select table names you must specify one and only one database name.

Use <PF4> to drill on the field to obtain a list of available Table names. Use a prefix and % to limit the displayed list to all Tables with that prefix.
E.g. XYZ% for all table names starting XYZ.

EZ-DB2 will display a list of the table names available from the selected subsystem for the selected database name and will indicate any which have been selected.

If EZ-DB2 is unable to access the subsystem then a default panel will display showing any table names currently selected. You can manually enter table names using this panel.

If you have table names selected and you change the selected database name then all selected table names will be deselected.

SQL

Select ALL or EXISTING SQL Only

This control allows you to filter the incoming SQL from a Trace to process only those SQL already in the Workload. This enables you to analyze the performance of a distinct application by pre-loading all the application SQL into the Workload, using the option to Load SQL From Catalog.

Having created the SQL Warehouse Workload with all the application SQL you may import many traces from EZ-Tracer or EZ-Cache and exclude any SQL from other applications.

Enter A to process ALL SQL or E to process only EXISTING SQL. The default value is A.

STATIC, DYNAMIC or ALL SQL This control allows you to filter the incoming SQL from a Trace to include just Static SQL, just Dynamic SQL or ALL SQL in the Warehouse Workload.

Enter **S** to include just Static SQL
Enter **D** to include just Dynamic SQL
Enter **A** to include All (both Static and Dynamic) SQL in the Workload.

The default value is A.

Once you have entered any required FILTER parameters for the data to be imported, press <PF3> to return to the Workload Analysis SQL and Statistics Controls panel. The following additional parameters may be specified:-

EXPLAIN the SQL against a Warehouse Database If you have Database(s) in the Warehouse environment with matching structure to the Live Database(s) then the SQL may be EXPLAINed in the Warehouse Environment. This is not required for EZ-Alerts and EZ-Impact Analyzer but is mandatory for EZ-Index Analyzer and EZ-XOP.

Note the Analysis Databases do not require the same DBNAMEs or Creator names as these values will be automatically translated.

Enter Y to EXPLAIN the SQL or N not to.

REPLACE all existing Workload Data with this Import This option allows you to replace the existing data in the Workload with the data from this import. This means you do not have to delete and recreate the workload if you no longer want the current data and wish to reuse the workload to import new data.

Enter Y to REPLACE the exist data in the Workload. Enter N to keep the existing data.

Run RUNSTATS against the Warehouse Database This option allows you to choose whether or not to run RUNSTATS against the EZ-DB2 Warehouse Database.

Over a period of time the amount of data in a warehouse workload may grow in size and it is recommended that you run RUNSTATS in order to optimize the various EZ-DB2 process that access the warehouse workload.

Enter Y to run RUNSTATS, enter N not to run RUNSTATS

If EXPLAINing the SQL do you want to:

Use the Statistics from the Live Environment

When the SQL is explained it is desirable to use the Statistics taken from the environment from which the SQL is captured, particularly if it was from a Trace in production. You may manually control which Statistics are used through EZ-Stats, or let the software apply the Statistics it has gathered.

If this workload is being created for EZ-XOP (Index Optimizer), then you MUST copy the Live Statistics to the EZ-XOP Clone database in the analysis environment. This option should be set to Y and ensure that the statistics in the Live environment are up-to-date, and that you have run RUNSTATS for ALL columns.

Enter Y to use the Live Statistics or N to use existing Statistics. The default value is Y.

Preserve the Current Statistics

If you opt to let the software update the Statistics then you may opt to preserve the Statistics currently in place by archiving them at the start of the process and restoring them once the process is complete.

Enter Y to Archive/Restore current Statistics or N if this is not required. The default value is Y.

If this workload is being created for EZ-XOP (Index Optimizer) then the database will be an EZ-XOP Clone database. You should use the Live Statistics as described above, but there is no need to Preserve the Current Statistics. This option should therefore be set to N.

Explain ALL the SQL, or just NEW SQL

If you are analyzing a very large Workload with thousands of SQL then it may take some time to EXPLAIN the entire Workload. When updating the Workload from new traces it is not necessary to EXPLAIN the existing SQL again, only new SQL, unless you have made some change to the EXPLAIN Databases which might affect access paths. Explaining only new SQL will speed up the process when updating the Workload with new trace data.

Enter A to explain ALL SQL or N to explain NEW SQL only. The default value is A.

When loading Statistics for the EXPLAINs do you want to:

Ignore missing Objects in the Explain Database

If you have opted to use Statistics from the Live environment then the procedures will update Statistics for all Objects in the selected Databases. If the Database(s) against which the SQL are EXPLAINed do not exactly match the Live Database(s) then any Objects missing in the Explain Database will be reported. For example a Tablespace may be defined as Partitioned in the Live environment but not in the Explain Database, which may be acceptable. However

you should use this option with caution as missing Tables may cause SQL to fail the 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.

Enter any required changes then press <**enter**> to display the Select Databases for Processing panel as described in the next section.

Select Databases for Processing

After specifying the Workload Analysis SQL and Statistics Controls the following panel is displayed:-

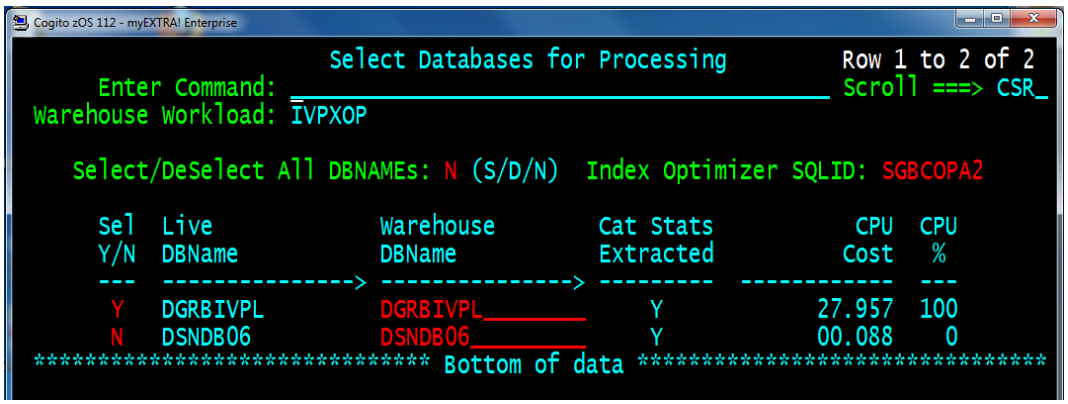

Figure 11 Select Databases for Processing

This panel displays an entry for each database encountered in the Selected SQL Trace Files sorted in descending CPU Cost.

You may select the databases for which you wish the SQL Warehouse Manager to load SQL into the current SQL Warehouse Workload for subsequent analysis by other EZ-DB2 tools.

For each database selected, If you wish to Explain the SQL then an Explain database must exist in the warehouse environment, mirroring the Live definition of that database. The DBNAME of the Explain database may be different to the Live DBNAME, as may the Creator names used in the DDL, but the Explain database should contain all the same Tablespaces, Tables, Views, Indexes and Foreign Keys. If any of these entities are missing the analysis may be incomplete.

If this Workload is being created for EZ-XOP (Index Optimzer) analysis, then the Warehouse database will be the name of the corresponding EZ-XOP Clone database. Refer to the Appendix Creating the Explain Database for further information.

The following fields may be specified on this panel:-

Select/DeSelect All DBNames(S/D/N) — Enter S to select all Databases, D to reset the selection indicator to N for all Databases, or N to manually select specific Databases.

The default value is N.

Index Optimizer SQLID — An AUTHID for the Index Optimizer (EZ-XOP) to use when creating Candidate Indexes for evaluation. This is not required for any other product.

This is usually the Authid used when creating the EZ-XOP Clone database and must NOT be the same as the Authid used to create the EZ-DB2 Warehouse Database. Prior to the CREATE INDEX a SET CURRENT SQLID will be issued using the specified SQLID.

Sel (Y/N)

Specify Y to select this database. The SQL that references tables in the database will be loaded into EZ-DB2 for analysis.

Specify N to de-select this database for EZ-DB2 processing.

If you specified DBNAMEs in the Workload Trace parameters, the applicable DBNAMEs will be set to Y. All other DBNAMEs encountered in the Trace will be set to N.

If you did not specify any DBNAMEs in the Workload Trace Parameters, all of the DBNAMEs encountered in the workload will be set to N. You may Select or De-Select DBNames as required.

Live DBName

The DBNAME of the database encountered in the SQL Trace Files.

Warehouse DBName

The DBNAME of the Explain database in the EZ-DB2 Warehouse subsystem corresponding to the Live database, against which SQL in the Workload can be Explained.

For EZ-XOP workloads, this will be the name of the EZ-XOP Clone database, Refer to Appendix Creating the Explain Database, and the *EZ-XOP User Guide* for further information.

CATSTATS Extracted

If Y then Catalog Statistics have already been extracted. If you select a database for processing where this is set to N then you will be prompted to submit the job to extract the missing statistics.

CPU Cost

The total CPU consumed by SQL in the Workload that run against this Database.

CPU %

The percentage of the total CPU for the Workload consumed by the SQL that run against this Database.

Press **<enter>** to save any changes and to generate the job stream to load the SQL Trace data into the current SQL Warehouse Workload.

- Type SUBMIT in the command line and press <**PF3**> to return to the SQL Warehouse manager menu.

Job step summary

The generated job contains the following internal job steps:-

Job Step	Description
ALLOCDS	Allocate required datasets
COMPRESS[1]	Compress Statistics Archive File
RUNSTATD[1]	Archive Current Statistics
RUNSTATU[2]	Load Live Catalog Statistics
RUNELOG	Update Warehouse Workload with SQL data
RUNIMPE	Extract SQL data from the Warehouse database and populate the VSAM files ready for online report processing.
SUMMARY	Accumulate Workload Summary.
RUNSTATU[1]	Restore Catalog statistics from Archive
DELGRPTS[3]	Existing report files for the workload are deleted before they are recreated by the next step.
REPORTS[3]	Generate Reports for Warehouse
DELGRPTS[3]	Delete Reports for Alerts and Impact Analyzer for Version O Original Costs.
REPORTS[3]	Generate Reports for Alerts Version O
REPORTS[3]	Generate Reports for Impact Analyzer Version O
DELGRPTS[4]	Delete Reports for Alerts for Version E Live EXPLAIN
REPORTS[4]	Generate Reports for Alerts Version E
REPORTS[4]	Generate Reports for Impact Analyzer Version E

For further details of these job steps refer to the Appendix.

---

1 Only if Preserve the Current Statistics is selected

2 Only if Use the statistics from the Live environment is selected.

3 Only if Explain the SQL against Analysis environment is selected

4 Only if Live Explain (Version E) access path was created by Tracer.

6.2 Update Qualifiers (if required)

This option is required if the Select and Import SQL Traces process fails in the RUNELOG step because it has failed to perform the *Live* to *Explain* table qualifier translation. In most cases, this step will not be necessary.

Refer to the Appendix section Load Catalog Statistics - Qualifier Translation Error for further information.

Select option 6.2 to display the Update Qualifiers panel as shown below:-

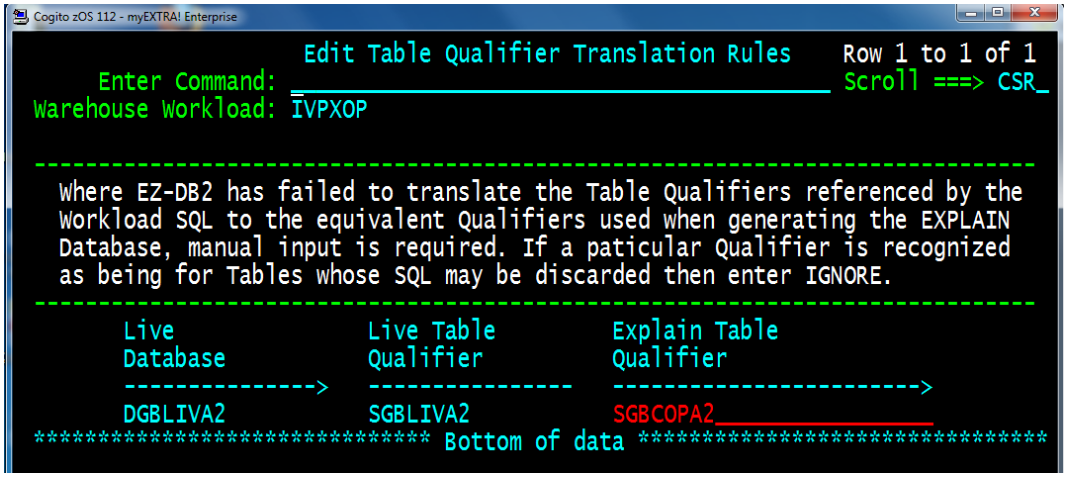

Figure 12 Edit Table Qualifiers Control File

This panel lists the *Live Table Qualifier(s)* and corresponding *Explain Database Table Qualifiers* as automatically derived by EZ-DB2 for the current SQL Warehouse Workload. It should not normally be necessary to modify these values.

The following fields are displayed on this panel:-

Live Database	The List of Live DBNAMEs that have been referenced by SQL in the current SQL Warehouse Workload
Live Table Qualifier	The table qualifier as applicable on the Live database.
Explain Table Qualifier	The table qualifier as applicable on the Explain database. Where EZ-DB2 has been unable to determine the Explain database table qualifier, the Explain database Table Qualifier will be set to '????????'. Set this to the required value. If the table is not required in the analysis, set to IGNORE.

Press **<enter>** to save any changes and to generate the job stream to restart the Select and Import SQL Traces job stream from the RUNELOG step.

- Type SUBMIT in the command line and press **<PF3**> to return to the SQL Warehouse manager menu.

6.3 Load Application SQL from Catalog

You may use this option to load SQL for specific application(s) from the System Catalog directly into an SQL *Warehouse Workload*.

See also EZ-Tracer (option 2.3). You may also use this option to initially populate a Trace Workload with SQL for a specific application, and then use the Trace Filter parameters so that the trace only monitors the pre-loaded SQL and ignores any other trace activity.

The SQL is loaded into the SQL Workload with the Execution and Fetch Frequencies set to 1.

You may subsequently use the SQL Warehouse Manager function Manage SQL and Refine Frequencies (option 6.5) to explicitly set the Execution and Fetch Frequencies for each SQL.

However, by pre-loading the SQL from the Catalog, and then using the Trace facility to monitor the pre-loaded SQL, you will collect actual SQL Frequency and Cost information for the application.

This combination of processing options allows you to not only analyse and optimise existing production applications, but also to model planned new applications or simulate predicted workload growth.

Note you can use the % sign in the same manner as a LIKE condition to load SQL from groups of PLANS or PACKAGES.

Select option 6.3 to select the Load Application SQL from Catalog option panel as shown in the following figure:-

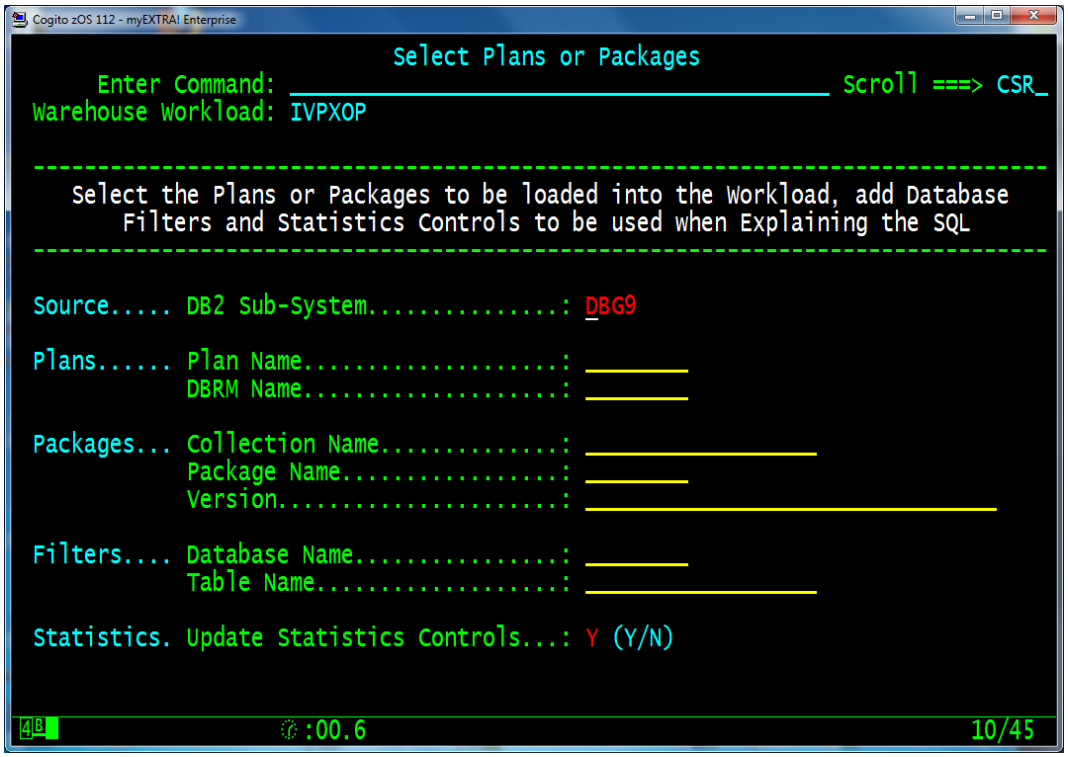

Figure 13 Select Plans or Packages

The following parameters may be specified on this panel.

Source

DB2 Sub-System

Specify the DB2 subsystem ID of the system of the System Catalog you wish to load SQL from

Plans

Plan Name

Specify the Plan Name and DBRM for the program that you want to extract the SQL from.

NOTE: You can use the % sign in the same manner as a LIKE condition to load SQL from groups of PLANS or PACKAGES.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

NOTE: You can Drill (PF4) on the Plan Name or DBRM Name fields to add up to 8 Plan/DBRM name pairs

DBRM Name

Specify the Plan Name and DBRM for the program that you want to extract the SQL from.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

Packages

Collection

Specify the Collid, Package name and optionally the Version for the program that you want to extract the SQL from. If Version is omitted, then Warehouse Manager will use the latest version of the Package in the catalog.

NOTE: You can extract SQL from either a Plan or a Package but not both in the same run.

NOTE: You can use the % sign in the same manner as a LIKE condition to load SQL from groups of PLANS or PACKAGES

NOTE: You can Drill(PF4) on the Collection, Package or Version fields to add up to 8 Collection/Package/Version sets

Package

See Collection.

Version

See Collection.

Databases

Database Name

You may provide Database Filters so the only SQL selected from the nominated Plans or Packages are those that reference the specified Database Name(s).

This field will contain the database name when only a single database name has been selected. When more than one is selected the field will display the word **MULTIPLE**.

You may select up to 8 database names. However, if you also select Table Names only one database name may be selected.

Use <PF4> to drill on the field to obtain a list of available Database names. Use a prefix and % to limit the displayed list to all databases with that prefix.
E.g. ABC% for all database names starting ABC.

EZ-DB2 will display a list of the database names available from the selected subsystem and will indicate any which have been selected.

If EZ-DB2 is unable to access the subsystem then a default panel will display showing any database names currently selected. You can manually enter database names using this panel.

Table Name

You may also provide Table Name Filters so the only SQL selected from the nominated Plans or Packages are those that reference the specified Table Name(s).

This field will contain the Table name when only a single table name has been selected. When more than one is selected the field will display the word **MULTIPLE**.

You may select up to 50 table names. However, when you select table names you must specify one and only one database name.

Use <PF4> to drill on the field to obtain a list of available Table names. Use a prefix and % to limit the displayed list to all Tables with that prefix.
E.g. XYZ% for all table names starting XYZ.

EZ-DB2 will display a list of the table names available from the selected subsystem for the selected database name and will indicate any which have been selected.

If EZ-DB2 is unable to access the subsystem then a default panel will display showing any table names currently selected. You can manually enter table names using this panel.

If you have table names selected and you change the selected database name then all selected table names will be deselected.

Statistics

Update Statistics Controls

You may elect to control which Statistics are used when the SQL loaded from Catalog are Explained.

If loading from a DB2 Subsystem other than the EZ-DB2 Warehouse Subsystem then the process will automatically capture the Statistics from the other environment.

Enter Y to link to the panel where you may choose to use the source or current system statistics.

If you have selected to update statistics controls, the Workload Analysis SQL Statistics Controls panel is displayed as shown in the following figure:-

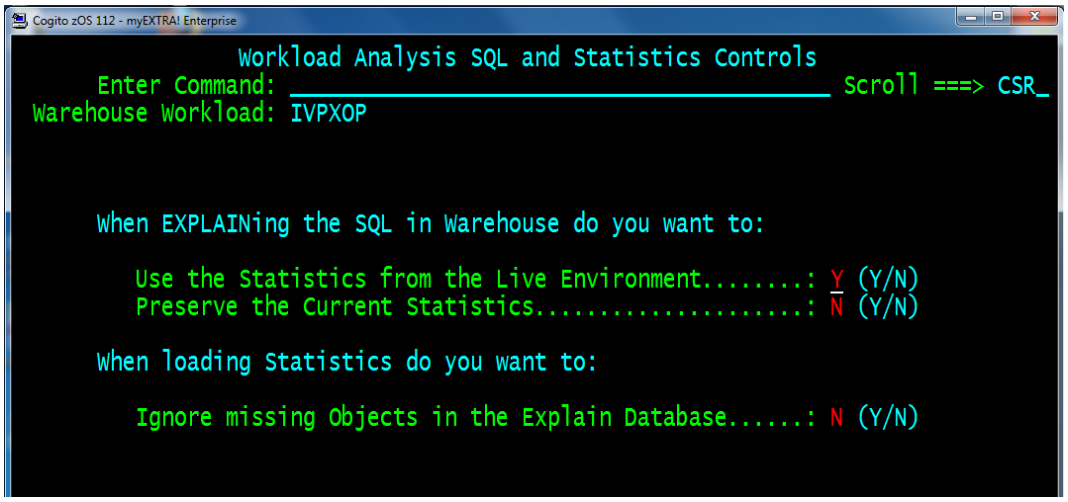

Figure 14 Workload Analysis SQL and Statistics Controls

The following Options may be selected on this panel:-

When EXPLAINing the SQL do you want to:

Use the Statistics from the Live Environment

When the SQL is explained it is desirable to use the Statistics taken from the environment from which the SQL is captured, particularly if it was from a Trace in production. You may manually control which Statistics are used through EZ-Stats, or let the software apply the Statistics it has gathered.

Enter Y to use the Live Statistics or N to use existing Statistics.

The default value is N.

Preserve the Current Statistics

If you opt to let the software update the Statistics then you may opt to preserve the Statistics currently in place by archiving them at the start of the process and restoring them once the process is complete.

Enter Y to Archive/Restore current Statistics or N if this is not required.

The default value is N.

When Loading the statistics do you want to:-

Ignore missing Objects in the Explain Database

If you have opted to use Statistics from the Live environment then the procedures will update Statistics for all Objects in the selected Databases. If the Database(s) against which the SQL are EXPLAINed do not exactly match the Live Database(s) then any Objects missing in the Explain Database will be reported. For example a Tablespace may be defined as Partitioned in the Live environment but not in the Explain Database, which may be acceptable. However you should use this option with caution as missing Tables may case SQL to fail the 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 if missing objects are detected when loading statistics.

The default value is N.

Press **<enter>** to return to the Select Plans or Packages panel.

Press **<enter>** again to continue. The following panel is displayed:-

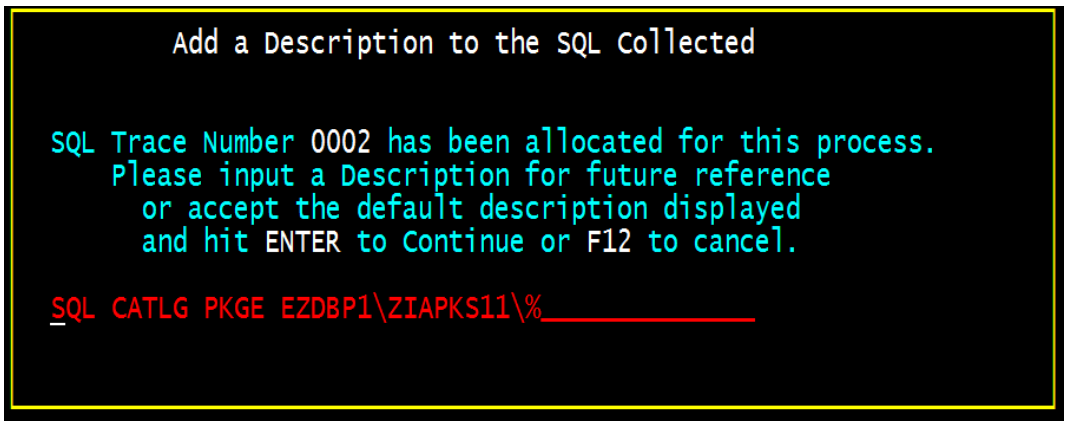

Figure 15 Add an SQL Input Description

The panel displays the number of the next available SQL Trace Files that will be created to contain the applicable SQL loaded from the catalog.

You are prompted to enter a Description for this trace file. Enter a Description and press <**enter**> to display the generated job stream to load the SQL from the selected Plans or Packages.

- Type SUBMIT in the command line to submit the job for execution.
- Press the <PF3> key to return to the SQL Warehouse Manager menu.

Job Step Summary

The generated job contains the following internal job steps:-

Job Step	Description
	Following job steps executed on live subsystem LPAR
ALLOCDS	Allocate required datasets
ALOCRPTV	Allocate VSAM Report datasets
RUNVSAM3	Initialize VSAM Report datasets
ALOQSQLV	Allocate EZ-DB2 SQL VSAM Work Datasets
RUNASQL	Load SQL from PLANs and Packages
XOPTRC2	Analyze SQL data
SUMMARY	Accumulate Workload Summary.
RUNLOGO	Create SQLLOGO File. This is a consolidated SQLLOG file that is used when importing data into the SQL Warehouse.
ENDSQLV	Delete VSAM and Log datasets
COMPRESS[5]	Compress CATSTATS dataset in preparation for running RUNSTATD
RUNSTATD[5]	Extract System Catalog statistics for applicable databases to the statistics Warehouse.
	Remaining job steps executed on analysis subsystem LPAR
COMPRESS[6]	Compress CATSTATS Achive
RUNSTATD[6]	Archive current catalog statistics
RUNSTATU[5]	Load Live Catalog statistics
RUNELOG	Update Warehouse with SQL data
RUNIMPE	Extract SQL data from the Warehouse tables and populate the VSAM files ready for online report processing.
SUMMARY	Accumulate Workload Summary Data

[5] Only if Use Statistics from Live environment is selected

[6] Only if Preserve the Current Statistics is selected

RUNSTATU[6]	Restore Catalog Statistics from Archive
DELGRPTS	Existing report files for the workload are deleted before they are recreated by the next step.
REPORTS	Generate Reports for Warehouse

For further details of these job steps refer to the Appendix.

6.4 Load Additional Free Form SQL

This panel submits a job that will add additional SQL from a free-form SQL text file to the SQL Warehouse workload. Additional SQL can either be added to an existing workload or can be used to build a new workload, for example, where an application is still under development.

Put the SQL Text in a member in a PDS dataset and specify the PDS dataset name and member on the panel. When you specify a PDS dataset in the panel you can drill on the Member name field for a list of members to select from.

The SQL must be valid and be fully qualified or preceded by an appropriate SET CURRENT SQLID. You may specify multiple SQL in the same member but each SQL must be terminated by a semicolon. Comment lines (- - in cols 1and 2) may be specified for your own use.

You may optionally specify –PROGNAME *prog-name* to define the program name to be associated with the subsequent SQL statements. If No –PROGNAME parameter is specified, the program name defaults to SQLTEXT.

If there are live statistics available for the current workload, then those statistics may be the ones used for the analysis. Otherwise the statistics from the Explain database will be used.

If you want to build a workload consisting of a mixture of traced SQL activity and freeform SQL added by this panel, start building the SQL Warehouse workload by loading a Trace Workload first. Then add freeform SQL from this panel. After the Live SQL trace data has been loaded, more trace data and freeform SQL can be subsequently added in any order.

Select option 6.4 to display the Load Freeform SQL Panel as shown in the following figure:-

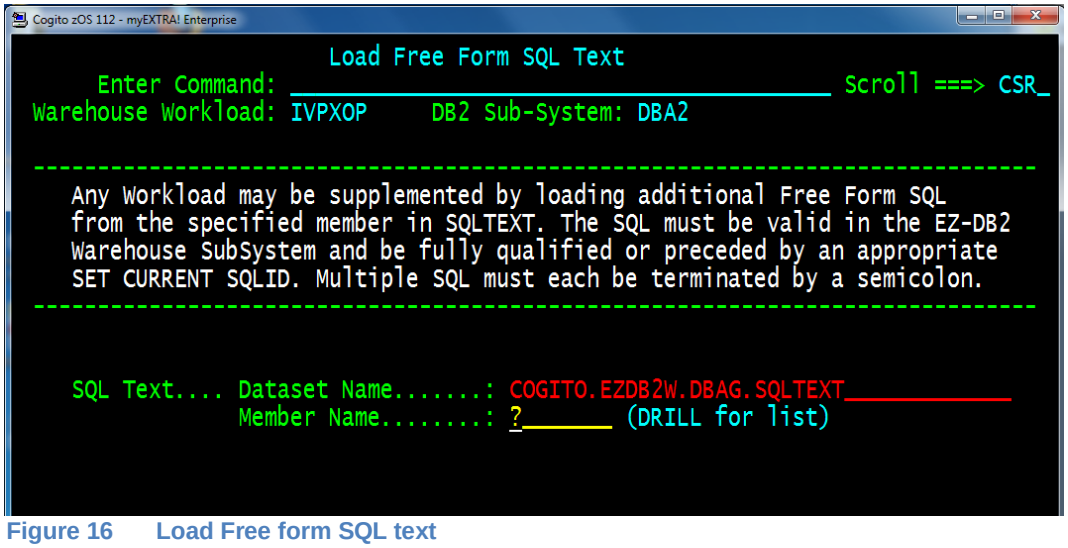

Figure 16 Load Free form SQL text

The following parameters may be specified on this panel.

SQLTEXT Dataset Name Specify the name of the PDS that holds the member which contains the SQL to be loaded

Member Name

Specify the member of the PDS that contains the fully qualified SQL to be loaded.

Delimit each SQL statement with a semicolon or a comment line (-- in cols 1 and 2). Additional comment lines may be specified for your own use and will be ignored (except for the optional –PROGNAME parameter)

Enter a '?' to display a list of members in the applicable dataset from which you can make your selection.

You may select multiple members.

You can use the L(ocate) command to find the member name in the list

See Figure 18 for an example of an SQLTEXT member.

Press **<enter>** to continue. The following panel is displayed:-

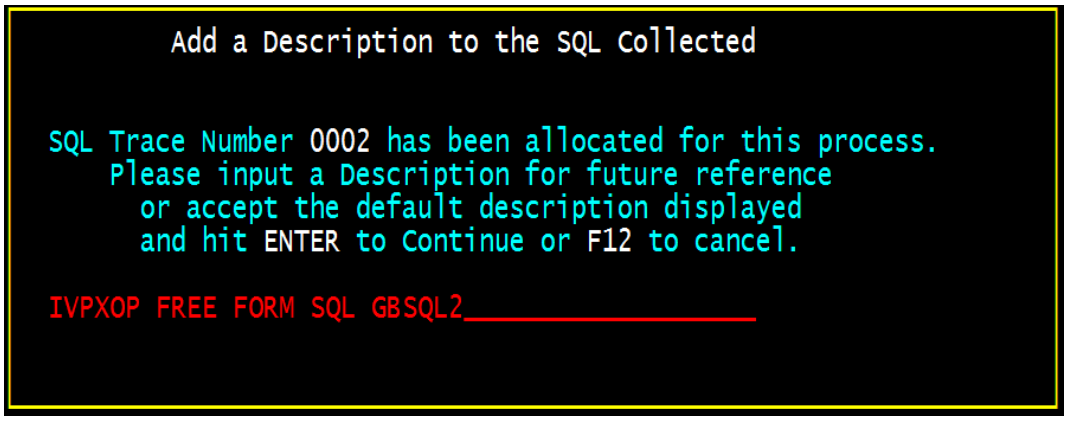

Figure 17 Add an SQL Input Description

The panel displays the number of the next available SQL Trace Files that will be created to contain the applicable SQL entered in free form..

You are prompted to enter a Description for this trace file. Enter a Description and press <**enter**> to display the generated job stream to load the Free form SQL text..

- Type SUBMIT in the command line to submit the job for execution.
- Press the <PF3> key to return to the SQL Warehouse Manager menu.

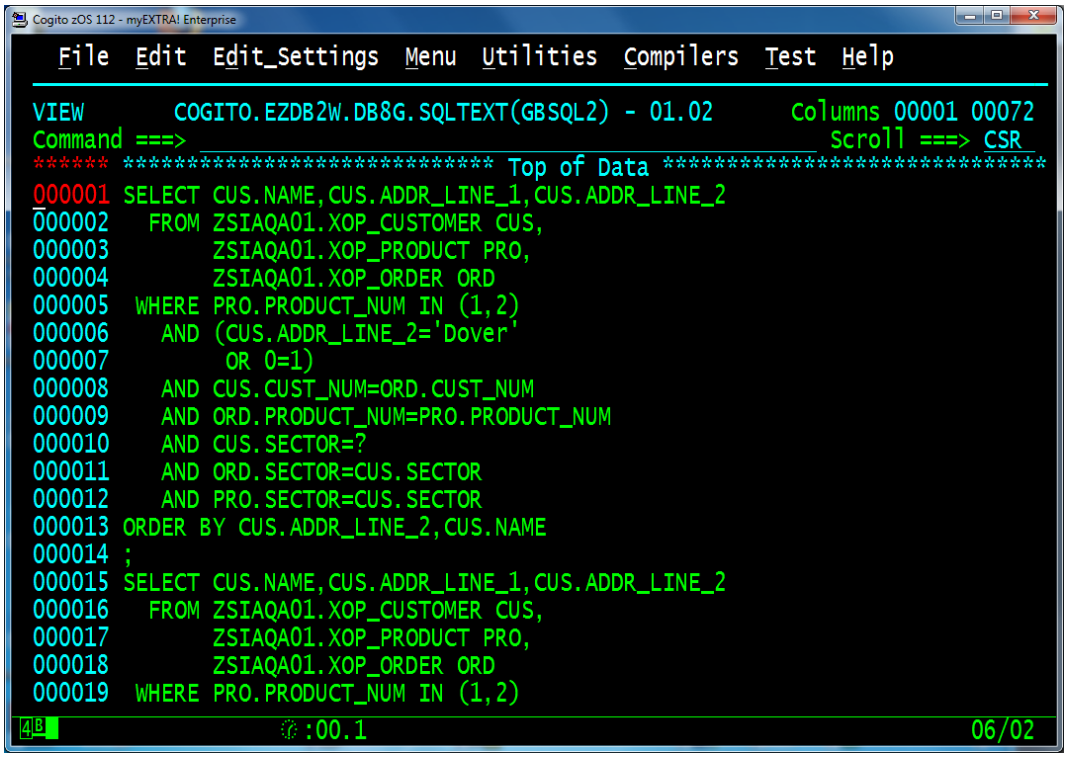

Figure 18 Example of SQLTEXT member

Job Step Summary

The generated job contains the following internal job steps:-

Job Step	Description
ALLOCDS	Allocate required datasets
RUNLSQL	Load additional free form SQL text
COMPRESS[7]	Compress CATSTATS dataset in preparation for running RUNSTATD
RUNSTATD[7]	Extract System Catalog statistics for applicable databases to the statistics Warehouse.
RUNELOG	Update Warehouse with SQL data
RUNIMPE	Extract SQL data from the Warehouse tables and populate the VSAM files ready for online report processing.
SUMMARY	Accumulate Workload Summary Data
DELFILES	Delete Files
DELGRPTS	Existing report files for the workload are deleted before they are recreated by the next step.
REPORTS	Generate Reports for Warehouse

For further details of these job steps refer to the Appendix.

[7] If you specified Use Statistics from Live environment.

6.5 Manage SQL and Refine Frequencies

The SQL Warehouse Manager allows the user to manipulate the SQL Warehouse workload in order to perform "*what-if*" analysis, or to adjust workloads traced in a test environment to more closely resemble a Production Workload.

This facility is particularly useful if using EZ-DB2 in a test environment where the SQL frequencies are not representative of the production environment.

For example, before a new program is available in production, you could

- Use EZ-Tracer to trace the program running in a test environment, or
- Use EZ-Tracer to trace the SQL submitted from SPUFI or
- Use SQL Warehouse Manager to extract the SQL from from the DB2 Catalog (the frequency will be set to one)
- n SQL Warehouse Manager, enter the SQL as free form text (the frequency will be set to one)

Whatever option you have used to load the SQL into the SQL Warehouse Workload, you can now modify the frequency to simulate the anticipated production frequency. EZ-DB2 analysis and optimization components can then be run to see what affect changing the relative frequencies of particular statements has upon, for example, the SQL Alerts or choice of indexes recommended by EZ-XOP.

Select option 6.5 to display the Manage SQL and Refine Frequencies panel as shown in the following figure:-

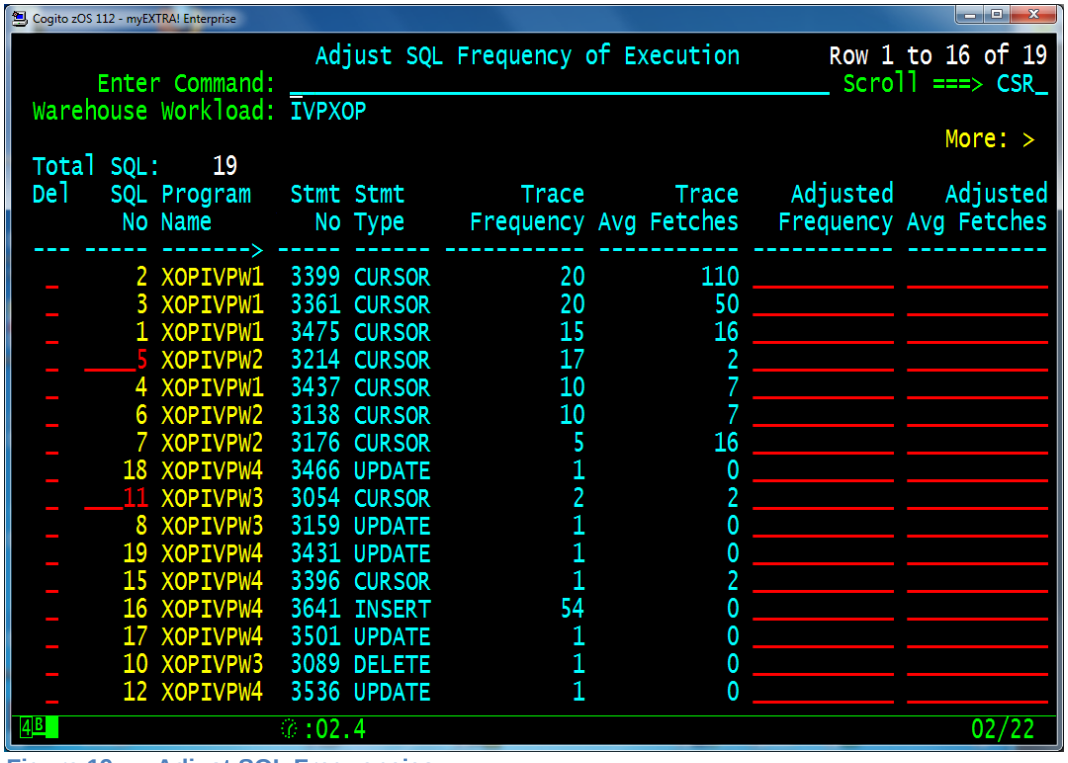

Figure 19 Adjust SQL Frequencies

This displays the total frequencies for each SQL and, for Open Cursor, the average number of Fetches per Open Cursor. You may manipulate the workload by entering your own values for Adjusted Frequency and/or Adjusted Average Fetch Frequency to increase or decrease the footprint of the corresponding SQL in the SQL Warehouse Workload.

The following information is displayed.

DEL

Enter D to delete this SQL from the Workload.

You can also enter 'X' on a particular SQL and this will populate the DEL column of that SQL and of all the other SQL rows below it. This will then generate a job to delete all these SQL from the workload. This can be useful , for example, if you wanted to only keep the 10 most frequently executed SQL in the workload. In this instance you would sort by execution frequency and then put an 'X' in the DEL column of the 11th SQL and press enter. This will then put 'D' in the DEL columns of all SQL from 11th to the end of the workload. The job generated from this will then delete these SQL from the workload leaving only the top 10 most frequently executed SQL's.

Note: Press <PF11> to display additional columns, such as Total/Avergage CPU/Clock/Getpages/imerons. You may also sort on these additional columns.

SQL no

An EZ-DB2 internal identifier used to represent each distinct SQL statement being processed for a particular WORKLOAD.

Program Name

The name of the program containing this SQL statement.

Stmt No

The statement number within the program

Stmt Type

The Statement type, i.e.,

OPEN
UPDATE
INSERT
DELETE

Trace Executions

The actual total number of executions of this statement in the traced workload or 1 if the SQL was extracted from the catalog or loaded in freeform.

Trace Avg Fetchs

The average number of Fetches per Open curser for this SQL statement.

Adjusted Executions

Overwrite this value with the desired number if you wish to modify the trace executions for this SQL statement.

This value is initially set to the actual trace executions or 1 if the SQL was extracted from the catalog or loaded in freeform.

Adjusted Avg Fetchs	Over write this value with the desired number if you wish to modify the Trace average fetchs for this SQL statement. This value is initially set to the actual Trace average Fetch Frequency.
Total CPU	The Total CPU time for the SQL.
Total Clock	The Total Clock time for the SQL.
Total Getpages	The Total Getpages for the SQL.
Total Timerons	The Total Cost for the SQL in Timerons.
Average CPU	The Average CPU for the SQL.
Average Clock	The Average Clock for the SQL.
Average Getpages	The Average Getpages for the SQL.
Average Timerons	The Average Timerons for the SQL.

The values you enter here will be saved in a work-file (*xopwprfx*.XOPUPDT) and applied by submitting the update job presented when you exit the panel.

After modifying any Trace Executions or Trace Average Fetchs values as required, press **<enter>** to display the generated job stream to modify the frequencies in the EZ-DB2 database.

- Type SUBMIT in the command line to submit the job for execution.
- Press the **<PF3>** key to return to the SQL Warehouse Manager menu.

Job Step Summary

The generated job contains the following internal job steps:-

Job Step	Description
RUNSQLU	Update SQL Frequencies and Weightings
RUNIMPE	Extract SQL data from the Warehouse tables and populate the VSAM files ready for online report processing.
DELGRPTS	Existing report files for the workload are deleted before they are recreated by the next step.
REPORTS	Generate Reports for Warehouse

For further details of these job steps refer to the Appendix.

6.6 SQL Warehouse Workload Summary Reports

After an SQL Warehouse Workload has been created, you may use option 6.6 to display the Warehouse Workload summary reports.

The following panel is displayed:-

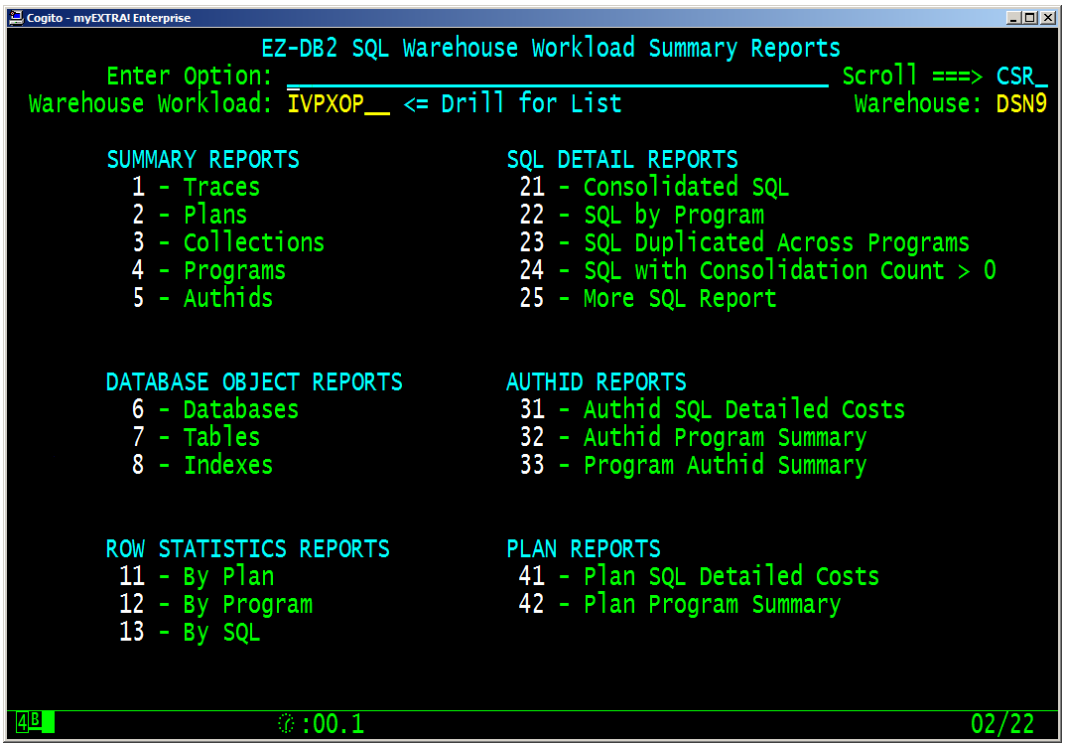

Figure 20 Reports menu

Refer to the EZ-DB2 Tracer, Cache & Warehouse Reports Guide for further information.

6.7 SQL Text and Access Paths after Bind

Select option 6.7 to display SQL detail information after the BIND of an application.

This option displays SQL text, either from the Catalog or from a DBRMLIB. If necessary, the text will be automatically translated from Unicode.

If it exists, a BIND Plan Table can be provided and the access path and estimated costs for each SQL will be retrieved and displayed.

These reports can be viewed interactively (online) or the panel may be used to generate a job stream that may be run in batch. This job stream can be added to existing Compile, Link and Bind processes to generate useful batch reports.

The following panel is displayed when selecting option 6.7.

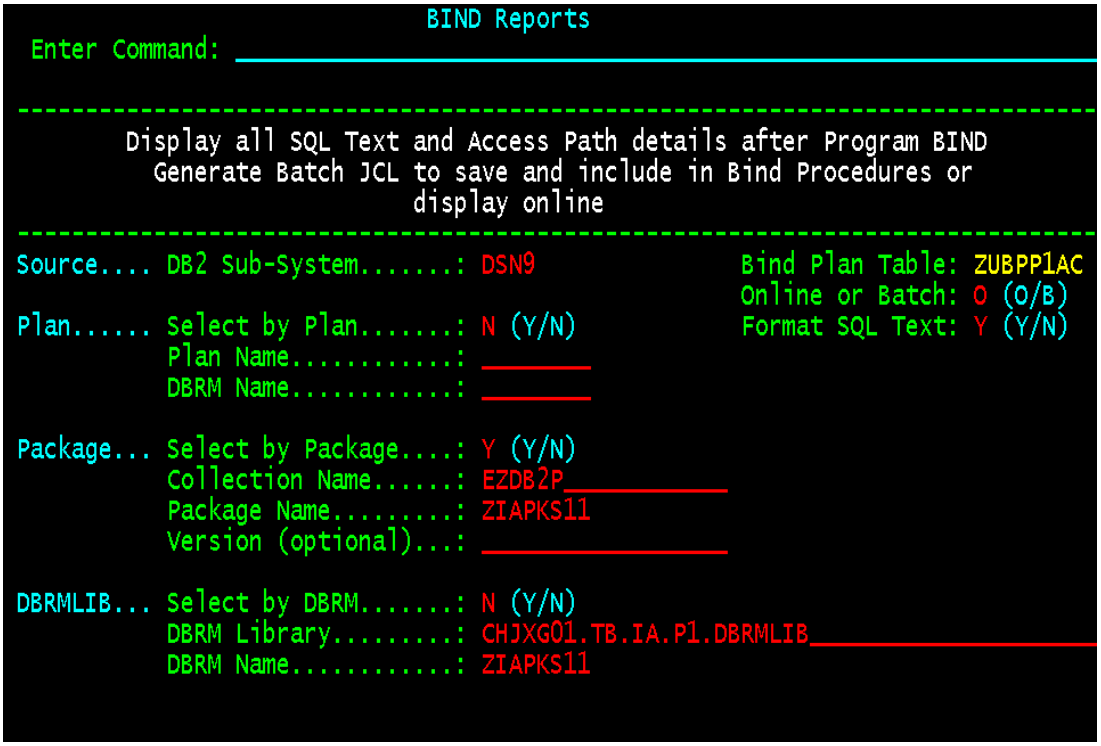

Figure 21 Generate BIND reports

The following information is displayed on the panel.

DB2 Sub-System Specify the name of the DB2 subsystem.

Bind Plan Table Specify the Owner of the Bind Plan table used for the Bind. You can DRILL on this field for a list of available Plan Tables on the specified DB2 subsystem.

A Plan table is required in order to display the access paths and estimated costs.

Online or Batch (O/B) O Generate Online Reports

	B	Generate Batch Job Stream
Format SQL Text (Y/N)	Y	The SQL text will be formatted in the reports.
	N	The SQL text will not be formatted.

Plan

Select by Plan (Y/N)	Y	Select SQL for the specified Plan Name. If specified, you must specify a PLAN and DBRM name.
	N	Do NOT Select SQL by Plan
Plan Name	The name of the Plan	
DBRM Name	The name of the DBRM	

Package

Select by Package (Y/N)	Y	Select SQL by Package/Collection Name. If Specified you must specify a Collection and Package Name.
	N	Do NOT select by Package name.
Collection Name	The name of the Collection.	
Package Name	The name of the Package	
Version (Optional)	You may optionally specify the package version number.	

DBRMLIB

Select by DBRM (Y/N)	Y	Select SQL for a DBRMLIB member. If specified, you must specify the DBRMLIB name and the member name.
	N	Do NOT select by DBRM.
DBRM Library	Specify the DBRM library name. This is the fully qualified name without quotes.	
DBRM name	Specify the DBRM library member name.	

The following figure shows an example of the SQL text display.

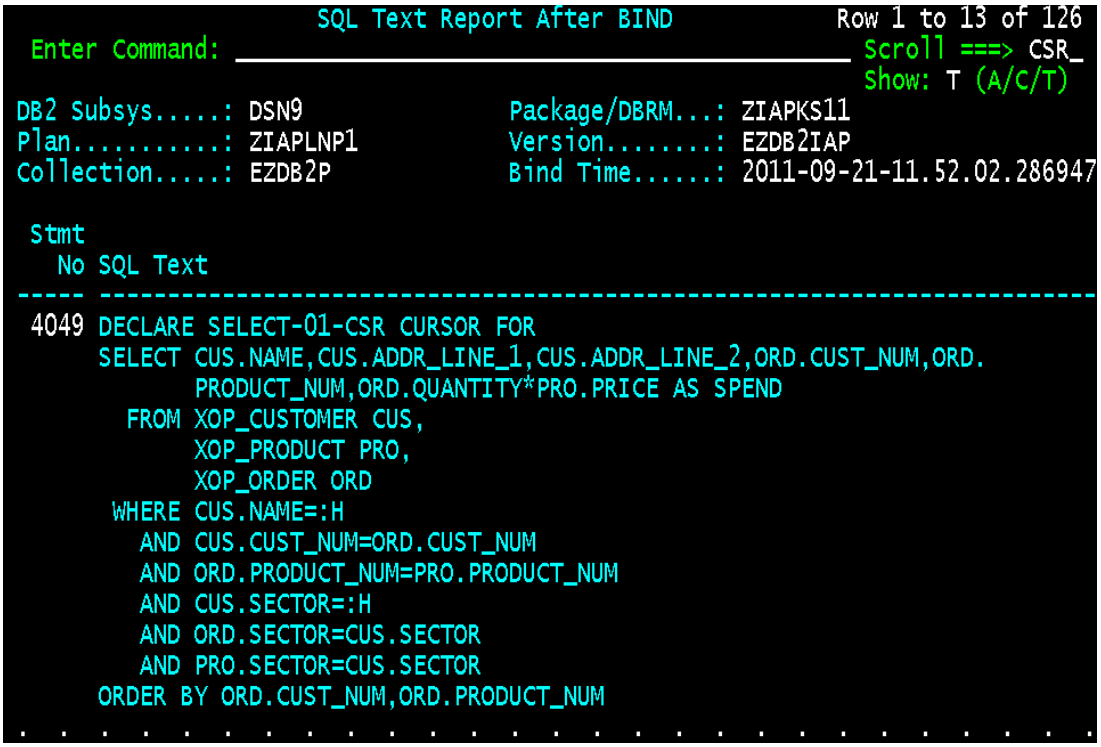

Figure 22 SQL Text Report after BIND

The report will show all of the SQL text for the selected Plan, Package or DBRM.

In the example in Figure 22 the SQL text has been formatted.

Use Page down (PF8) to scroll through the SQL.

This example shows the online version of the reports. When viewing the online reports, you can switch between the Text, Costs and Access Path displays by changing the SHOW to T, C or A respectively.

The following figure shows the Costs display:-

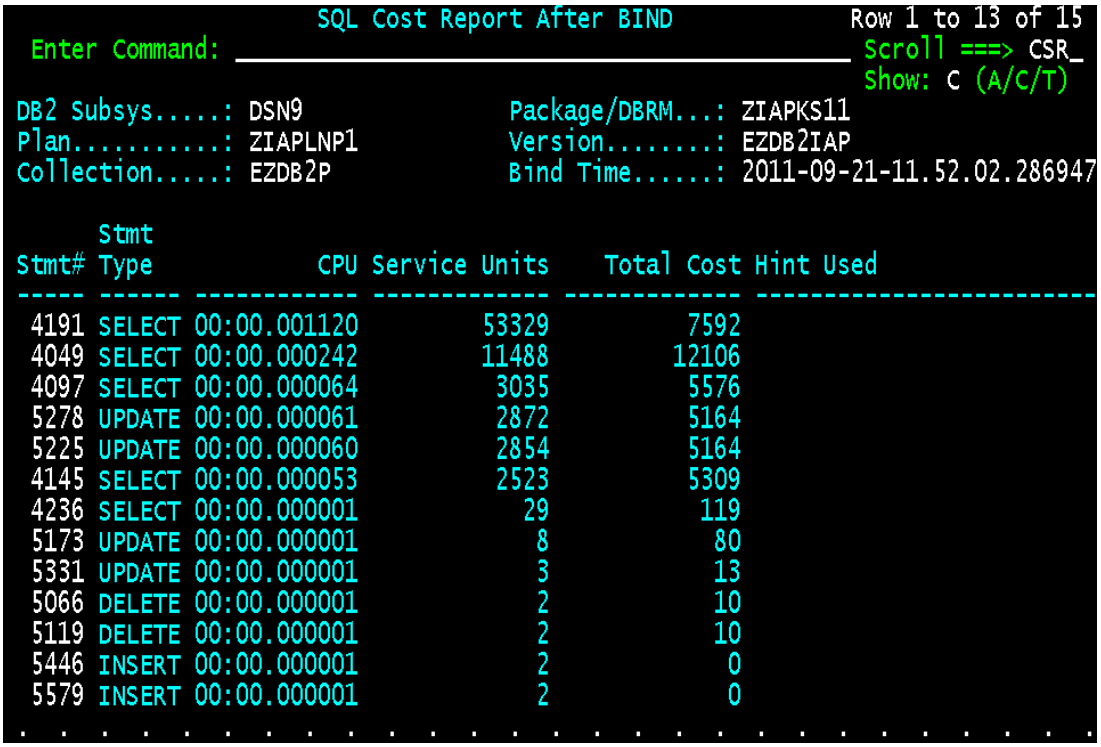

Figure 23 SQL Cost Report After BIND

And the following figure shows the Access Path Report:-

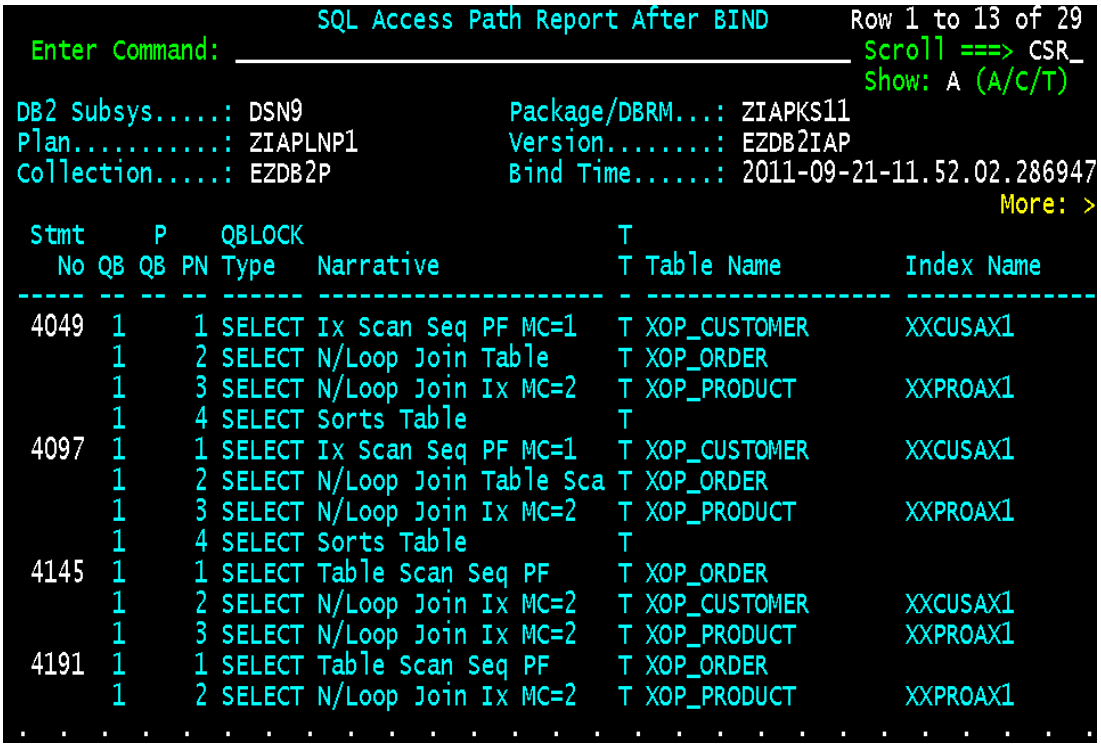

Figure 24 SQL Access Path Report After BIND

To view additional Access Path details, page right <PF11>.

Note that if you use the batch version of this process, the SQL Text, Costs and Access Path reports are written to three separate SYSOUT files, named XOPBINDT, XOPBINDC and XOPBINDA respectively.

6.8 Environment Virtualization

For effective access path evaluation and testing, it has long been standard practise to copy production statistics to the test or sand-box environment. See IBM Documentation Managing Performance (SC19 - 2978-06) Chapter 54 – *Modelling your Production system statistics in a test subsystem*.

This can be easily accomplished using EZ-Stats. Refer to the *EZ-Stats User Guide* for further information.

However, it is recognised that copying statistics does not guarantee accurate simulation of production access paths in test, as there are other factors that affect the access path selection.

These other factors are:-

- CPU speed and number of processors
- Sort Pool, RID Pool and Buffer Pool Sizes
- DSNZPARM optimization settings such as STARJOIN, NPGTHRSH, and PARAMDEG

It is now possible to SIMULATE these values in a test environment using a combination of DSNZPARM settings and Monitoring PROFILES as described below.

Note: This functionality was implemented in IBM APARS PM26475 for DB2 9 and PM26973 for DB2 10. For future releases, it is expected that this functionality will be part of the base product.

For further information about the concepts of **Environment Virtualization** refer to the Appendix Section Concepts of Environment Virtualization.

EZ-DB2 provides a user friendly interface to take advantage of this IBM functionality.

To invoke the Environment Virtualization panel select option 6.8. The following panel is displayed:-

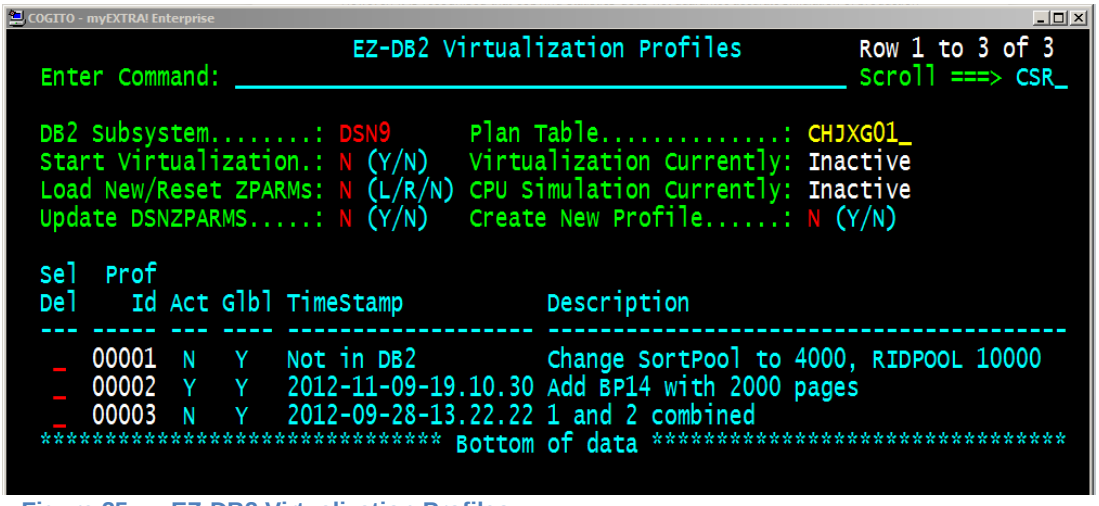

Figure 25 EZ-DB2 Virtualization Profiles

The following information is displayed on the panel.

DB2 Sub-System Specify the name of the DB2 subsystem.

Plan Table Specify the Owner of the Bind Plan table to be used by EZ-DB2 to issue an Explain in order to retrieve the Current values for CPU_SPEED, CPU_COUNT, RIDPOOL and SORT_POOL Size.

Refer to section Obtaining the Current Values.

Start/Stop Virtualization

After creating one or more PROFILES enter Y to start the Profile(s).

EZ-DB2 issues the START PROFILE command.

All PROFILES defined as “active” will be started. It is not possible to start a specific profile.

Once Profiles are started, the literal on this panel changes to “Stop Virtualization”. Enter “Y” to issue the STOP PROFILE command. Any active Profiles will be stopped.

Virtualization Currently

Inactive Enter Y in the Start Virtualization option to START Profiling.

Active Enter Y in the Stop Virtualization option to STOP profiling.

Load New/Reset ZPARMS

After generating a new DSNZPARM module (see Update DSNZPARMS) you must issue the DB2 command to activate the new DSNZPARM settings. This option is used to activate the new DSNZPARM settings, or to reset them to the Startup values.

L Load the new DSNZPARM settings. EZ-DB2 issues the DB2 command:-

- SET SYSPARM LOAD(EZVZPARM)

R Reset the DSNZPARM parameters back to the original settings. EZ-DB2 issues the DB2 Command:-

-SET SYSPARM STARTUP

CPU Simulation Currently

The current status of the DSNZPARM simulation.

Active DSNZPARM CPU simulation is currently active.

Enter R in the Load New/Reset ZPARMS option to set the DSNZPARM settings back to the startup values.

Inactive DSNZPARM simulation is currently inactive.

Enter L in the Load New/Reset ZPARMS option to set the DSNZPARM settings to the simulated values.

Update DSNZPARMS

Enter Y to display the Update DSNZPARMS panel.

The Update DSNZPARM panel will be displayed as described below. See Update DSNZPARMS.

Create New Profile

Enter Y to display the Create PROFILE panel.

The Create New Profile Panel will be displayed as described below. See Create New Profile.

Select Delete

S Select a currently defined PROFILE. The Create Profile panel is displayed.

D Delete a currently defined PROFILE.

Prof ID

The PROFILE Identifier.

Act

Whether this PROFILE is defined as ACTIVE.

Y The PROFILE is defined as ACTIVE. This PROFILE will be started when you set the Start \virtualization option to Y.

N This PROFILE is not currently defined as ACTIVE.

Glbl

Global Indicator.

Y This is a GLOBAL Profile.

N Indicates that this PROFILE is not a Global Profile, rather it is associated with a particular Authid, Plan Name, Collid, Package Name or IPAddr.

These Profiles will still be affected by any START/STOP PROFILE commands issued by EZ-DB2.

Timestamp

The TIMESTAMP of the last update to this PROFILE.

Description

A description for this Profile as entered on the CREATE Profile panel.

Update DSNZPARMS

The following figure shows the panel displayed when you select the Update DSNZPARMS option:-

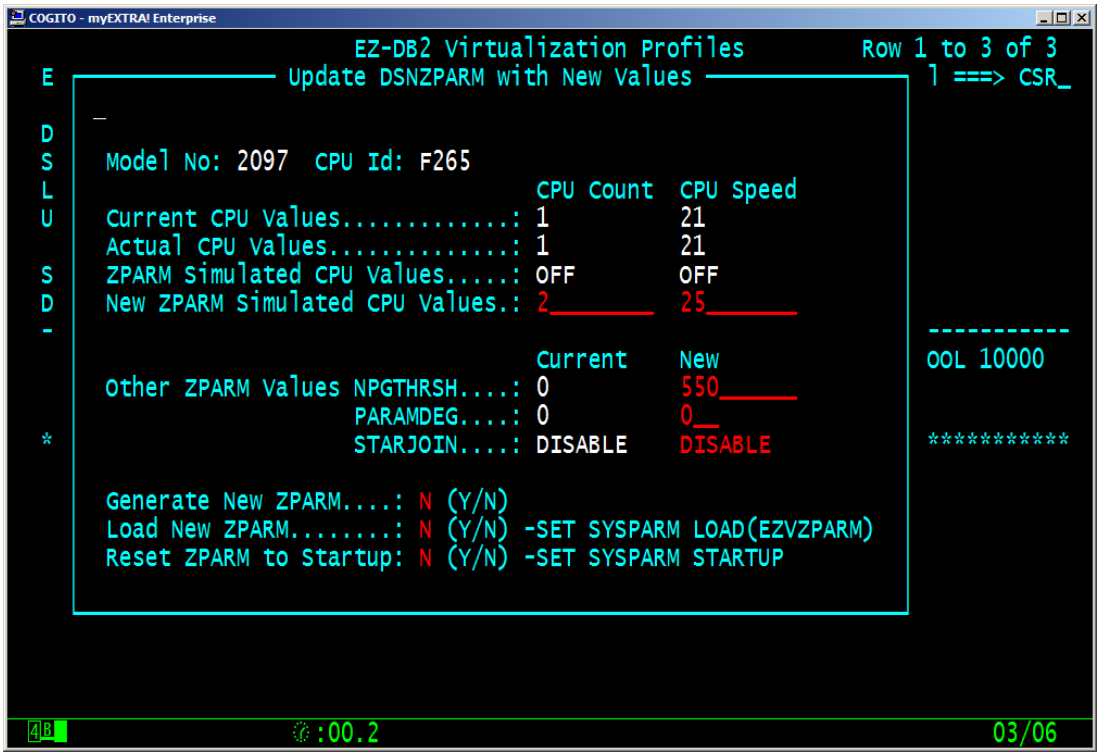

Figure 26 Update DSNZPARM with new values

This panel is used to define the required parameters for DSNZPARM.

Once the required DSNZPARM Parameters are specified you must regenerate and then re-load the DSNZPARM module. See the option Generate New ZPARM and Load New ZPARM. To reset the DSNZPARM back to the original setting, use the Reset ZPARM to Startup option.

The following information is displayed on the panel.

Model No		The current CPU Model is displayed.
CPU Id		The current CPU identifier is displayed.
Current CPU Values	**CPU Count**	The Current Number of CPUs.
	CPU Speed	The Current CPU Speed.
Actual CPU Values	**CPU Count**	The Actual Number of CPUs.
	CPU Speed	The Actual CPU Speed.
ZPARM Simulated CPU Values	**CPU Count**	The Simulated Value for CPU Count if the New ZPARM has been activated. See Load New ZPARM option.

		OFF if the new ZPARM has not been activated or has been reset to startup values, See Reset ZPARM to Startup.
	CPU Speed	The Simulated Value for CPU Speed if the New ZPARM has been activated. See Load New ZPARM option. OFF if the new ZPARM has not been activated or has been reset to startup values, See Reset ZPARM to Startup.
New ZPARM Simulated CPU Values	**CPU Count**	Set this to the CPU Count to be Simulated. This value is used as the DSNZPARM parameter:- SIMULATED_CPU_COUNT Generate the DSNZPARM and Load the New ZPARM for this simulated value to take effect. See the option Generate New ZPARM and Load New ZPARM.
	CPU Speed	Set this to the CPU Speed to be Simulated. This value is used as the DSNZPARM parameter- SIMULATED_CPU_SPEED Generate the DSNZPARM and Load the New ZPARM for this simulated value to take effect. See the option Generate New ZPARM and Load New ZPARM.
NPGTHTRSH	**Current**	The Current Value for NPGTHRSH
	New	The New Value for NPGTHRSH. Set this to the same value as the Production System you are Simulating. Generate the DSNZPARM and Load the New ZPARM for this simulated value to take effect. See the option Generate New ZPARM and Load New ZPARM.
PARAMDEG	**Current**	The Current Value for PARAMDEG
	New	The New Value for PARAMDEG. Set this to the same value as the Production System you are Simulating.

Generate the DSNZPARM and Load the New ZPARM for this simulated value to take effect.

See the option Generate New ZPARM and Load New ZPARM.

STARJOIN — **Current** — The Current Value for STARJOIN

New — The New Value for STARTJOIN. Set this to the same value as the Production System you are Simulating.

Generate the DSNZPARM and Load the New ZPARM for this simulated value to take effect.

See the option Generate New ZPARM and Load New ZPARM.

Generate New ZPARM — Enter Y to Generate the Job stream to assemble and link-edit the DSNZPARM module with the new (simulated) values. The modified DSNZPARM load module is generated with a name of EZVZPARM.

Load New ZPARM — Enter Y to load the new DSNZPARM. EZ-DB2 issues the command

-SET SYSPARM LOAD(EZVZPARM)

Reset ZPARM to Startup — Enter Y to reset the DSNZPARM back to the original settings. EZ-DB2 issues the command

-SET SYSPARM STARTUP

Create New Profile

The following figure shows the panel displayed when you select the Create New Profile option:-

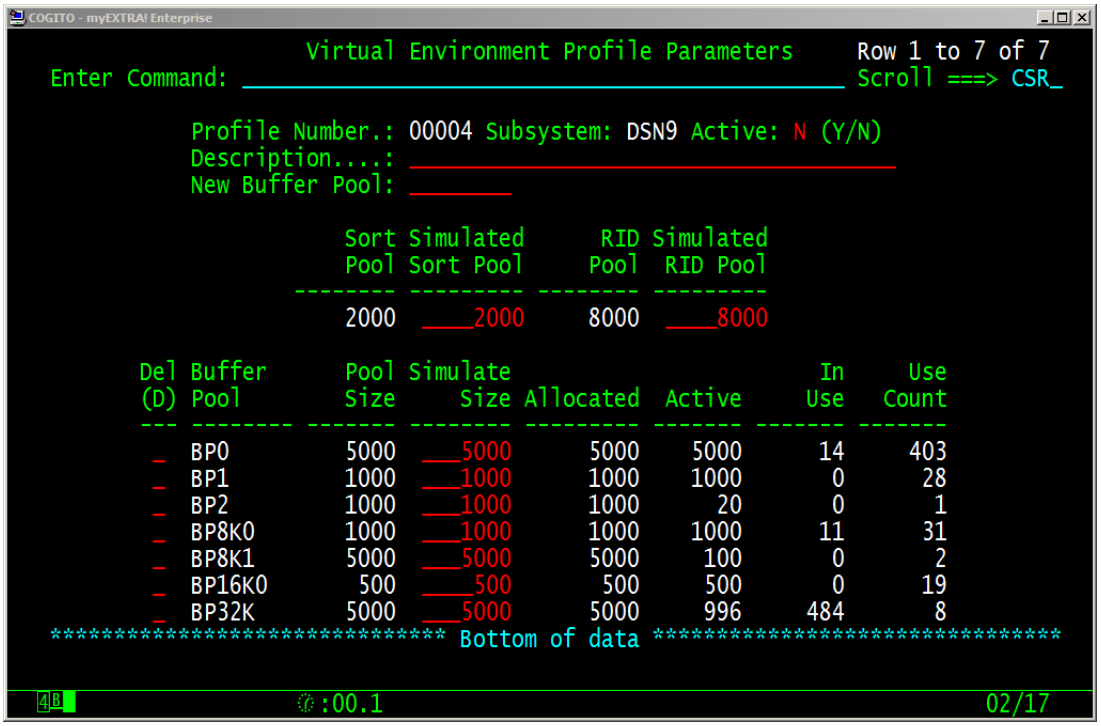

Figure 27 Virtual Environment Profile Parameters

The following information is displayed on the panel:-

Profile Number

The next available PROFILE ID is allocated.

The PROFILE ID is used as the key to link the DSN_PROFILE_TABLE to the DSN_PROFILE_ATTRIBUTES table entry that will be created for this profile. Refer to the IBM documentation Managing Performance for further information.

The PROFILE ID is also displayed in the REASON column in DSN_STATEMNT_TABLE to indicate that the PROFILE was active during an EXPLAIN process.

Subsystem

The DB2 subsystem Identifier

Active

Indicate whether this PROFILE is currently active.

Y The profile is currently active.

N The Profile is currently inactive.

After Creating a Profile you must use the option Start Virtualization to Start the Profiles. All Active Profiles will be started. It is not possible to Start Selected Profiles, therefore ensure that only the Profile(s) you mean to start are set to Active.

Description	You may enter a description for this Profile.
New Buffer Pool	Optionally enter the name of a new Buffer Pool to be created. The Buffer Pool will be added to the list of Buffer Pool described later.
Sort Pool	The Current Sort Pool Size on this DB2 subsystem.
Simulated Sort Pool	Optionally specify a new Sort Pool Size if you wish to simulate the Sort Pool Size with this Profile. The Simulated Sort Pool size will be used by any EXPLAINS on this DB2 Subsystem while this Profile is Active and Started.
Rid Pool	The Current RID Pool Size on this DB2 subsystem.
Simulated RID Pool	Optionally specify a new RID Pool Size if you wish to simulate the RID Pool Size with this Profile. The Simulated RID Pool size will be used by any EXPLAINS on this DB2 Subsystem while this Profile is Active and Started.

There follows a list of entries for each buffer pool defined on this subsystem, including any New Buffer Pool you may have created in this Profile. The following information is displayed:-

Del(D)	Enter D to Delete this Buffer Pool.
Buffer Pool	The Name of the Buffer Pool.
Pool Size	The Current Size of this Buffer Pool in K Bytes.
Simulate Size	Optionally specify a new Buffer Pool Size if you wish to simulate the Buffer Pool Size with this Profile. The Simulated Buffer Pool size will be used by any EXPLAINS on this DB2 Subsystem while this Profile is Active and Started.
Allocated	The number of allocated buffers in an active buffer pool.
Active	The number of active buffers in the buffer pool.
In Use	The number of currently active buffers in the buffer pool. These buffers cannot be stolen.
Use Count	The number of open table spaces or index spaces that use this buffer pool. A value of zero means the pool is inactive.

Appendix

A Appendix

A.1 Creating the Explain Database

With effect from EZ-DB2 R6.10 *it is no longer mandatory to have an Explain database in the Warehouse environment for all EZ-DB2 components. You may choose NOT to explain the SQL when importing the SQL statements into the SQL Warehouse. You might do this because you have already generated the access path information in the Live environment using the* ***EZ-Tracer*** *component. You do not need to generate the access paths again when importing the SQL into the Warehouse. You could use* ***EZ-Alerts*** *to view the Live access paths and Alerts for the SQL in the SQL Warehouse. You might also use* ***EZ-Impact Analyzer*** *to export the SQL in the Warehouse to other environments and generate other versions of the access paths in these other environments for subsequent comparison with the live access path. In this scenario, you will need Explain databases in the other environments as applicable, but you will not need to have explained the SQL in the Warehouse environment. See* ***EZ-Impact Analyzer Reference Guide*** *for further information.*

You DO need to Explain the SQL in the Warehouse environment to use ***EZ-Index Analyzer*** *and* ***EZ-Index Optimizer****.*

To enable EZ-DB2 components in the analysis environment to EXPLAIN and cost SQL captured in the traces of the Live Subsystem, a Version of the Live Database referred to as the EZ-DB2 *Explain database* must be created in the EZ-DB2 Warehouse subsystem. This *Explain database* should be created using the current DDL that defines the Live database so that all Tables, Views and Indexes match the Live configuration.

EZ-Alerts, EZ-Index Analyzer and EZ-Impact Analyzer can use an existing database in the Warehouse (Analysis) environment suitable for Explaining the applicable workload. That is, it must match the data definitions for the selected databases. No sample data is required. In the Workload Analysis SQL and Statistics Controls panel, you decide whether the Statistics on the Explain database will be updated with the Statistics unloaded from the Live database.

EZ-XOP Index Optimizer cannot use an existing test database; instead it requires its own dedicated copy of the Live database(s) to be optimized. This is referred to as the EZ-XOP *Clone* database. Initially, this is created using the Live database definitions only. No sample data is required until you have decided which tables are going to be included in the index optimization process. However, the Clone database used by EZ-XOP must be a database dedicated to EZ-XOP. An existing test/QA database used by other processes will not be suitable, as EZ-XOP will be dynamically creating and dropping indexes during the index evaluation process.

It is recommended that you use your site standard procedures and available 3rd Party tools to create the Explain database.

*Ensure that when you migrate the DDL, that PRIQTY and SECQTY are eliminated from the DDL streams****.***

The recommended method for defining the Explain database is as follows:-

- In the Live environment, generate DDL for each database required with all dependant objects WITHOUT Views and RI.
- In the EZ-DB2 analysis environment, Create the databases required, first eliminating the PRIQTY/SECQTY.
- In the Live environment, generate the DDL for each database required for ONLY the Views and RI.
- In the EZ-DB2 analysis environment, create the Views and RI for each database.

A.2 Job Step Details - Select and Import SQL Traces

ALLOCDS - Allocate required datasets

Dataset allocation step

Compress

This step is not invoked if you have not requested PRESERVE CURRENT STATISTICS. See Workload Analysis SQL and Statistics Controls.

The Statistics Archive file is compressed in preparation for the subsequent statistics archive step.

RUNSTATD - Archive Current Catalog Statistics

This step is not invoked if you have not requested PRESERVE CURRENT STATISTICS. See Workload Analysis SQL and Statistics Controls.

The Statistics for the selected databases are saved to the Statistics Warehouse Archive database. Refer to the EZ-Stats User Guide - XOPSTATD processing for further information.

RUNSTATU - Load Catalog Statistics

This step is not invoked if you have not requested USE STATISTICS FROM THE LIVE ENVIRONMENT. See Workload Analysis SQL and Statistics Controls.

This step loads the required System Catalog Statistics into the EZ-DB2 analysis environment. These catalog statistics were previously saved to the EZ-DB2 Statistics Warehouse when creating the Tracer Workload. Refer to the EZ-Tracer Reference for further information.

For further details about XOPSTATU processing refer to the EZ-Stats User Guide.

Load Catalog Statistics - Missing Objects

It is intended that the Explain database is identical in definition to the corresponding Live DBNAME in the Live environment. If that is not the case, the Load Catalog Statistics LOGFILE may include a MISSING OBJECTS report and the job will terminate with RC=16.

Refer to the EZ-Stats Guide for further details about Missing Objects.

In the event that you receive missing object reports, you must correct the database definition of the Explain database and restart the Select SQL from Traces process.

Load Catalog Statistics - Qualifier Translation Error

XOPSTATU calls an internal routine (XOPSQUAL) which performs the required qualifier translation permitting different qualifiers to be used between the Live and corresponding Explain databases.

Refer to the EZ-Stats User Guide for further details about qualifier translation.

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

In this case, the Load Catalog Statistics 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 28 XOPSQUAL - Qualifier Translation Failure

Typically, errors such as this would occur for example if

- The database definition of the Explain database is not in synch with the live system definition, e.g., there are missing tables and / or indexes.
- A table exists in the Explain 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 Update Qualifiers (if required) (option 6.2) to explicitly set the Explain Qualifier names as applicable before re-executing the Select and Import SQL Traces process.

Load Catalog Statistics - SAP Databases

When loading the Catalog Statistics for an SAP database (Refer to EZ-Tracer Reference - *SAP WORKLOAD = YES*), the Load Catalog Statistics process includes special provision for handling unavoidable inconsistencies between the SAP Live database and the analysis Explain database.

Refer to EZ-Stats User Guide - *Processing SAP Workloads* for further information.

RUNELOG - Update SQL Warehouse with SQL Trace data

The XOPELOG step will populate the SQL Warehouse Workload with the selected SQL data.

The following figure (Figure 29) shows a part of a sample report written to Job log file.

The report includes:-

- **SQL Trace Events Processed** - Number of distinct SQL statements processed for the selected SQL Trace Files.

- **SQL Trace Events Inserted** - Number of distinct SQL statements inserted - i.e do not already exist from any previous Trace Files processed.

- **SQL Trace Text Inserted** - Number of SQL texts lines inserted. This only applies to the SQL Trace Files Events Inserted, where each event can have 1 or more text lines.

- **SQL Trace Events Updated** - Number of distinct SQL statements updated - i.e existed from previous SQL Trace Files, so only frequencies are updated.

```
XOPPARM  08.38.48 : EZ-DB2 Release 9.1 APAR Level: CX910163
XOPELOG  08.38.48 : R9.10 Starting Call Attach to DB2 Subsystem DBA2
XOPELOG  08.38.48 : DB2 Version 10.1.0 NFM
XOPELOG  08.38.49 :  Starting Extract, Store & Analyse SQL
XOPELOG  08.38.49 :
XOPELOG  08.38.49 : Trace Numbers To Load:
XOPELOG  08.38.49 :  *ALL TRACES*
XOPELOG  08.38.49 :
XOPELOG  08.38.49 :
XOPELOG  08.38.49 : Processing SQL For Active DBNAMEs
XOPELOG  08.38.49 : DGBLIVA2
XOPELOG  08.38.49 :
XOPSQUAL 08.38.49 : ****************************************************
XOPSQUAL 08.38.49 :  Starting Live to Image Qualifier translation
XOPSQUAL 08.38.49 : ****************************************************
XOPSQUAL 08.38.49 :
XOPSQUAL 08.38.49 : Selected DBNAMEs:
XOPSQUAL 08.38.49 :
XOPSQUAL 08.38.49 : Live DBN                 Image DBN
XOPSQUAL 08.38.49 : ------------------------ ------------------------
XOPSQUAL 08.38.49 : DGBLIVA2                 DGBCOPA2
XOPSQUAL 08.38.49 :
XOPSQUAL 08.38.49 : Existing Qualifiers from QUALS file:
XOPSQUAL 08.38.49 :
XOPSQUAL 08.38.49 : Live               Image              Live
XOPSQUAL 08.38.49 : Qualifier          Qualifier          DBName
XOPSQUAL 08.38.49 : ------------------ ------------------ ----------------------
XOPSQUAL 08.38.49 : SGBLIVA2           SGBCOPA2           DGBLIVA2
XOPSQUAL 08.38.50 :
XOPSQUAL 08.38.50 : Updated QUALS file output:
XOPSQUAL 08.38.50 :
XOPSQUAL 08.38.50 : Live               Image              Live
XOPSQUAL 08.38.50 : Qualifier          Qualifier          DBName
XOPSQUAL 08.38.50 : ------------------ ------------------ ----------------------
XOPSQUAL 08.38.50 : SGBLIVA2           SGBCOPA2           DGBLIVA2
XOPSQUAL 08.38.50 :
XOPSQUAL 08.38.50 : ****************************************************
XOPSQUAL 08.38.50 :  End of Live to Image Qualifier translation
XOPSQUAL 08.38.50 : ****************************************************
XOPSQUAL 08.38.50 :
XOPELOG  08.38.50 :
XOPELOG  08.38.50 : Will Translate qualifier SGBLIVA2 to SGBCOPA2
XOPELOG  08.38.50 :
XOPELOG  08.38.50 : ----------------------------------------
XOPELOG  08.38.50 :  All Data Loaded from Traces
XOPELOG  08.38.51 :
XOPELOG  08.38.51 :  Loaded Trace : IVP Trace Date : 2012-04-18 Trace Time : 03. 39.48
XOPELOG  08.38.51 :
XOPELOG  08.38.51 :  End   Load/Update of Trace Data
```

```
XOPELOG  08.38.51 : ------------------------------------------
XOPELOG  08.38.51 :
XOPEXP2  08.38.52 : 12/06/28 Started
XOPEXP2  08.38.52 : Start of SQL Explain and INSERT Processing
XOPEXP2  08.38.52 :  Summarize & Explain New SQL
XOPEXP2  08.38.52 :
XOPPLAN  08.39.05 :  Starting Load of PLAN Data
XOPPLAN  08.39.16 :  XOP_PLAN_TABLE Rows Deleted :         0
XOPPLAN  08.39.16 :  XOP_PLAN_TABLE Rows Inserted:        44
XOPPLAN  08.39.16 :  End of Load of PLAN Data
XOPEXP2  08.39.16 :  Analyse Tables Hit
XOPEXP2  08.39.16 :  Inserts Analysis
XOPEXP2  08.39.16 :  ----------------------------------------
XOPEXP2  08.39.16 :  Insert Statements Processed = 000000000003
XOPEXP2  08.39.16 :  Insert Subselects Found     = 000000000000
XOPEXP2  08.39.16 :  SQL_TABLE           Inserts = 000000000037
XOPEXP2  08.39.16 :  SQL Explained Successfully  = 000000000019
XOPEXP2  08.39.16 :  SQL Explain Failures        = 000000000000
XOPEXP2  08.39.16 :  ----------------------------------------
XOPEXP2  08.39.16 : End of   SQL Explain and INSERT Processing
XOPELOG  08.39.16 :  Processing Completed
XOPELOG  08.39.16 :
XOPELOG  08.39.16 : ------------------------------------------
XOPELOG  08.39.16 : SQL Trace Events   Processed = 000000000019
XOPELOG  08.39.16 : SQL Filtered on DBNAME       = 000000000000
XOPELOG  08.39.16 : SQL Filtered on QUALIFIER    = 000000000000
XOPELOG  08.39.16 : SQL Filtered on PLAN/DBRM    = 000000000000
XOPELOG  08.39.16 : SQL Filtered on COLL/PACK/VER= 000000000000
XOPELOG  08.39.16 : SQL already loaded           = 000000000000
XOPELOG  08.39.16 : SQL over text size limit     = 000000000000
XOPELOG  08.39.16 : SQL Trace Events    Inserted = 000000000019
XOPELOG  08.39.16 : SQL Trace Text      Inserted = 000000000028
XOPELOG  08.39.16 : SQL Trace Events     Updated = 000000000000
XOPELOG  08.39.16 : SQL Authid Entries  Inserted = 000000000019
XOPELOG  08.39.16 : SQL Authid Entries   Updated = 000000000000
XOPELOG  08.39.16 : SQL Statistics      Inserted = 000000000072
XOPELOG  08.39.16 : SQL Statistics       Updated = 000000000000
XOPELOG  08.39.16 : SQL/Auth Statistics Inserted = 000000000072
XOPELOG  08.39.16 : SQL/Auth Statistics  Updated = 000000000000
XOPELOG  08.39.16 : ------------------------------------------
XOPELOG  08.39.16 :
XOPELOG  08.39.16 : R9.10  End of Extract, Store & Analyse SQL
XOPELOG  08.39.16 :
XOPELOG  08.39.16 : R9.10 Stopping Call Attach
XOPELOG  08.39.16 :  End of Extract, Store & Analyse SQL
XOPELOG  08.39.16 : 12/06/28 Ending            RETURN CODE - 0000
******************************** BOTTOM OF DATA ********************************
```

Figure 29 XOPELOG log file

XOPELOG return codes

RC	Description

0	Processing completed successfully
8	No rows found for database name used.
12	Invalid XOPPARMS parameter. Check LOGFILE for diagnostic information.
16	SQL error. Check LOGFILE for diagnostic information.
20	IO error. Check LOGFILE for diagnostic information.
XX	RUNSTATS error code. Check LOGFILE for diagnostic information.

A.3 Job Step Details - Load Application SQL From Catalog

Extract SQL from Catalog (XOPASQL)

Figure 30 shows a section of a LOGFILE produced by XOPASQL.

```
XOPPARM  09.09.43 : EZ-DB2 Release 9.1 APAR Level: CX910163
XOPASQL  09.09.43 : R9.10 Starting Call Attach to DB2 Subsystem DBA2
XOPASQL  09.09.43 : DB2 Version 10.1.0 NFM
XOPASQL  09.09.43 : ********************************************************
XOPASQL  09.09.43 : * Starting Extract Of SQL Text                          *
XOPASQL  09.09.43 : ********************************************************
XOPASQL  09.09.43 :
XOPASQL  09.09.43 :  Filter parameters:
XOPASQL  09.09.43 :
XOPASQL  09.09.43 :  COLLID           =EZDBP1
XOPASQL  09.09.43 :  PACKAGE          =ZIA%
XOPASQL  09.09.44 :
XOPASQL  09.09.44 : Static SQL extracted for input parameters:
XOPASQL  09.09.44 :  COLLID           =EZDBP1
XOPASQL  09.09.44 :  PACKAGE          =ZIA%
XOPASQL  09.09.44 :
XOPASQL  09.09.44 :  COLLID              PACKAGE   # SQL    VERSION
XOPASQL  09.09.44 :  ------------------  --------  -------  --------------------
XOPASQL  09.09.44 :  EZDBP1              ZIAPKM15       13  EZDB2IAP
XOPASQL  09.09.44 :  EZDBP1              ZIAPKS11       15  EZDB2IAP
XOPASQL  09.09.44 :  EZDBP1              ZIAPKS12       15  EZDB2IAP
XOPASQL  09.09.44 :  EZDBP1              ZIAPKS17       14  EZDB2IAP
XOPASQL  09.09.44 :  EZDBP1              ZIAPKS18       16  EZDB2IAP
XOPASQL  09.09.45 :  EZDBP1              ZIAPKS1A       15  EZDB2IAP
XOPASQL  09.09.45 :  EZDBP1              ZIAPKS1B       15  EZDB2IAP
XOPASQL  09.09.45 :  EZDBP1              ZIAPKS1C       15  EZDB2IAP
XOPASQL  09.09.45 :  ------------------  --------  -------  --------------------
XOPASQL  09.09.45 :
XOPASQL  09.09.45 :  Total SQL written to SQLLOG =      118
XOPASQL  09.09.45 :
XOPASQL  09.09.45 : R9.10 Stopping Call Attach
XOPASQL  09.09.45 :  End of SQL Extract
XOPASQL  09.09.45 : 12/06/28 Ending            RETURN CODE - 0000
```

Figure 30 XOPASQL logfile

The report echoes the Run Time Parameters specified, and reports details of the number of SQL statements that were selected for the requested Plans and Packages.

Expanding the SQL data (XOPTRC2)

After running the SQL extract process (XOPASQL), XOPTRC2 is run to resolve the trace data. This process reads the trace data sequential file (SQLLOG) created by XOPASQL and accesses the DB2 Catalog to resolve SQL statements with the SQL text. The resolved SQL data is written to SQLLOGO.

Figure 31 shows a section of the LOGFILE produced by XOPTRC2.

```
XOPPARM  09.09.46 : EZ-DB2 Release 9.1 APAR Level: CX910163
XOPTRC2  09.09.46 : R9.10 Starting Call Attach to DB2 Subsystem DBA2
XOPTRC2  09.09.46 : DB2 Version 10.1.0 NFM
XOPTRC2  09.09.46 : (9.1152)  Starting SQL Consolidation
XOPTRC2  09.09.46 :
XOPTRC2  09.09.47 : Processing a single trace log file
XOPTRC2  09.09.47 : Consolidate SQL
XOPTRC2  09.09.47 : Consolidate AUTHIDs for non-distributed SQL
XOPTRC2  09.09.47 : Consolidate, first adding table qualifiers
XOPTRC2  09.09.47 : Consolidate differing length IN lists
XOPTRC2  09.09.47 : Consolidate DDF Program Names
XOPTRC2  09.09.47 :
EZX13005 09.09.49 :
EZX13005 09.09.49 : Trace Summary
EZX13005 09.09.49 : -------------
EZX13005 09.09.49 : Total Trace Events read         =        118
EZX13005 09.09.49 : Trace Events for SQL            =        118
EZX13005 09.09.49 :
EZX13005 09.09.49 : SQL processed this workload     =        118
EZX13005 09.09.49 : SQL processed this trace        =        118
EZX13005 09.09.49 : New SQL added to workload       =        118
EZX13005 09.09.49 :
EZX13006 09.09.49 : Tables Referenced By Selected SQL = 000000003
EZX13006 09.09.49 : TableName                        DBName
EZX13006 09.09.49 : ZSIAPRA2 . XOP_CUSTOMER         ZDIAPRA2
EZX13006 09.09.49 : ZSIAPRA2 . XOP_ORDER            ZDIAPRA2
EZX13006 09.09.49 : ZSIAPRA2 . XOP_PRODUCT          ZDIAPRA2
EZX13006 09.09.49 :
EZX13007 09.09.49 : DBNames Referenced By Selected SQL = 000000001
EZX13007 09.09.49 : ZDIAPRA2
EZX13007 09.09.49 :
XOPTRC2  09.09.49 : End of User Reports
XOPTRC2  09.09.49 : - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - - -
XOPTRC2  09.09.49 : Internals Reports
XOPTRC2  09.09.49 : Thread records  = 00000008
XOPTRC2  09.09.49 : High water mark = 00000008
XOPTRC2  09.09.49 : Allocation      = 00050000
XOPTRC2  09.09.49 : Reads           = 00000008
XOPTRC2  09.09.49 : F-reads         = 00000008
XOPTRC2  09.09.49 : Writes          = 00000008
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 : SQL slots       = 00000010
XOPTRC2  09.09.49 : Reallocate x2   = 00000008
XOPTRC2  09.09.49 : Reallocate x3   = 00000000
XOPTRC2  09.09.49 : Reallocate >3   = 00000000
XOPTRC2  09.09.49 : Reallocate HWM  = 00000002
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : SQL prep. recs  = 00000000
XOPTRC2  09.09.49 : SQL exec. recs  = 00000118
XOPTRC2  09.09.49 : In use at end   = 00000118
XOPTRC2  09.09.49 : High water mark = 00000118
XOPTRC2  09.09.49 : Allocation      = 00150000
XOPTRC2  09.09.49 : Reads           = 00000118
```

```
XOPTRC2  09.09.49 : F-reads         = 00000300
XOPTRC2  09.09.49 : Writes          = 00000118
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : SQL stats recs  = 00000118
XOPTRC2  09.09.49 : High water mark = 00000118

XOPTRC2  09.09.49 : Allocation      = 00150000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : SQL prep. recs  = 00000000
XOPTRC2  09.09.49 : SQL exec. recs  = 00000118
XOPTRC2  09.09.49 : In use at end   = 00000118
XOPTRC2  09.09.49 : High water mark = 00000118
XOPTRC2  09.09.49 : Allocation      = 00150000
XOPTRC2  09.09.49 : Reads           = 00000118
XOPTRC2  09.09.49 : F-reads         = 00000300
XOPTRC2  09.09.49 : Writes          = 00000118
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : SQL stats recs  = 00000118
XOPTRC2  09.09.49 : High water mark = 00000118
XOPTRC2  09.09.49 : Allocation      = 00150000
XOPTRC2  09.09.49 : Reads           = 00000118
XOPTRC2  09.09.49 : F-reads         = 00000118
XOPTRC2  09.09.49 : Writes          = 00000118
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : Plan/Pack recs  = 00000008
XOPTRC2  09.09.49 : High water mark = 00000008
XOPTRC2  09.09.49 : Allocation      = 00050000
XOPTRC2  09.09.49 : Reads           = 00000000
XOPTRC2  09.09.49 : F-reads         = 00000008
XOPTRC2  09.09.49 : Writes          = 00000008
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : PPH key recs    = 00000145
XOPTRC2  09.09.49 : High water mark = 00000145
XOPTRC2  09.09.49 : Allocation      = 00500000
XOPTRC2  09.09.49 : Reads           = 00000236
XOPTRC2  09.09.49 : F-reads         = 00000145
XOPTRC2  09.09.49 : Writes          = 00000145
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : Synonym recs    = 00000003
XOPTRC2  09.09.49 : High water mark = 00000003
XOPTRC2  09.09.49 : Allocation      = 00020000
XOPTRC2  09.09.49 : Reads           = 00000163
XOPTRC2  09.09.49 : F-reads         = 00000003
XOPTRC2  09.09.49 : Writes          = 00000003
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : Temp. Tab recs  = 00000000
XOPTRC2  09.09.49 : High water mark = 00000000
XOPTRC2  09.09.49 : Allocation      = 00050000
XOPTRC2  09.09.49 : Reads           = 00000000
XOPTRC2  09.09.49 : F-reads         = 00000000
XOPTRC2  09.09.49 : Writes          = 00000000
XOPTRC2  09.09.49 : Deletes         = 00000000
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : Tokens HWM      = 00000098
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : Pre-normalization hits  =                 0
XOPTRC2  09.09.49 : Post-normalization hits =                91
```

```
XOPTRC2  09.09.49 : Normalization misses    =                27
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : *** : XOPSQLCP File Records Written:              118
XOPTRC2  09.09.49 : *** : XOPSQLFP File Records Written:              237
XOPTRC2  09.09.49 : *** : XOPSQLFP File Records Rewritten:            237
XOPTRC2  09.09.49 : *** : XOPSQLTP File Records Written:               28
XOPTRC2  09.09.49 : *** : XOPSQLTP File Records Rewritten:            145
XOPTRC2  09.09.49 : *** : XOPSQLXP File Records Written:              195
XOPTRC2  09.09.49 : *** : XOPSQLVP File Records Written:                0
XOPTRC2  09.09.49 : *** : XOPSQLSP File Records Written:                0
XOPTRC2  09.09.49 : *** : XOPSQLDP File Records Written:                0
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : *** : XOPSQLCP File Records Read   :                0
XOPTRC2  09.09.49 : *** : XOPSQLFP File Records Read   :              354
XOPTRC2  09.09.49 : *** : XOPSQLTP File Records Read   :              146
XOPTRC2  09.09.49 : *** : XOPSQLXP File Records Read   :              106
XOPTRC2  09.09.49 :
XOPTRC2  09.09.49 : R9.10 Stopping Call Attach
XOPTRC2  09.09.49 :  End of SQL Consolidation
XOPTRC2  09.09.49 : 12/06/28 Ending            RETURN CODE - 0000
XOPTRC2  09.09.49 :
```

Figure 31 **XOPTRC2 logfile**

The report shows the number of SQL trace events processed. The number of SQL statements that were not expanded for various reasons is also shown.

The names of the DB2 tables and database names referenced by the selected SQL statements are also reported.

Extract System Catalog Statistics (XOPSTATD)

This step is not invoked if you have not requested USE STATISTICS FROM THE LIVE ENVIRONMENT. See Workload Analysis SQL and Statistics Controls.

If you have specified Use Statistics From Live Environment then XOPSTATD is executed to extract the System Catalog statistics for all applicable databases in the "live" environment to the Statistics Warehouse.

For further details about XOPSTATU processing refer to the EZ-Stats User Guide.

Archive System Catalog Statistics (XOPSTATD)

This step is not invoked if you have not requested PRESERVE THE CURRENT STATISTICS. See Workload Analysis SQL and Statistics Controls.

If you have specified Preserve the Current Statistics then XOPSTATD is executed to archive the statistics for all applicable databases in the analysis environment to the Statistics Warehouse archive.

For further details about XOPSTATU processing refer to the EZ-Stats User Guide.

Load System Catalog Statistics (XOPSTATU)

This step is not invoked if you have not requested USE STATISTICS FROM THE LIVE ENVIRONMENT. See Workload Analysis SQL and Statistics Controls.

If You have specified Use Statistics From Live Environment then XOPSTATU is executed to load the System Catalog statistics for all applicable databases from the Statistics Warehouse to the analysis environment.

Update SQL Warehouse with SQL from Catalog (XOPELOG)

The XOPELOG step will populate the SQL Warehouse Workload with the selected SQL data.

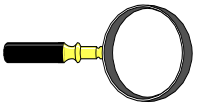

XOPELOG will perform automatic qualifier translation when loading the SQL into the SQL Warehouse. That is, the table name and other qualifiers will be converted to the corresponding names as applicable for the Explain database.

The Explain database is a copy of the "live" database created in the analysis environment against which EZ-DB2 can explain the SQL workload.

The XOPELOG log file will report

- The qualifier translation it has performed, showing the "live" and corresponding analysis system qualifiers.
- Statistics showing the number of SQL statements loaded into the SQL Warehouse, number of statements successfully Explained, and number of Explain failures.

The report also includes

- **SQL Trace Events Processed** - Number of distinct SQL statements processed for the selected SQL Trace Files.
- **SQL Trace Events Inserted** - Number of distinct SQL statements inserted - i.e do not already exist from any previous Trace Files processed.
- **SQL Trace Text Inserted** - Number of SQL texts lines inserted. This only applies to the SQL Trace Files Events Inserted, where each event can have 1 or more text lines.
- **SQL Trace Events Updated** - Number of distinct SQL statements updated - i.e existed from previous SQL Trace Files, so only frequencies are updated.

Where Explain failures have occurred, the log will also include detailed information about the statement and the Explain.

The following figure shows a section of the log file written by XOPELOG:-

```
XOPPARM  09.10.06 : EZ-DB2 Release 9.1 APAR Level: CX910163
XOPELOG  09.10.06 : R9.10 Starting Call Attach to DB2 Subsystem DBA2
XOPELOG  09.10.06 : DB2 Version 10.1.0 NFM
```

```
XOPELOG  09.10.06 :  Starting Extract, Store & Analyse SQL
XOPELOG  09.10.06 :
XOPELOG  09.10.06 : Trace Numbers To Load:
XOPELOG  09.10.06 : 0001
XOPELOG  09.10.06 :
XOPELOG  09.10.06 :
XOPELOG  09.10.06 : Processing SQL For Active DBNAMEs
XOPELOG  09.10.06 : ZDIAPRA2
XOPELOG  09.10.06 :
XOPSQUAL 09.10.06 : ******************************************************
XOPSQUAL 09.10.06 :  Starting Live to Image Qualifier translation
XOPSQUAL 09.10.06 : ******************************************************
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Selected DBNAMEs:
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Live DBN                 Image DBN
XOPSQUAL 09.10.06 : ------------------------ ------------------------
XOPSQUAL 09.10.06 : ZDIAPRA2                 ZDIAPRA2
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Existing Qualifiers from QUALS file:
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Live                Image               Live
XOPSQUAL 09.10.06 : Qualifier           Qualifier           DBName
XOPSQUAL 09.10.06 : ------------------- ------------------- ----------------------
XOPSQUAL 09.10.06 : ZSIAPRA2            ZSIAPRA2            ZDIAPRA2
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Updated QUALS file output:
XOPSQUAL 09.10.06 :
XOPSQUAL 09.10.06 : Live                Image               Live
XOPSQUAL 09.10.06 : Qualifier           Qualifier           DBName
XOPSQUAL 09.10.06 : ------------------- ------------------- ----------------------
XOPSQUAL 09.10.06 : ZSIAPRA2            ZSIAPRA2            ZDIAPRA2
XOPSQUAL 09.10.07 :
XOPSQUAL 09.10.07 : ******************************************************
XOPSQUAL 09.10.07 :  End of Live to Image Qualifier translation
XOPSQUAL 09.10.07 : ******************************************************
XOPSQUAL 09.10.07 :
XOPELOG  09.10.07 :
XOPELOG  09.10.07 : ----------------------------------------
XOPELOG  09.10.07 :  Start Load/Update of Trace Data
XOPELOG  09.10.07 :
XOPELOG  09.10.07 :  Start of Trace :  Trace Date : 2012-06-28 Trace Time : 09.09.44
XOPELOG  09.10.09 :
XOPELOG  09.10.09 :  End   Load/Update of Trace Data
XOPELOG  09.10.09 : ----------------------------------------
XOPELOG  09.10.09 :
XOPEXP2  09.10.09 : 12/06/28 Started
XOPEXP2  09.10.09 : Start of SQL Explain and INSERT Processing
XOPEXP2  09.10.09 :  Summarize & Explain New SQL
XOPEXP2  09.10.09 :
XOPPLAN  09.10.21 :  Starting Load of PLAN Data
XOPPLAN  09.10.24 :  XOP_PLAN_TABLE Rows Deleted :         0
XOPPLAN  09.10.24 :  XOP_PLAN_TABLE Rows Inserted:       226
XOPPLAN  09.10.24 :  End of Load of PLAN Data
XOPEXP2  09.10.24 :  Analyse Tables Hit
XOPEXP2  09.10.25 :  Inserts Analysis
XOPEXP2  09.10.25 :  ----------------------------------------
XOPEXP2  09.10.25 :  Insert Statements Processed = 000000000024
XOPEXP2  09.10.25 :  Insert Subselects Found     = 000000000000
XOPEXP2  09.10.25 :  SQL_TABLE           Inserts = 000000000195
XOPEXP2  09.10.25 :  SQL Explained Successfully  = 000000000118
XOPEXP2  09.10.25 :  SQL Explain Failures        = 000000000000
XOPEXP2  09.10.25 :  ----------------------------------------
```

```
XOPEXP2  09.10.25 : End of   SQL Explain and INSERT Processing
XOPELOG  09.10.25 :  Processing Completed
XOPELOG  09.10.25 :
XOPELOG  09.10.25 : ----------------------------------------
XOPELOG  09.10.25 : SQL Trace Events   Processed = 000000000118
XOPELOG  09.10.25 : SQL Filtered on DBNAME       = 000000000000
XOPELOG  09.10.25 : SQL Filtered on QUALIFIER    = 000000000000
XOPELOG  09.10.25 : SQL Filtered on PLAN/DBRM    = 000000000000
XOPELOG  09.10.25 : SQL Filtered on COLL/PACK/VER= 000000000000
XOPELOG  09.10.25 : SQL already loaded           = 000000000000
XOPELOG  09.10.25 : SQL over text size limit     = 000000000000
XOPELOG  09.10.25 : SQL Trace Events    Inserted = 000000000118
XOPELOG  09.10.25 : SQL Trace Text      Inserted = 000000000156
XOPELOG  09.10.25 : SQL Trace Events     Updated = 000000000000
XOPELOG  09.10.25 : SQL Authid Entries  Inserted = 000000000118
XOPELOG  09.10.25 : SQL Authid Entries   Updated = 000000000000
XOPELOG  09.10.25 : SQL Statistics      Inserted = 000000000000
XOPELOG  09.10.25 : SQL Statistics       Updated = 000000000000
XOPELOG  09.10.25 : SQL/Auth Statistics Inserted = 000000000000
XOPELOG  09.10.25 : SQL/Auth Statistics  Updated = 000000000000
XOPELOG  09.10.25 : ----------------------------------------
XOPELOG  09.10.25 :
XOPELOG  09.10.25 : R9.10  End of Extract, Store & Analyse SQL
XOPELOG  09.10.25 :
XOPELOG  09.10.25 : R9.10 Stopping Call Attach
XOPELOG  09.10.25 :  End of Extract, Store & Analyse SQL
XOPELOG  09.10.25 : 12/06/28 Ending            RETURN CODE - 0000
```

Figure 32 **XOPELOG - Logfile**

Restore System Catalog Statistics (XOPSTATU)

This step is not invoked if you have not requested PRESERVE THE CURRENT STATISTICS. See Workload Analysis SQL and Statistics Controls.

If you have specified Preserve the Current Statistics then XOPSTATU is executed at the end of the process to restore the statistics from the Statistics Warehouse archive to the analysis environment.

A.4 Concepts of Environment Virtualization

Obtaining the Current Values

The IBM documentation suggests using SQL as below to obtain the Production values for CPU Speed, number of processors, number of RID blocks and Sort Pool Size.

```
SET CURRENT DEGREE='ANY';
      EXPLAIN ALL SET QUERYNO=6475 FOR
      SELECT * FROM SYSIBM.SYSDUMMY1;

      SELECT
      HEX(SUBSTR(IBM_SERVICE_DATA,17,2)) AS CPU_COUNT,
      HEX(SUBSTR(IBM_SERVICE_DATA,69,4)) AS CPU_SPEED,
      HEX(SUBSTR(IBM_SERVICE_DATA,13,4)) AS RIDPOOL,
      HEX(SUBSTR(IBM_SERVICE_DATA,9,4))  AS SORT_POOL
      FROM PLAN_TABLE
      WHERE QUERYNO=6475;
```

Note though that the CPU_COUNT will only be populated if the explained SQL uses parallelism (as this is the only time that the number of CPUs becomes relevant). Of course, the simple explain shown here will NEVER be executed in parallel – the IBM recommendation is to find an SQL statement in PLAN_TABLE where parallelism IS used, and re-explain that!

Alternatively, ask your Systems Programmer how many CPU engines there are.

- You can use the DISPLAY BUFFERPOOL command to get information about the Buffer Pool configuration in the production environment.

- You can use the EZ-DB2 Option 9.8 to display the current DSNZPARM settings. Refer to the *EZ-DB2 Housekeeping and Controls User Guide* for further information.

DSNZPARM Simulation

It is now possible to simulate the CPU Speed and number of Processors using the new DSNZPARM parameters SIMULATED_CPU_SPEED and SIMULATED_CPU_COUNT.

- SIMULATED_CPU_SPEED

  Specify the value that was captured from the Production system as CPU_SPEED

- SIMULATED_CPU_COUNT

  Specify the value that was captured from the Production system as CPU_COUNT

- NPGTHRSH

  Specify the same value that is used in the Production environment

- PARAMDEG

Specify the same value that is used in the Production environment

- STARJOIN

  Specify the same value that is used in the Production environment

The DSNZPARM needs to be reassembled. There is no need to recycle the DB2 subsystem as the are online changeable parameters. The DSNZPARM settings can be refreshed using the command

-SET SYSPARM

As you will see later, all of this is handled by the EZ-DB2 Environment Virtualisation functionality.

Using PROFILES

The Sort Pool, RID Pool and Buffer Pool sizes may be simulated using Monitoring PROFILES.

To make use of the PROFILE functionality a complete set of profile tables and related indexes must exist in the test system. The complete set of profile tables and related indexes includes the following objects:-

- SYSIBM.DSN_PROFILE_TABLE
- SYSIBM.DSN_PROFILE_HISTORY
- SYSIBM.DSN_PROFILE_ATTRIBUTES
- SYSIBM.DSN_PROFILE_ATTRIBUTES_HISTORY
- SYSIBM.DSN_PROFILE_TABLE_IX_ALL
- SYSIBM.DSN_PROFILE_TABLE_IX2_ALL
- SYSIBM.DSN_PROFILE_ATTRIBUTES_IX_ALL

These are created when you run job DSNTIJSG during the DB2 installation or migration.

In addition, a complete set of EXPLAIN tables must exist with the SYSIBM qualifier in the production subsystem and the test subsystem.

Such a profile consists of a single row in DSN_PROFILE_TABLE which defines under what circumstances a profile will be considered by the optimiser and one or more rows in DSN_PROFILE_ATTRIBUTES which define the specific attributes to be used.

These tables are user tables (like the other extended explain tables) so must be created before they are first used. DDL is provided in the manual Managing Performance (SC19-2978).

DSN_PROFILE_TABLE

A profile is created by an INSERT statement. Any unique PROFILEID value can be used. The INSERT statement creates a global profile for a single DB2 subsystem that is active when profile monitoring is on:

```
INSERT INTO SYSIBM.DSN_PROFILE_TABLE (
                PROFILEID , PROFILE_ENABLED)
                VALUES (4713, 'Y');
```

As this is a global profile, no other data is required

DSN_PROFILE_ATTRIBUTES_TABLE

- The specific attributes are then defined by inserting rows into the DSN_PROFILE_ATTRIBUTES table:
- Insert BUFFER POOL values to model into DSN_PROFILE_ATTRIBUTES for the global parameter profile that was just defined.
- Buffer Pool assignments in Test should be consistent with the buffer pool assignments in production.

Let's suppose, BP0 will use a value of 25000 and BP8K0 will use a value of 2500 and override the actual bufferpool sizes defined in the subsystem when determining the access path. Those actual bufferpool sizes are not changed. The bufferpool assignments for tables in the test system need to be consistent with bufferpool assignments in production. It is not important that the exact bufferpool used on the test system is the same as the production system. However, if a table is assigned to a bufferpool with VPSIZE 10000 in production, it should also be assigned to a bufferpool modeled at size 10000 in the modeled environment. This is also true for indexes.

- Use the Production SORT_POOL and RIDPOOL size values to Insert RID POOL and Sort Pool values into the DSN_PROFILE_ATTRIBUTES table for the Global Profile that was previously defined.

For example

```
INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
        (PROFILEID,KEYWORDS,ATTRIBUTE1,ATTRIBUTE2)
        VALUES
        (4713, 'BP0',NULL, 25000);
        INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
        (PROFILEID,KEYWORDS,ATTRIBUTE1,ATTRIBUTE2)
        VALUES
        (4713, 'BP8K0',NULL, 2500);
        INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
        (PROFILEID,KEYWORDS,ATTRIBUTE1,ATTRIBUTE2)
        VALUES
        (4713, 'SORT_POOL_SIZE',NULL, 307200);
        INSERT INTO SYSIBM.DSN_PROFILE_ATTRIBUTES
        (PROFILEID,KEYWORDS,ATTRIBUTE1,ATTRIBUTE2)
        VALUES
        (4713, 'MAX_RIDBLOCKS',NULL, 300);
```

Before these changes can be considered by the optimizer, the profile ust be started with the

-START PROFILE command

The PROFILE_ID is used ONLY to link the profile definition with its attributes

It is not possible to start a single profile. All Profiles that are defined (and have a value of "Y" in DSN_PROFILE_TABLE PROILE_ENABLED) will be started.

Validating the PROFILE has been used

Execute an EXPLAIN as usual. In DSN_STATEMNT_TABLE the column REASON will contain the value 'PROFILEID 4713' appended to the existing REASON value for that statement (where the actual PROFILEID will be the value for the profile that was used to provide the virtual environment information)

Of course, these virtualized settings will ONLY be honoured by the DB2 optimizer during an EXPLAIN, if a new plan or package is bound, the REAL values will be used instead

For further information refer to IBM Documentation Managing Performance (SC19 -2978-06) Chapter 53 – Modelling a production environment on a test subsystem.

Use of the subsystem virtualization has been implemented in EZ-DB2 in the SQL Warehouse component providing a simple interface to the IBM implementation.

A.5 Importing CA Detector® data

This section describes the process to load CA Detector® data into an EZ-DB2 Workload.

This process assumes that you have already externalized the data from CA Detector®.

Refer also to the CA Detector® documentation for information about externalizing the SQL data from CA Detector®.

In EZ-Tracer, create a new Workload and DRILL into the Trace Filters panel. In the following example we have created a Trace workload called Detector.

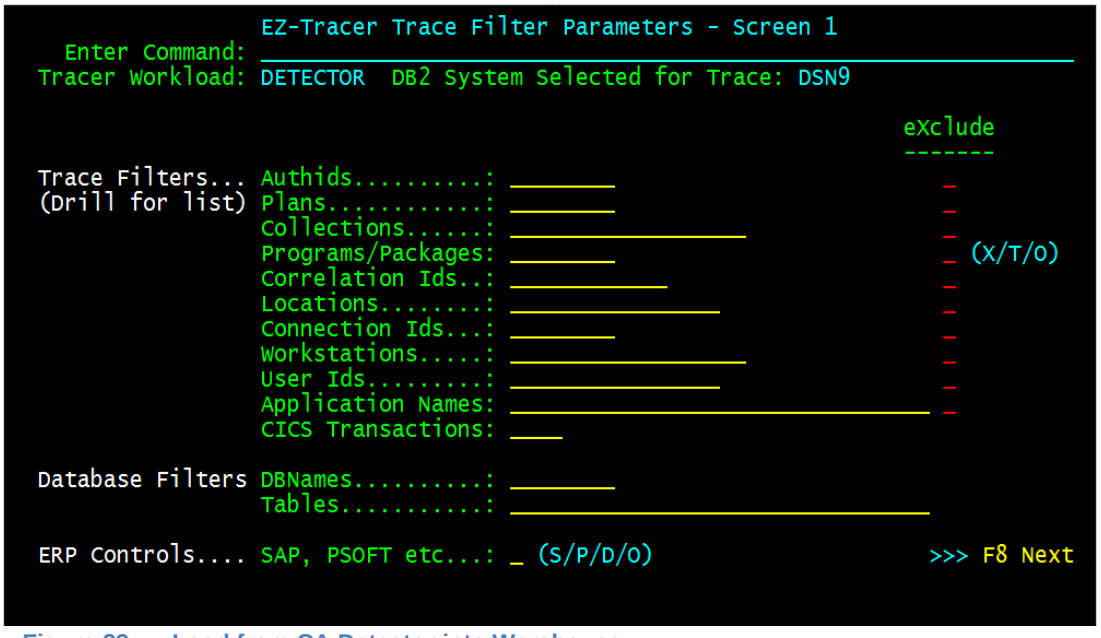

Figure 33 Load from CA Detector into Warehouse

In the ERP Controls, enter **D** for Detector. The following Pop-Up window is displayed:-

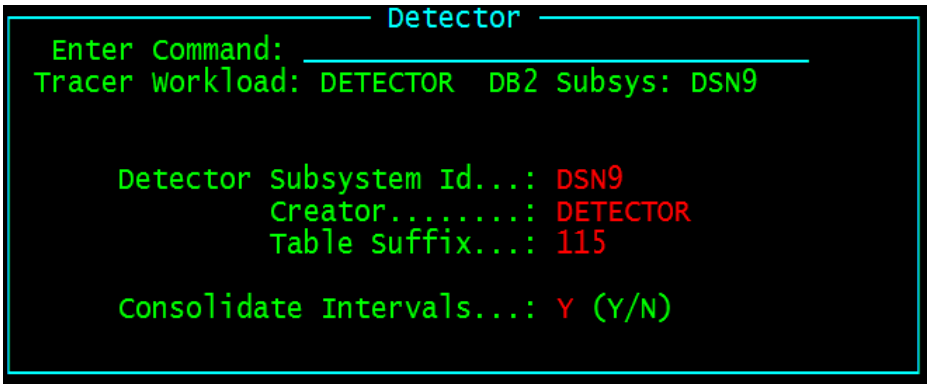

Figure 34 Load from CA Detector Pop-Up Window

The following information is displayed on the panel:-

Detector Subsystem ID — Specify the DB2 Subsystem ID where the CA Detector tables reside.

Creator — Specify the Creator used for the CA Detector® tables containing the externalized data, Refer to the CA Detector® documentation for information about externalizing the captured data.

Table Suffix — Specify an optional table suffix to be appended to the Detector Table name.

Consolidate Intervals — Perform SQL consolidate across Trace Intervals.
Y Consolidate Intervals
N Do not consolidate intervals

Note: In the Trace Parameters, you can enter all of the normal trace parameters, for example, FILTERS and SQL Consolidation Controls. Trace Time and Duration etc. are ignored. The DB2 Subs-system specified on the Start Trace panel should be the DB2 subsystem the SQL relates to. Refer to *the EZ-Tracer User Guide* for further information.

Once you have entered the data on this panel, continue to the normal tracer job submission and submit the job. The JCL contains the new step XOPLDET to load the CA Detector® data into the trace workload.

When the job has completed, you can view the Tracer workload using the EZ-Tracer reports menu in the normal way

You can also import the trace workload into an EZ-SQL Warehouse workload in the normal way, where it can be processed by EZ-Index Analyzer and /or EZ-XOP (Index Optimizer).