precisely

Syncsort Optimize DB2
(EZ-DB2)

Housekeeping and Controls User Guide
Release 9.10

Version 6.11

July 2014

Syncsort Optimize DB2 Housekeeping and Controls User Guide
(formerly EZ-DB2 Housekeeping and Controls 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.

ACKNOWLEDGEMENTS:

EZ-DB2 and the **EZ-DB2** family of products are registered trademarks of Cogito Limited.

Z/OS, OS/390, **DB2** and **IBM** are proprietary trademarks of International Business Machines. All other trademarks acknowledged.

EZ-XOP has been developed under a patent license from British Telecommunications plc. European Application No: 96944107.0 including its foreign counterparts.

TRSMAIN is a licensed product of IBM.

CONDITIONS:

EZ-DB2 is designed to be used by licensed users of DB2 as a supplement to this product.

PREFACE:

EZ-DB2 is a program product of Cogito consisting of EZ-Tracer, EZ-Alerts, EZ-Impact Analyzer, EZ-Index Analyzer and EZ-XOP Index Optimizer.

This document describes the EZ-DB2 Housekeeping and Controls functions which in many cases are common to the other EZ-DB2 Components.

RELATED DOCUMENTATION:

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

TABLE OF CONTENTS

FIGURES

New Features

Import Plan Table data from a remote system using DRDA

A new option has been added to allow the import of Plan table data from a remote system using DRDA.

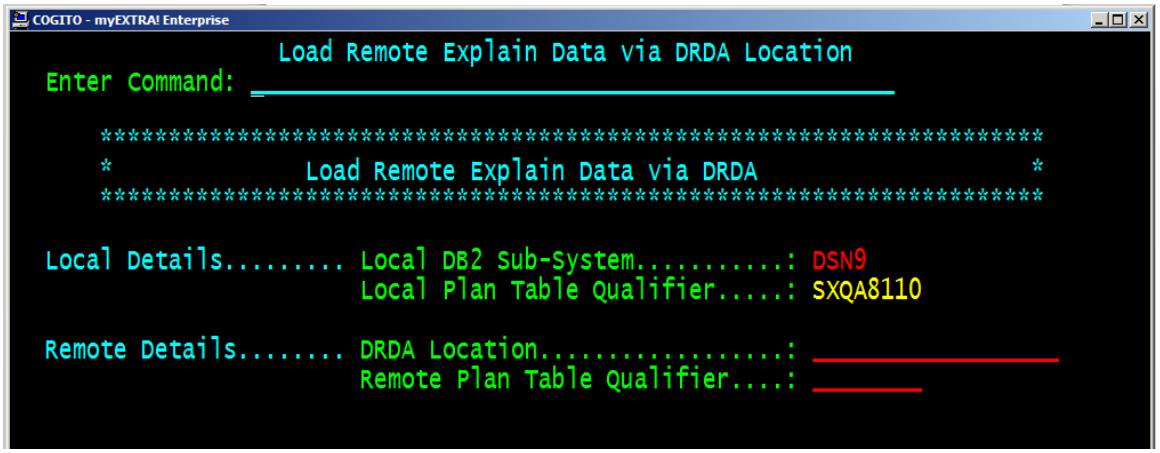

Refer to option 9.7 for further information.

Housekeeping and Controls

9 Housekeeping and Controls

Select option 9 to display the Housekeeping and Controls menu as shown in the following figure:-

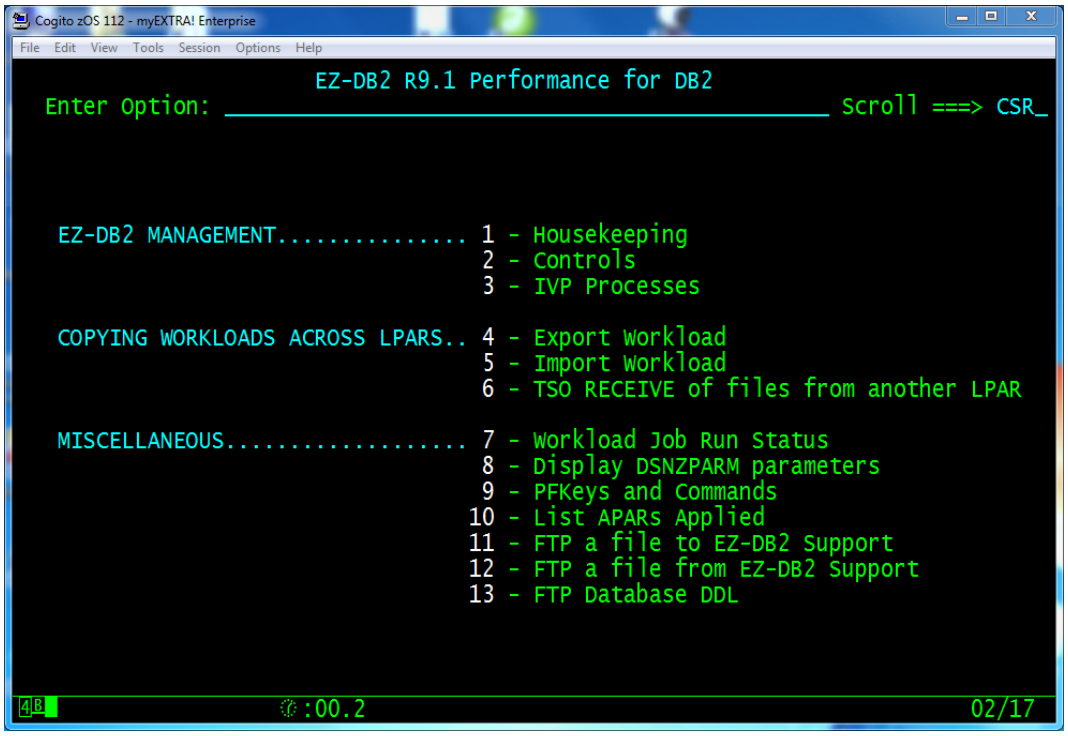

Figure 1 EZ-DB2 Housekeeping and Controls menu

EZ-DB2 Management

9.1 Housekeeping Functions

Housekeeping functions allow you to maintain the EZ-DB2 environment and databases.

From the House Keeping menu you can delete entire sets of workload data from EZ-DB2 tables and report files, generate RUNSTATS for both the EZ-DB2 and Live databases, setup and populate IVP databases, re-bind and re-link EZ-DB2 programs, add EZ-DB2 databases, and edit various EZ-DB2 control files.

Select option 9.1 to display the Housekeeping Menu as shown in the following diagram:-

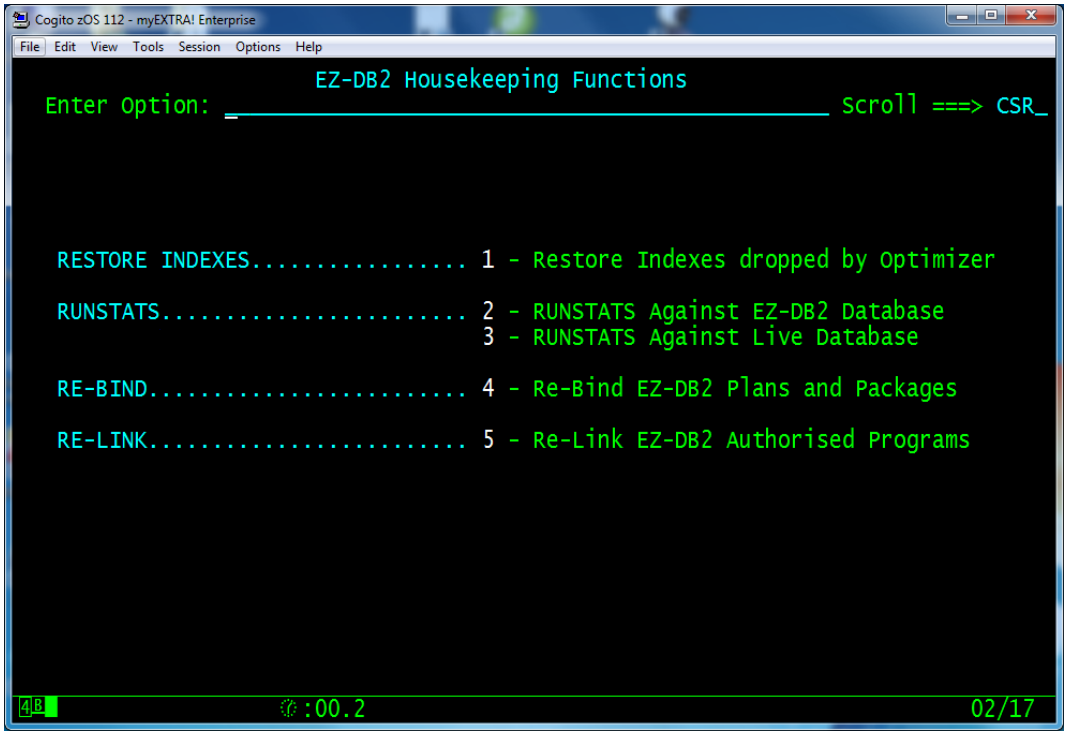

Figure 2 EZ-DB2 Housekeeping Functions

9.1.1 Restore User Indexes Dropped by Optimizer

Each run of the EZ-XOP Optimizer will drop all user indexes not required for the analysis and create the set of indexes it calculates to be the optimum for the Workload.

If the same Workload is processed again, then the Prepare for Index Optimization process (option 8.1) will drop the EZ-XOP indexes and restore all the original user indexes to put everything back to its original state.

However, if you wish to process any of the same tables using a *different* Workload, then you must first select this option to drop the EZ-XOP indexes and restore the User indexes.

Note that this option must be selected with Workload and Version set to the same values as when the EZ-XOP Optimizer was run and the user indexes dropped.

Select option 9.1.1 to generate the job stream to restore the User Indexes Dropped by EZ-XOP.

X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the Housekeeping menu.

9.1.2 Run RUNSTATS against EZ-DB2 Database

As more and more traces are run, EZ-DB2 performance will be improved by occasionally running RUNSTATS against the EZ-DB2 Warehouse database.

This option generates the JCL to run RUNSTATS against the EZ-DB2 Warehouse database.

Select the Option 9.1.2 to generate the job stream to run RUNSTATS against the EZ-DB2 database.

9.1.3 Run RUNSTATS against Live Database

This option allows you to run 'RUNSTATS COLUMN ALL' for all tables being analyzed by EZ-DB2.

RUNSTATS must be run against the Live databases to be analyzed prior to unloading catalog statistics and after implementing EZ-XOP recommended indexes.

Select Option 9.1.3 to display the following panel:-

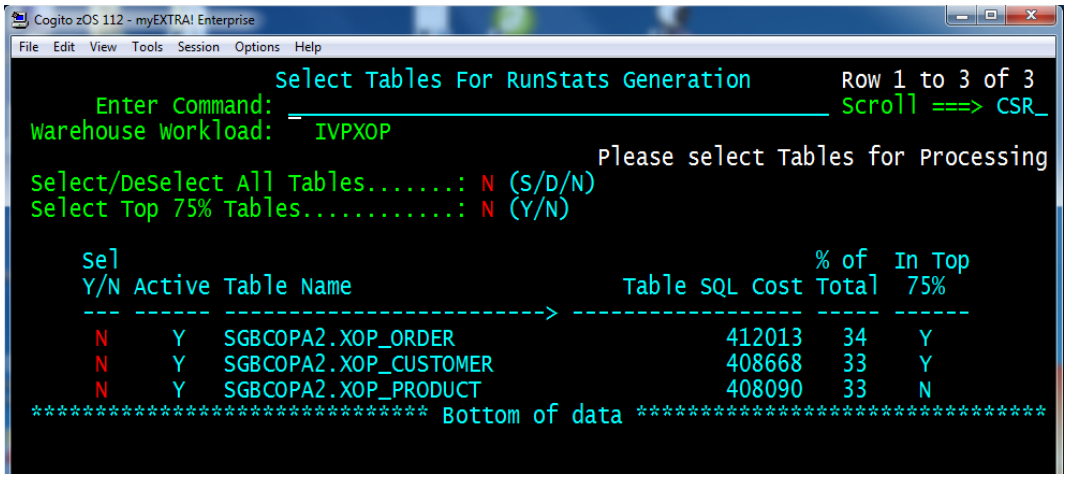

Figure 3 EZ-DB2 Select Tables for RUNSTATS Generation

The following parameters may be specified on this panel.

Select/DeSelect All Tables(S/D/N)

Set the Select Status to Yes or No for all Tables.

S - Set the Select indicator for all tables to Y (Selected)

D - Set the Select indicator for all tables to N (DeSelected)

N -No action

The default value is N.

Select Top 75% Table Names(Y/N)?

Specify Y if you wish that all the tables that comprise the top 75% of the Workload cost should be Selected.

Sel Y/N

Specify Y if wish to generate RUNSTATS for this table.

Table Name

The list of tables which have been referenced by one or more SQLs in the Workload sorted in descending order of Table SQL Cost.

Table SQL Cost

The total weighted Timeron Cost of the SQL in the Workload that reference the table. Where SQL reference more than one table its cost will be included against each table.

% of Total	The total cost attributed to this table represented as a percentage of the total cost of all tables for the current workload.
In Top 75%	Set to Y for all tables that are included in those that comprise 75 percent of the total Workload Cost.

Press <PF3> to generate the RUNSTATS job stream.

The names of the tables for which RUNSTATS parameters are required are determined from the contents of the XOPTBNAM file. However, the table names in the XOPTBNAM file at this time reflect the names as used on the EZ-DB2 analysis DB2 subsystem, which may be qualified differently to the Live DB2 subsystem. For this reason, there is a preliminary job step which invokes program XOPTABC. This program reads XOPDBNAM, XOPTBNAM and XOPQUALS, translates the table qualifiers back to the Live values where applicable, and writes the modified XOPTBNAM out to a temporary dataset.

X Type SUBMIT in the command line to submit the job.
X Press <**PF3**> to return to the Housekeeping menu.

9.1.4 Re-Bind EZ-DB2 Plans and Packages

Option 9.1.4 generates a job stream to bind all the EZ-DB2 supplied programs.

Use this option if you need to BIND the EZ-DB2 supplied programs.

You do not need to run the BIND job following supplied maintenance to EZ-DB2.
If BINDS are required to implement the maintenance, then these will have been incorporated into the APAR apply job itself. Refer to the EZ-DB2 Installation Reference Guide for more information about applying maintenance.

9.1.5 Re-Link EZ-DB2 Authorized Programs

EZ-DB2 requires that some of the programs are run authorized.

You specified the EZ-DB2 Authorized Loadlib DSN Qualifier as part of the install procedure.

Use option 9.1.5 to generate the job stream to manually re link the EZ-DB2 authorized programs.

This option should only be required if you need to recreate the authorized dataset. Any supplied maintenance that affects the authorized programs, will automatically re link the authorized Loadlib. Refer to the EZ-DB2 Installation Reference Guide for more information about applying maintenance

9.2 Controls

The EZ-DB2 Controls option, allows you to maintain the Environmental details that were originally set up as part of the install procedures. You can also maintain your Workload Profiles, and GRANT EZ-DB2 permissions to users.

Select Option 9.2 to display the following menu panel:-

Figure 4 EZ-DB2 Controls menu

9.2.1 Jobcard

Select Option 9.2.1 - Jobcard to edit the default job card as required to reflect your installation standards. This job card will be used in all subsequent installation job steps, and all runtime job steps when using the EZ-DB2 system.

Note that the values shown in lower case will be substituted at run time and should not be changed. The runtime job name is substituted with the default job name as specified here. *cntlmemb* is substituted with the JCL member name in the CNTL dataset. This makes it easier to identify jobs on the output queue.

The Job Card can be refreshed from the installation default by entering **Y** in Refresh Job Card from Default field.

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

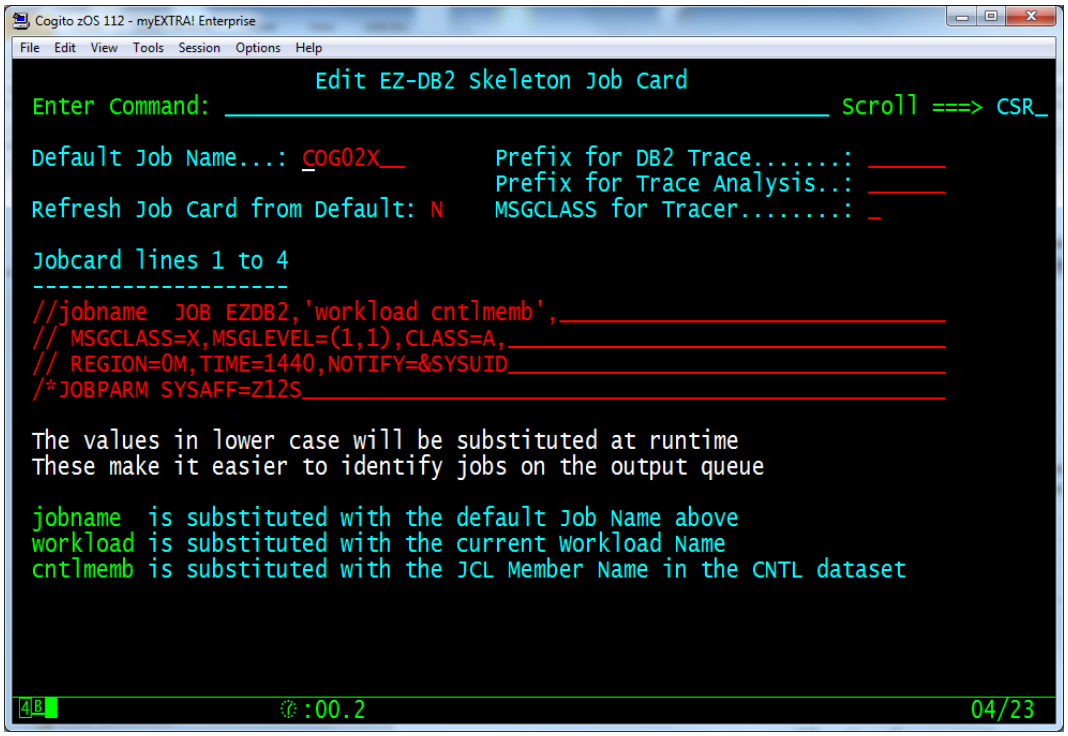

Figure 5 EZ-DB2 Skeleton Job Card

Press **<PF3>** to save your changes.

9.2.2 EZ-DB2 Configuration Parameters

Select Option 9.2.2 - EZ-DB2 Configuration Parameters to display the EZ-DB2 configuration parameters panel as shown in the following figure:-

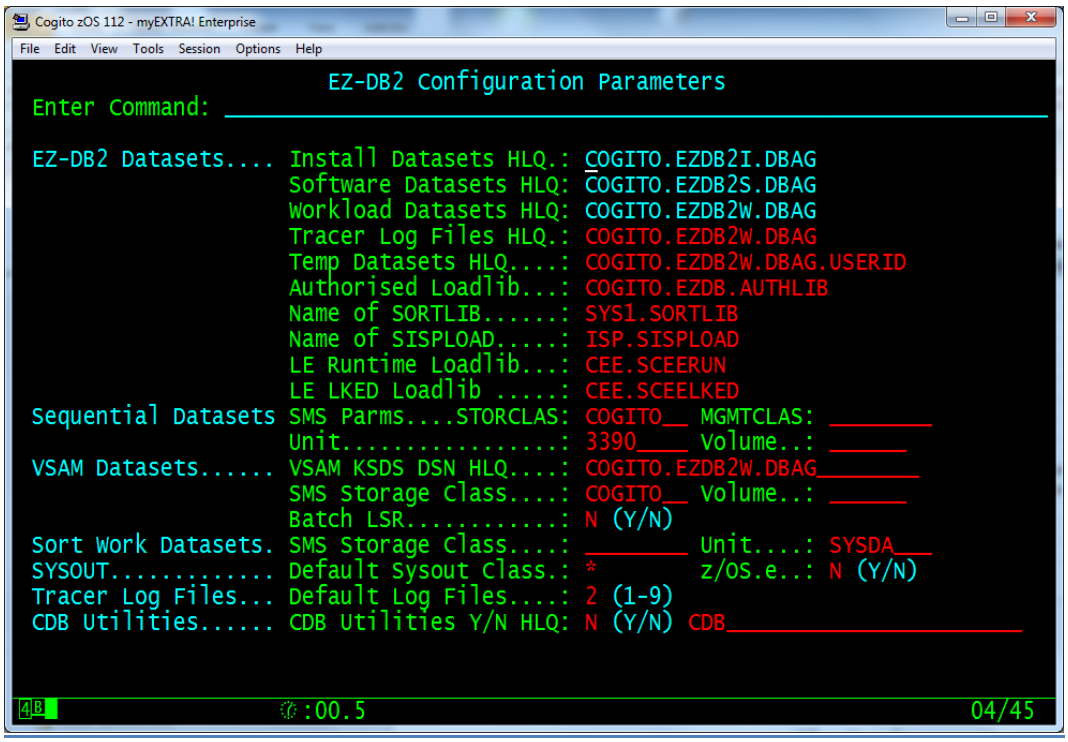

Figure 6 EZ-DB2 Configuration Parameters

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

After completing entry of the required parameters, press **<PF3>** key to save the parameters and return to the EZ-DB2 controls menu.

9.2.3 EZ-DB2 DB2 Subsystems

Select option 9.2.3 to display the EZ-DB2 DB2 Subsystems menu as shown in the following figure:-

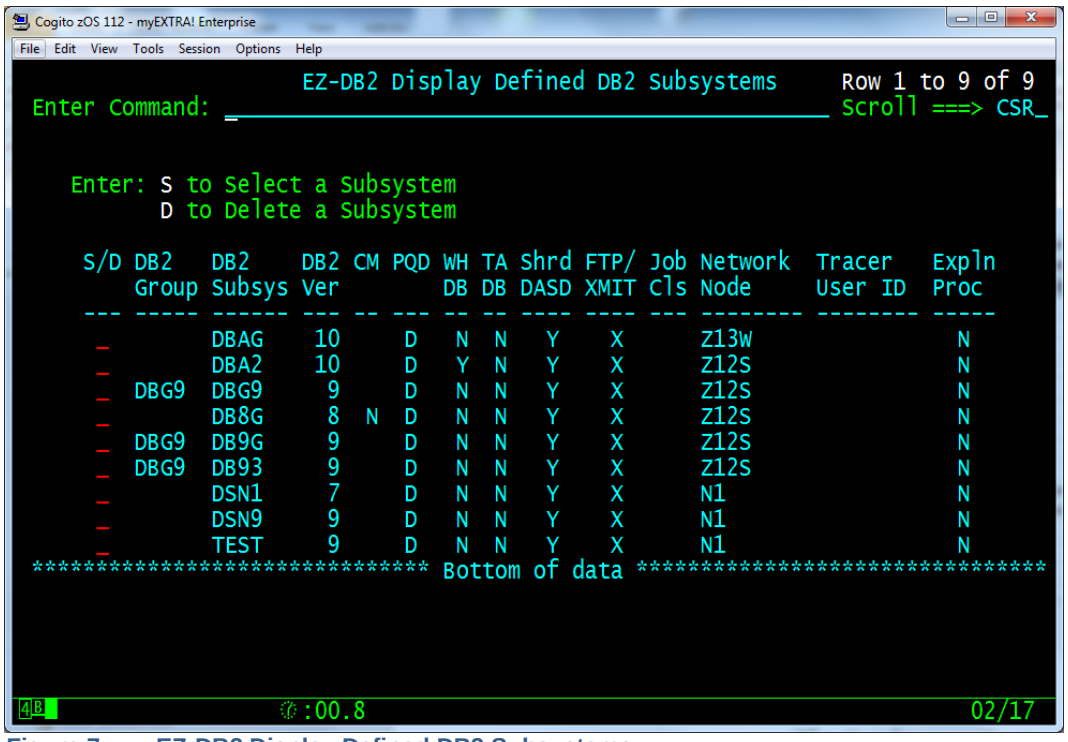

Figure 7 EZ-DB2 Display Defined DB2 Subsystems

This panel displays all DB2 Subsystems currently defined to EZ-DB2 and their main attributes.

You may select a DB2 Subsystem name to display the subsystem parameters or add new subsystem definitions.

The following fields are displayed on this panel:-

S/D	Enter S or Y to display the DB2 Subsystem Parameters panel for the subsystem as shown below. Enter D to delete the subsystem
DB2 Group	The name of the DB2 Group if the subsystem is a member of a data sharing group.
DB2 Subsys	The name of the DB2 Subsystem
DB2 Ver	The DB2 Version
CM	Whether this DB2 Subsystem is defined to EZ-DB2 as running in compatibility Mode. DB2 V8 Only.

PQD	Whether this subsystem is Production (P), QA (Q) or Development (D).
WH DB	Y indicates that an EZ-DB2 Warehouse is installed on this Subsystem.
TB DB	Y indicates that a Trend Analysis Database is installed on this Subsystem.
Shared DASD	Y if the Subsystem shares DASD with the EZ-DB2 Analysis Subsystem.
FTP/XMIT	F if FTP is the preferred method of file transfer to and from the EZ-DB2 Analysis Subsystem. X if XMIT is the preferred method of file transfer to and from the EZ-DB2 Analysis Subsystem.
Job Cls	The default job class to be used.
Network Node	The Network Node of the LPAR on which this DB2 subsystem resides.
Tracer Userid	RACF user ID to be used for Tracer.
Expln Proc	Y indicates EZ-DB2 will use its Stored Procedure when processing explains.

The following panel is displayed when you select a particular DB2 Subsystem allowing you to edit any parameters as required.

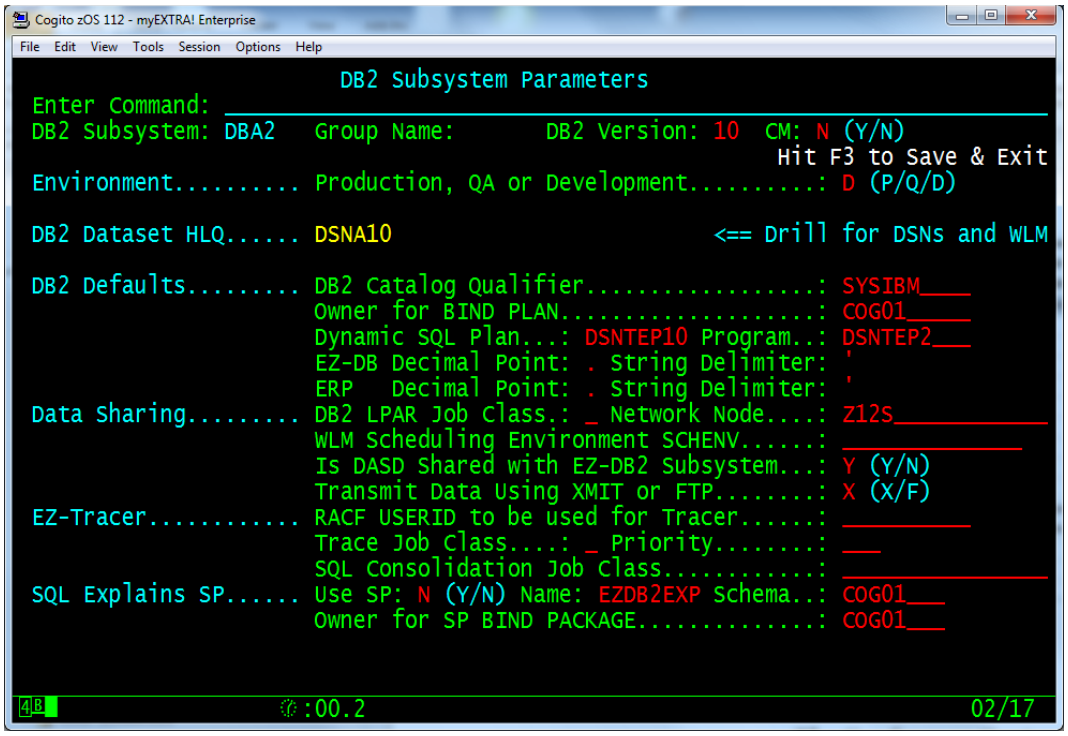

Figure 8 DB2 Subsystem Parameters

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

9.2.4 EZ-Impact Analyzer Workloads

Select Option 9.2.4 - EZ-Impact Analyzer Workloads to display and maintain EZ-Impact Analyzer Workloads.

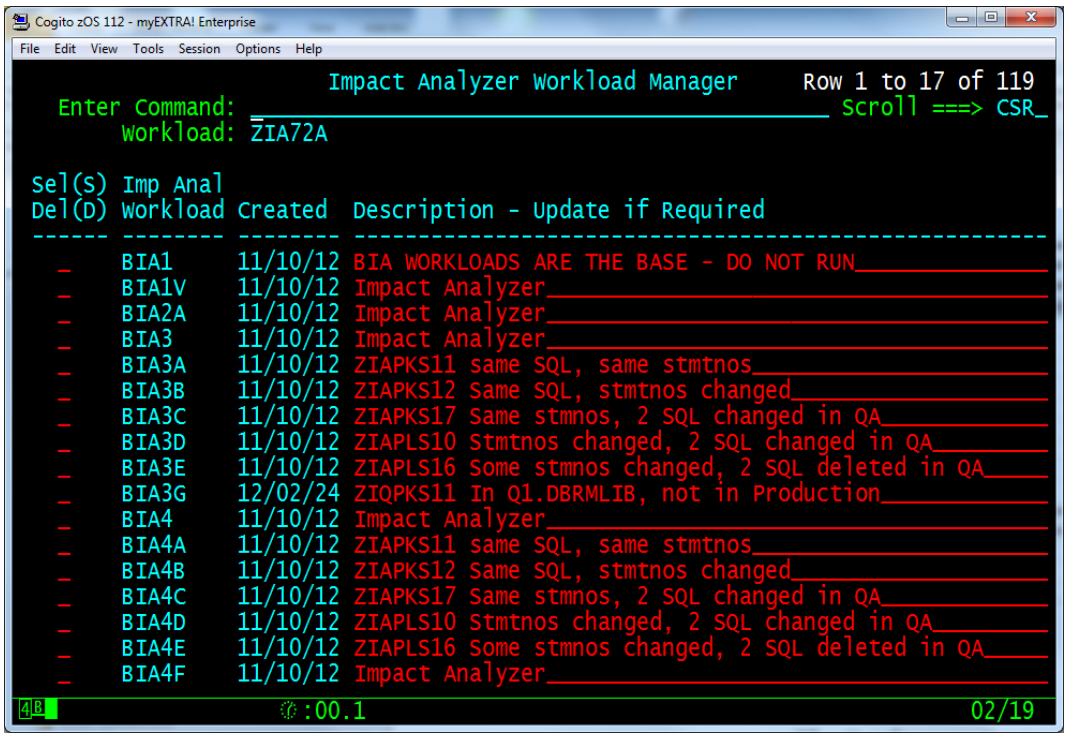

Figure 9 EZ-Impact Analyzer Workloads

The following parameters are displayed on this panel:-

Sel(S) **Del(D)**	Enter S to select this workload as the current Impact Analyzer Workload. Enter D to delete this workload
Imp Anal Workload	The Workload name.
Created	The date this workload was created.
Description	A description of this workload. You may update this description if required.

See Impact Analyzer User Guide for full details about Impact Analyzer Workloads.

9.2.5 EZ-Warehouse Workloads

Select Option 9.2.5 - EZ-Warehouse Workloads to display and maintain EZ-Warehouse Workloads.

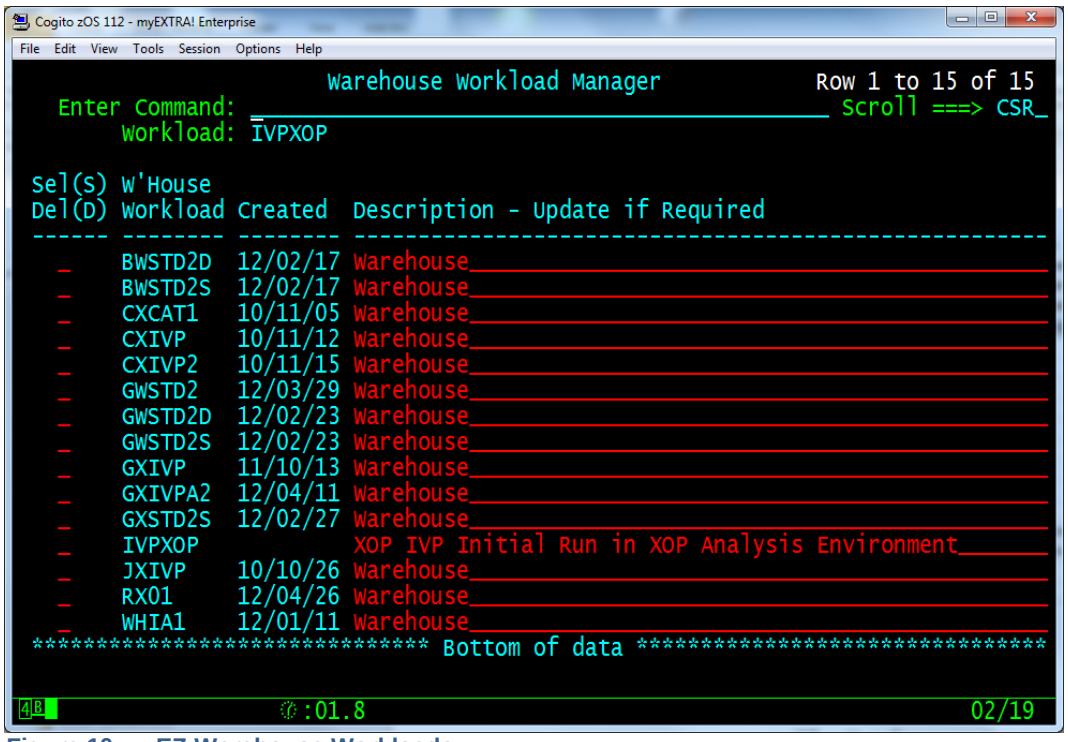

Figure 10 EZ-Warehouse Workloads

The following parameters are displayed on this panel:-

Sel(S) **Del(D)**	Enter S to select this workload as the current EZ-DB2 Warehouse Workload. Enter D to delete this workload
W=House Workload	The Workload name.
Created	The date this workload was created.
Description	A description of this workload. You may update this description if required.

See EZ-DB2 SQL Warehouse User Guide for full details about EZ-DB2 Warehouse Workloads.

9.2.6 EZ-Tracer Workloads

Select Option 9.2.6 - EZ-Tracer Workloads to display and maintain EZ-Tracer Workloads.

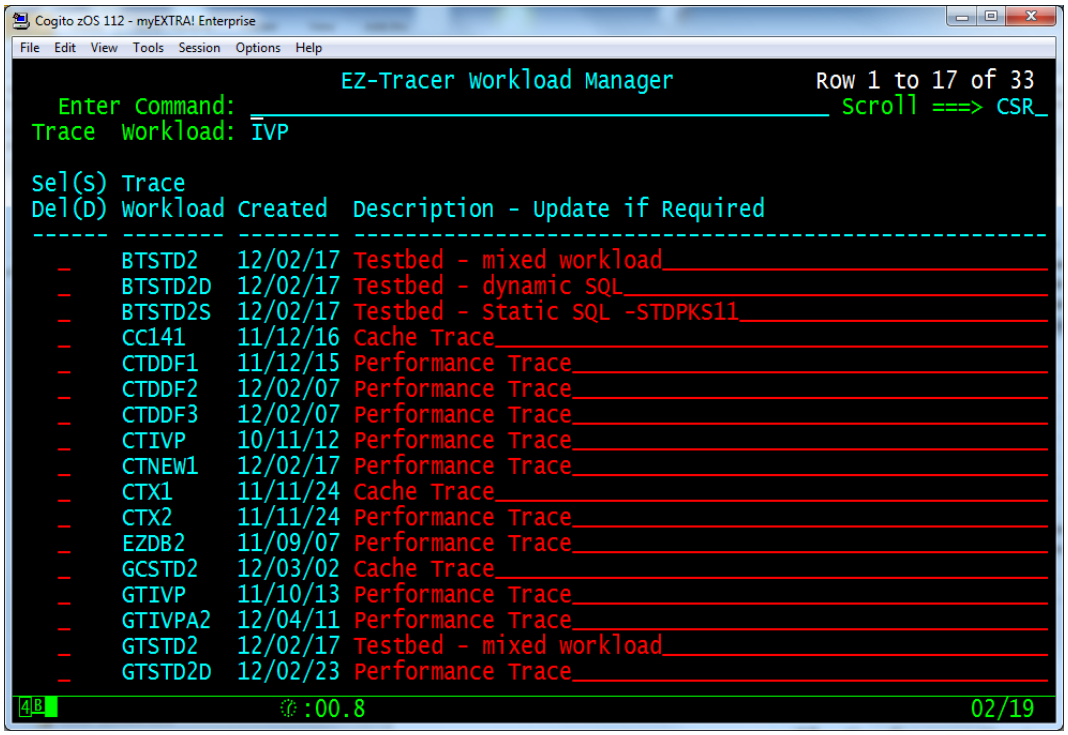

Figure 11 EZ-Tracer Workloads

The following parameters are displayed on this panel:-

Sel(S) **Del(D)**	Enter S to select this workload as the current EZ-Tracer Workload. Enter D to delete this workload
Trace Workload	The Workload name.
Created	The date this workload was created.
Description	A description of this workload. You may update this description if required.

See EZ-Tracer User Guide for full details about EZ-Tracer Workloads.

9.2.7 Trace Descriptions

Select option 9.2.7 to display the EZ-DB2 Trace Descriptions panel as shown in the following figure:-

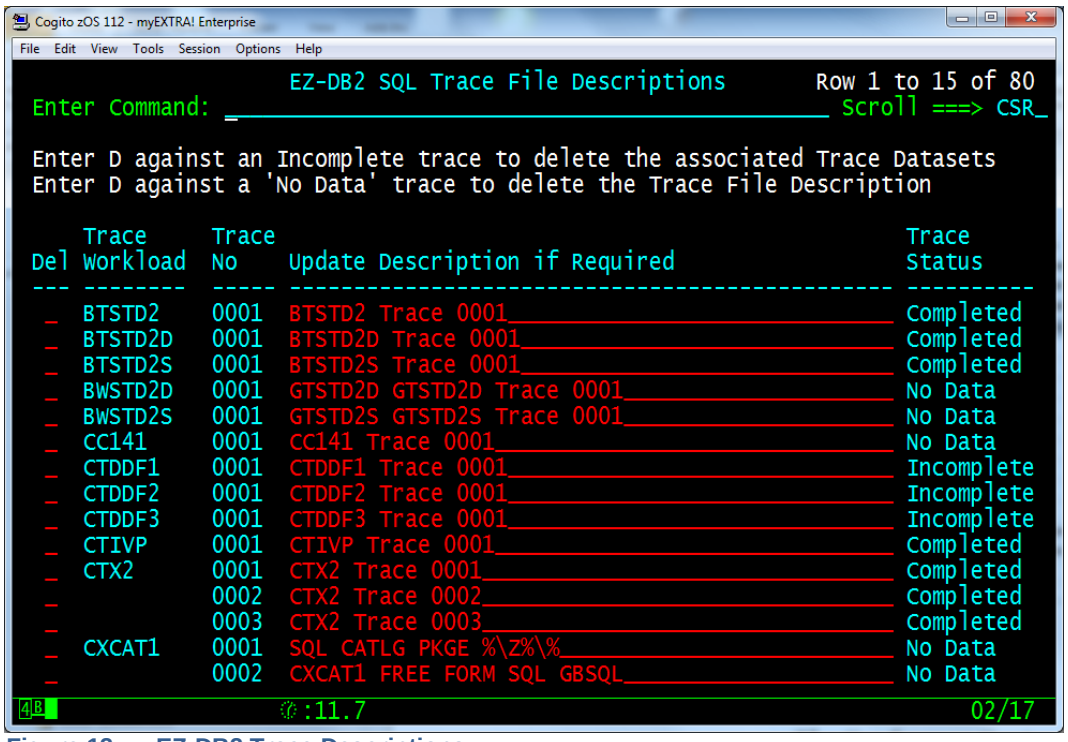

Figure 12 EZ-DB2 Trace Descriptions

The following parameters are displayed on this panel:-

Del	Enter D to delete this trace from the workload.
Trace Workload	The Workload name.
Update Description if Required	A description of this Trace. You may update this description if required.
Trace Status	The status of the trace.

Completed	The Trace completed successfully
Incomplete	The Trace did not complete successfully but there is some data.
No Data	The Trace did not complete successfully and there is no data.

See EZ-Tracer User Guide for full details about EZ-Tracer Workloads and traces.

9.2.8 GRANT EZ-DB2 Trace Permissions

The Grant EZ-DB2 Trace Permissions procedure is used to GRANT the necessary privileges including MONITOR for Users to be able to run DB2 Traces and the necessary EZ-DB2 Trace Programs on the DB2 Subsystem on which the trace will be executed.

This procedure must be submitted by a user with the appropriate authority with the applicable EZ-DB2 users' AUTHID as the recipient.

Select option 9.2.8 - GRANT EZ-DB2 Trace Permissions. The following panel is displayed:-

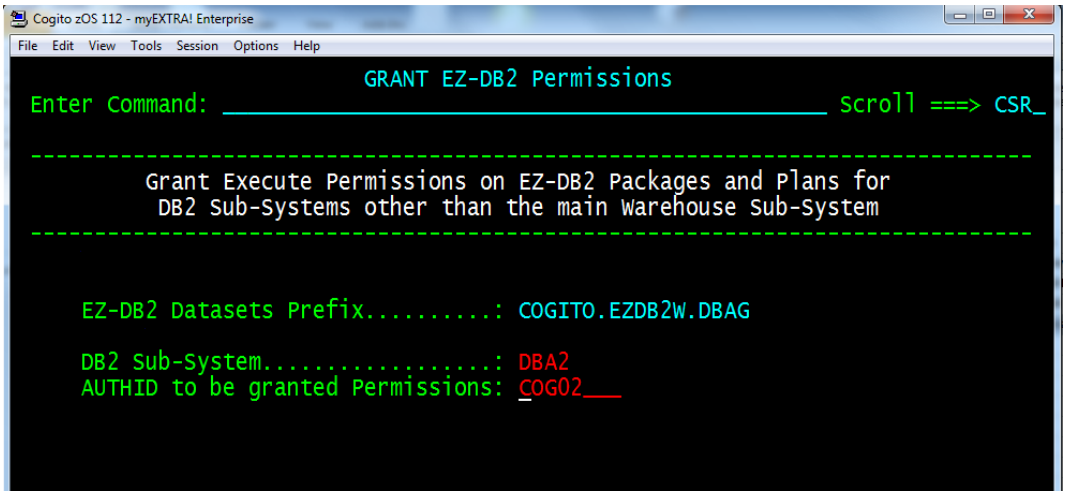

Figure 13 EZ-DB2 Grant Trace Permissions

The following parameters must be specified on this panel.

DB2 Subsystem	Specify the DB2 subsystem Id.
Authid to be Granted Permission	Specify the TSO User id corresponding to the User that will run the EZ-DB2 Trace facility.

When you have completed entry of the required parameters press **<ENTER>** to view the generated job stream.

If you are satisfied with the generated job stream, submit the job for execution.

Note: This procedure may be repeated for additional Authids as required. If you do not wish to authorize multiple users, you may specify a Tracer Userid for each applicable DB2 subsystem, which is a RACF userid to be substituted as the current userid when submitting the trace.

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

9.2.9 GRANT EZ-DB2 Analysis Permissions

The Grant EZ-DB2 Analysis Permissions procedure is used to GRANT the necessary privileges for Users to be able to run the EZ-DB2 Analysis process on the EZ-DB2 analysis DB2 subsystem. This procedure must be submitted by a user with the appropriate authority with one of the EZ-DB2 users' AUTHIDs as the recipient.

Select option 9.2.9 - GRANT EZ-DB2 Analysis Permissions. The following panel is displayed:-

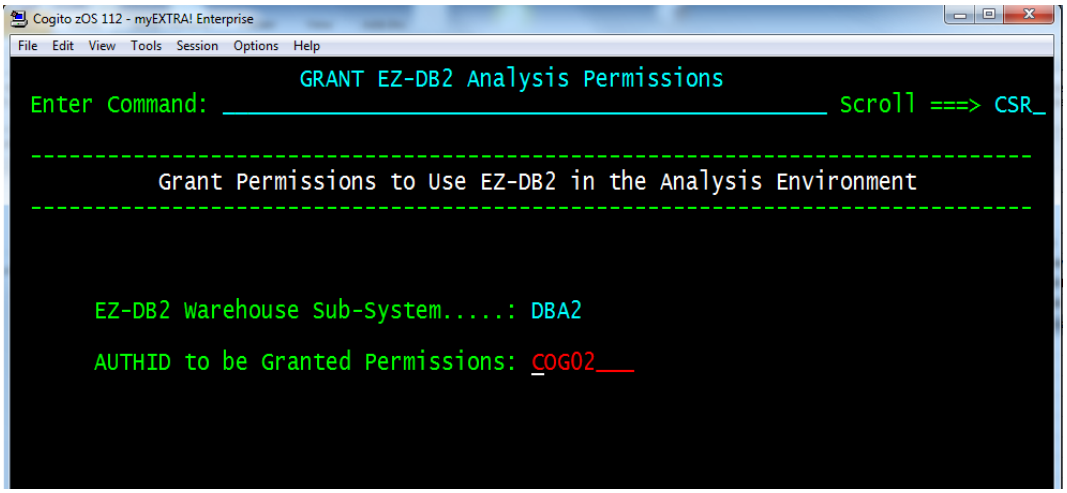

Figure 14 EZ-DB2 Grant Analysis Permissions

The following parameters must be specified on this panel.

EZ-DB2 Warehouse Sub-System	The DB2 subsystem ID of the EZ-DB2 analysis system.
AUTHID to be Granted Permissions	Specify the AUTHID corresponding to the User that will run EZ-DB2 on the 'EZ-DB2 Analysis' DB2 subsystem.

When you have completed entry of the required parameters press **<ENTER>** to view the generated job stream.

If you are satisfied with the generated job stream, submit the job for execution.

Note: This procedure may be repeated for additional AUTHIDs as required.

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

9.2.10 GRANT EZ-DB2 Create Permissions

The Grant EZ-DB2 Create Permissions procedure is used to GRANT the necessary privileges for Users to be able to create the EZ-DB2 Database on the EZ-DB2 Analysis DB2 subsystem (e.g. Grant Use of Stogroups, Buffer Pools, Catalog Tables).

This procedure must be submitted by a user with the appropriate authority with one of the EZ-DB2 users' AUTHIDs as the recipient.

Select option 9.2.10 - GRANT EZ-DB2 Create Permissions. The following panel is displayed:-

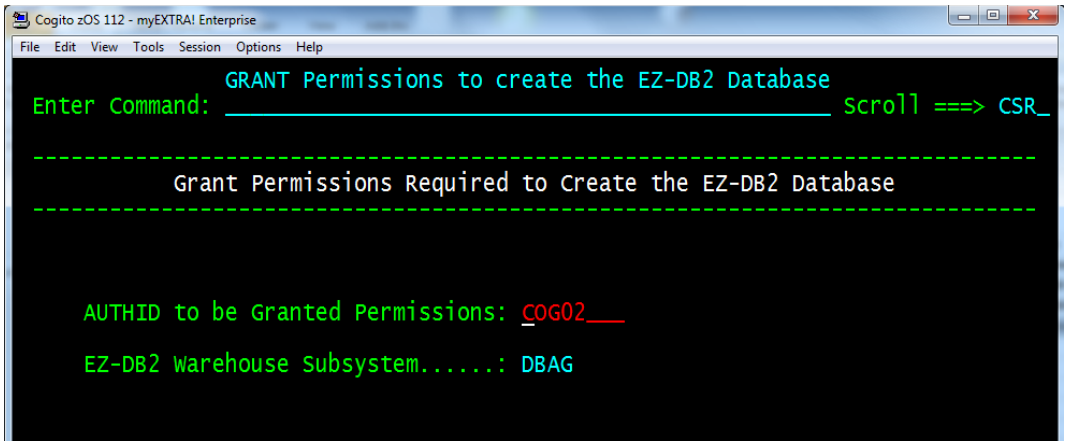

Figure 15 EZ-DB2 Grant Create Permissions

The following parameters must be specified on this panel.

AUTHID to be Granted Permissions	Specify the AUTHID corresponding to the User that will create the EZ-DB2 database on the EZ-DB2 analysis DB2 subsystem.
EZ-DB2 Warehouse Subsystem	The DB2 subsystem ID of the EZ-DB2 analysis system.

When you have completed entry of the required parameters press **<ENTER>** to view the generated job stream. If you are satisfied with the generated job stream, submit the job for execution.

Note: This procedure may be repeated for additional AUTHIDs as required.

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

9.2.11 Create new Explain Plan Table

A Plan table for EZ-DB2 Explains may be required on DB2 Subsystems on which you will be using Tracer to Generate the Access Path data for a Tracer Workload, or when installing EZ-DB2 for Impact Analyzer on another DB2 subsystem.

In these cases, you can use an existing PLAN TABLE, or use this option to generate a new PLAN TABLE. If you use an existing table it must be fully defined for the current release of DB2.

When using this option, EZ-DB2 will generate the DDL for the appropriate DB2 release and will include all of the required extended plan table definitions, i.e. DSN_STATEMNT_TABLE, DSN_DETCOST_TABLE, DSN_STRUCT_TABLE and DSN_QUERY_TABLE.

Select option 9.2.11 - to display the Create new Explain Plan Table as shown in the following figure:-

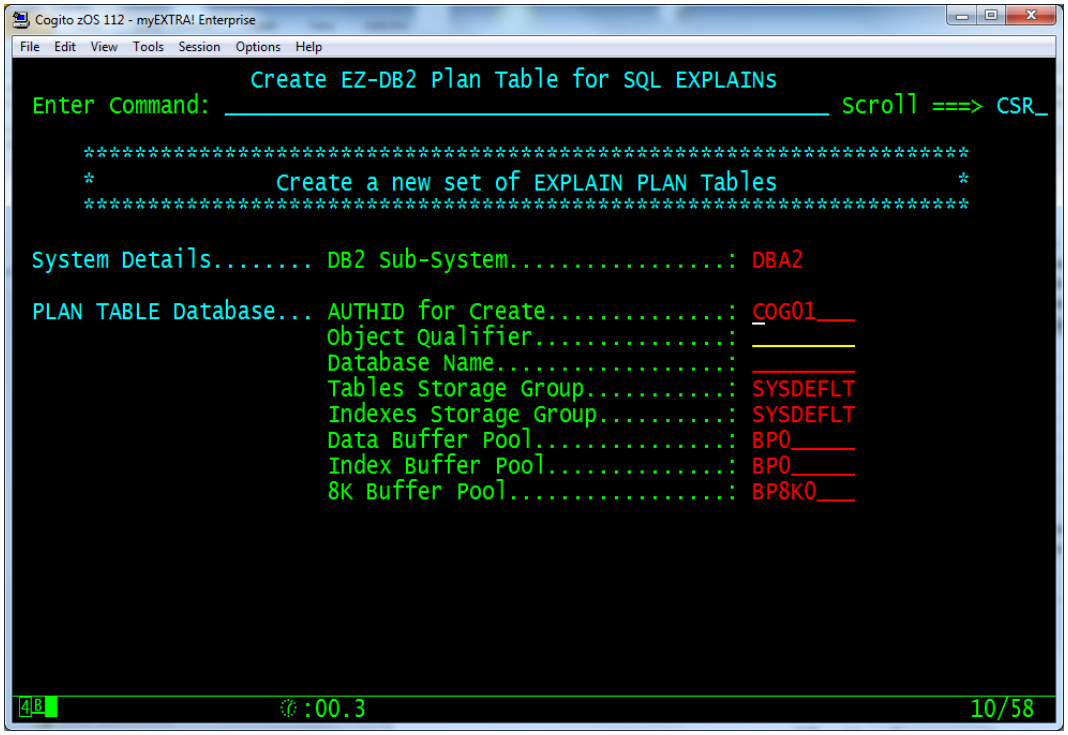

Figure 16 Create Plan Table for SQL Explains

The following fields are displayed on this panel.

System Details

DB2 Sub-System

The DB2 Subsystem on which the PLAN table is to be created.

Defaults to the current DB2 system on which the EZ-DB2 software is being installed.

Plan Table Database

Authid for Create	This is the SQLID to be used to Create the Plan Table.
Database Name	This is the DBNAME that will contain the PLAN_TABLE. This DBNAME will be created in the generated DDL.
Tables Storage Group	Specify the name of the Storage Group to be used for EZ-DB2 tables. By default EZ-DB2 assumes the storage group already exists. However, if you enter Catalog and Volume(s) for New Storage Groups, then EZ-DB2 will include the DDL to create the storage groups.
Indexes Storage Group	Specify the name of the Storage Group to be used for the EZ-DB2 Indexes. By default EZ-DB2 assumes the storage group already exists. However, if you enter Catalog and Volume(s) for New Storage Groups, then EZ-DB2 will include the DDL to create the storage groups.
Data Buffer Pool	Set this to the DB2 data buffer pool to be used by EZ-DB2. The default is BP1.
Index Buffer Pool	Set this to the DB2 index buffer pool to be used by EZ-DB2. The default is BP0.
8K Buffer Pool	Set this to the DB2 8K buffer pool to be used by EZ-DB2. The default is BP8K0.

When you have completed entry of the required parameters press **<enter>** to generate the job stream to create the Plan Table for Impact Analyzer.

Type SUBMIT in the command line and then <**PF3**> to return to the Activate EZ-DB2 on DB2 Subsystems menu.

9.3 IVP Processes

The Install Verification Process (IVP) is a sample database and workload supplied with EZ-DB2.

It enables the user to quickly run through the EZ-DB2 process from start to finish without having to find a suitable database just to try it out. Full details of installing and running the IVP can be found in the EZ-DB2 Installation Reference Guide and the EZ-DB2 Getting Started and IVP Guide.

9.3.1 Run IVP Workload to Generate SQL Trace Data

Select option 9.3.1 - Run IVP Workload To Generate SQL Trace Data to run the supplied IVP workload against the IVP Live Database.

This job will execute a predefined workload against the supplied IVP demo database. This job should execute in parallel with the previously submitted Workload Trace. You should check that the Trace job and sample workload jobs completed successfully, and that the trace job completed AFTER the sample workload job had completed. The default TRACE DURATION is only two minutes so, for example, if there was a delay before you started the workload after starting the trace, or there were other job scheduling delays, the trace may have terminated before completion of the workload. If this is the case, you should run both jobs again, perhaps specifying a longer TRACE DURATION.

The following panel is displayed:-

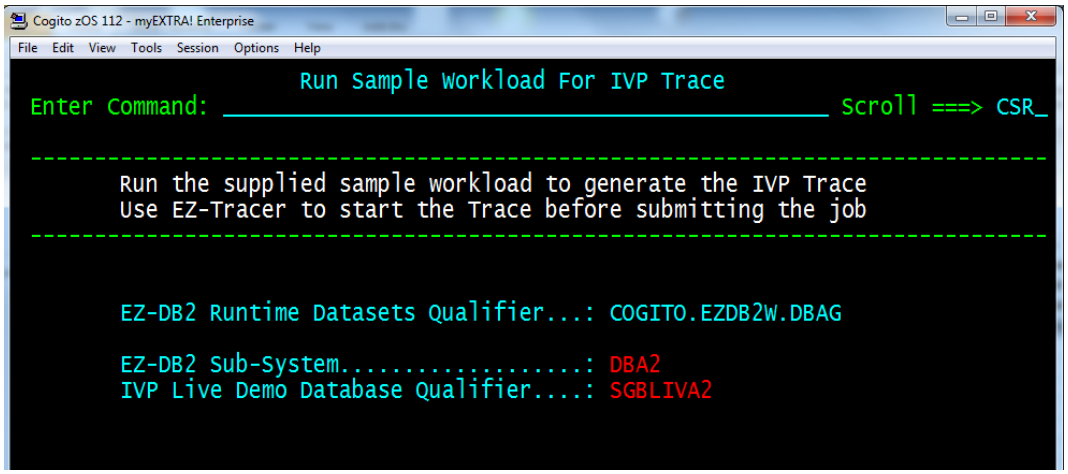

Figure 17 EZ-DB2 Run Sample Workload for IVP Trace

The following parameters may be specified on this panel.

IVP Live Demo Database Qualifier	Specify the qualifier that was used to create the IVP Live demo database. Refer to the EZ-DB2 Installation Reference Guide for further information.

Press **<enter>** to display the generated job stream for the sample workload job.

X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the Housekeeping menu.

9.3.2 Recreate IVP Live Database

The Recreate IVP Live Database option allows you to recreate the 'Live' copy of the supplied IVP Database. This is the IVP Database that you run the IVP Trace against.

Use this option if you wish to recreate the IVP Live Database after the IVP has been run as part of the EZ-DB2 Installation Process.

This step creates the database, populates it with some test data and runs Runstats Column(ALL) and Index(ALL) so that there are accurate production statistics available to the EZ-DB2 Workload Optimization process.

Select option 9.3.2 to display the Recreate IVP Live Database panel shown in the following figure:-

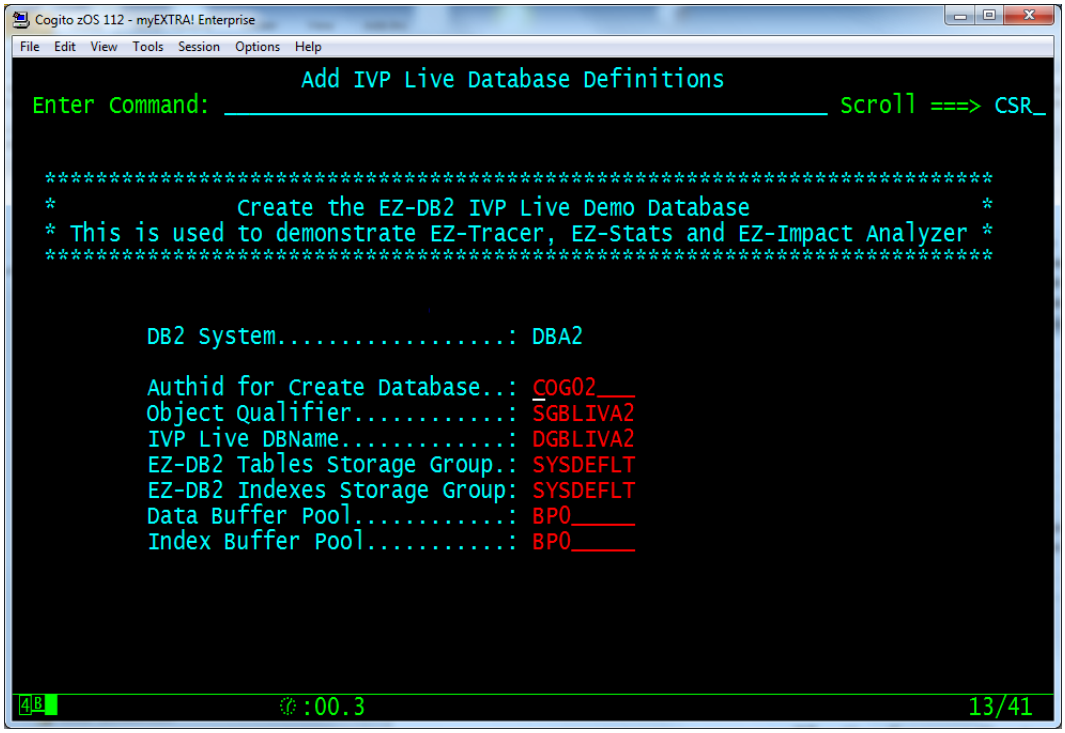

Figure 18 EZ-DB2 Recreate IVP Live Database

The following fields are displayed on this panel.

DB2 System

The SSID of the EZ-DB2 Analysis subsystem.

Defaults to the value specified for EZDBSSID.

Authid for Create Database

The Authid used when creating the IVP Live database.

This value will default to the Authid used to create the EZ-DB2 warehouse database. See EZ-DB2 Database Definitions parameter Authid to Create Database.

Object Qualifier	Specify the qualifier to be used to create the IVP Live demo database. This qualifier will be used as qualifier for the EZ-DB2 IVP Live demo database tables. This value will default to the Authid used to create the EZ-DB2 database. See EZ-DB2 Database Definitions parameter Authid to Create Database.
IVP Live DBName	Set this to the database name to be used when creating the IVP Live database.
EZ-DB2 Tables Storage Group	Specify the name of the Storage Group to be used for EZ-DB2 tables. Defaults to the value specified in EZ-DB2 Database Definitions.
EZ-DB2 Indexes Storage Group	Specify the name of the Storage Group to be used for the EZ-DB2 Indexes. Defaults to the value specified in EZ-DB2 Database Definitions.
Data Buffer Pool	Set this to the DB2 data buffer pool to be used by EZ-DB2. The default is BP1.
Index Buffer Pool	Set this to the DB2 index buffer pool to be used by EZ-DB2. The default is BP0.

When you have completed entry of the required parameters press **<enter>** to generate the job stream to create the IVP Database.

Type SUBMIT in the command line and then <**PF3**> to return to the IVP Process menu.

9.3.3 Recreate IVP Warehouse Database

The Recreate IVP Warehouse Database option allows you to recreate the 'Clone' copy of the supplied IVP Database. This is the IVP Database that you run the EZ-DB2 Analysis and Optimization against.

Use this option if you wish to recreate the IVP Clone Database after the IVP has been run as part of the EZ-DB2 Installation Process.

The IVP requires a Clone of the IVP Live Database to be defined to be used by the EZ-DB2 Workload Optimizer to evolve an Optimum set of indexes. This step just creates the database but does not load any data into it.

Select option 9.3.3 to display the Recreate IVP Warehouse Database panel shown in the following figure:-

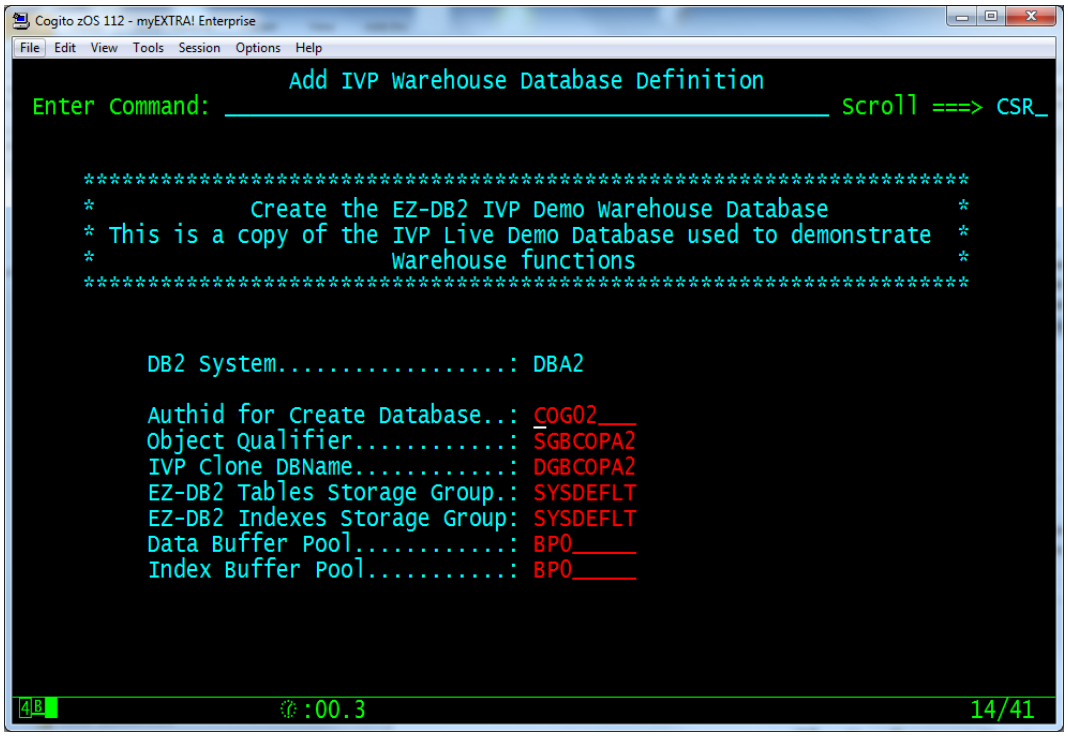

Figure 19 EZ-DB2 Recreate IVP Warehouse Database

The following parameters may be specified on this panel.

DB2 System

The SSID of the EZ-DB2 Analysis subsystem.

Defaults to the value specified for EZDBSSID.

Authid for Create Database

The Authid to be used when creating the IVP Clone database.

Note: This should be different to the Authid used to create the IVP Live database.

Object Qualifier	Specify the qualifier to be used to create the IVP Clone database. This qualifier will be used as qualifier for the EZ-DB2 IVP Clone database tables.
IVP Clone DBName	Set this to the database name to be used when creating the IVP (Installation Verification Process) Clone database
EZ-DB2 Tables Storage Group	Specify the name of the Storage Group to be used for IVP Clone Database tables.
EZ-DB2 Indexes Storage Group	Specify the name of the Storage Group to be used for IVP Clone Database Indexes.
Data Buffer Pool	Set this to the DB2 data buffer pool to be used by EZ-DB2. The default is BP1.
Index Buffer Pool	Set this to the DB2 index buffer pool to be used by EZ-DB2. The default is BP0.

X Press **<enter>** to display the generated job stream to create the IVP Clone database.
X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the IVP Process menu.

9.3.4 Add Indexes to Live Database

The Add Indexes to Live Database function is used by the IVP process. Its purpose is to add some additional indexes to the IVP Live database to demonstrate the functionality of EZ-Impact Analyzer.

Refer to the Getting Started and IVP Guide for further information.

Select option 9.3.4 to generate the job stream to add the Impact Analyzer Indexes. The job stream contains the following steps:-

X DDLIVNIX - add indexes to IVP Live database
X RUNSTATS - Run RUNSTATS to refresh Catalog Statistics for all tables in the IVP Live Database

Note the SQLID, table qualifiers etc default to the values used when the IVP Live database was originally created.

Note that the Database name, Authid and Object Qualifiers default to the values specified when the IVP Live database was originally created. Refer to option 9.3.2 for further details.

X Type Submit in the command line to submit the job.
X Enter <**PF3**> to return to return to the IVP Process menu.

9.3.5 Build XOP Defined Indexes on IVP Live Database

This option allows you to apply the index changes recommended by the Index Optimization process, to the IVP Live Database.

Select option 9.3.5 to display the following panel:-

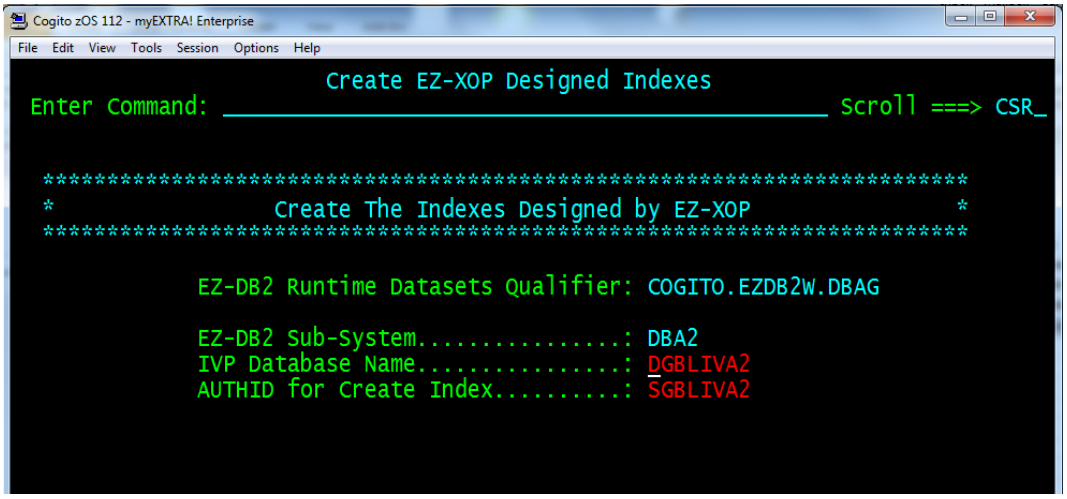

Figure 20 Create EZ-XOP Designed Indexes on the IVP database

The following parameters may be specified on this panel.

IVP Database Name Specify the DBNAME of the IVP 'Live' Database. This is the database that the index changes will be applied to.

Authid for Create Index Specify the Authid used as the qualifier for creating the 'Live' IVP Database. This is the Qualifier for the index changes to be applied.

X Press **<enter>** To display the generated job stream to apply the index changes..
X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the IVP Process menu.

Copying Workloads across LPARS

9.4 Export Workload

Use this panel to compress all data for a given Workload into a single Binary file to be sent to an LPAR that does not share DASD. The data may be sent by FTP, TSO XMIT or Connect:Direct.

If you wish to use Connect:Direct you must create a USRWKSNY member in the ...CNTL file by tailoring the EZ-DB2 supplied member XOPWKSNY.

If requested by Technical Support it may also be FTP'd to the support website. Specify SUPPORT as the FTP Node to generate JCL to FTP the compressed workload file directly to Technical Support.

If Technical Support has requested a Diagnostic Trace then enter Y against "Include Trace Diagnostics" to include an additional Trace diagnostic file. You can also elect to remove the SQL text from the Workload sent to support (for security), by specifying N against "Include SQL Text". These additional parameters are only displayed if you specify SUPPORT.

Enter the Workload Name and FTP, XMIT or Connect:Direct details.

If you wish to use the Statistics held on this LPAR, when explaining SQL in EZ-DB2 on the remote LPAR, then the statistics must be migrated separately using the EZ-Stats menu options.

Select option 9.4 to display the Export Workload panel as shown in the following figure:-

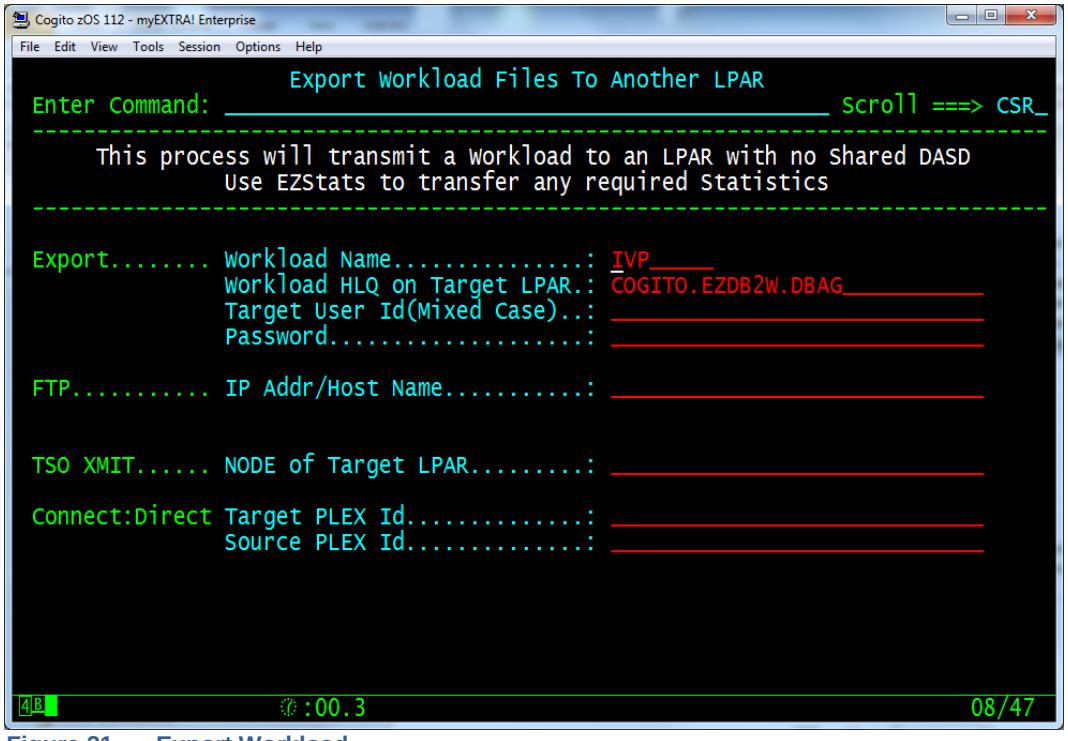

Figure 21 Export Workload

The following parameters may be specified on this panel.

Export

Workload Name
Set this field to the name of the Workload to be packaged and sent to the remote LPAR.

Workload HLQ on Target LPAR
Set this field to the high level qualifier for the EZ-DB2 Workload files on the remote LPAR.

This is used for FTP to 'PUT' the packaged workload file into 'tgtprfx.WORKLOAD.EXPORT' or to specify the &FILEOUT1 parameter for Connect:Direct.

Target User ID (Mixed Case)
Set this field to the username to sign onto the remote FTP server or to the Target User Id for the TSO XMIT.

For mainframe FTP, this is normally the remote TSO userid.

Password
Set this field to the password to use for signing onto the remote FTP server.

For mainframe FTP, this is normally the TSO Userid's password.

Note: This password is NOT stored within EZ-DB2.

FTP To

IP Addr/Host Name
Set this field to the IP address or Host Name of the remote LPAR. This field is only required when using FTP for the transfer.

If this field is set to SUPPORT (all other FTP fields can be left blank), then the generated JCL will automatically transfer the packaged workload to Technical Support.

Include Trace Diagnostics (Y/N)
If you are sending the workload to Technical Support, and have been requested to create a 'Diagnostic Trace'. Then specify 'Y' to include the diagnostic COTRAP file in the packaged workload.

This parameter is only displayed if you specify SUPPORT as the Host Name.

Include SQL text (Y/N)
For Security purposes you may wish to remove the SQL text from the diagnostic data sent to Technical Support.

Specify N to remove SQL Text.

This parameter is only displayed if you specify SUPPORT as the Host Name

TSO XMIT

NODE of Target LPAR — Set this field to the JESS Node of the remote LPAR, if you wish to use XMIT to send the packaged workload to the remote LPAR.

Connect:Direct

Target PLEX Id — Set this to the Id of the Target PLEX if you are using Connect:Direct to send the packaged workload file.

This value is substituted for the literal **dctgt** in the USRWKSNY CNTL member.

Source PLEX Id — Set this to the Id of the Source PLEX if you are using Connect:Direct to send the packaged workload file.

This value is substituted for the literal **dcsrc** in the USRWKSNY CNTL member.

X Press **<enter>** to display the generated job stream to Export the Workload.
X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the Housekeeping and Controls menu.

9.5 Import Workload

Use this panel to generate a job to rebuild all the Workload files from the single binary file received from the remote LPAR. The data may be have been sent by FTP or TSO XMIT.

Enter the Workload Name, the HLQ of the source system's VSAM files and the volume to allocate the imported VSAM files to.

Optionally, you may also need to specify override values for the SMS Storageclass and Managementclass to be used by the VSAM Import.

Note you may have to define to EZ-DB2 the subsystem referenced by the Workload on the source LPAR. Use EZDB2I Option 3.1 to add the Subsystem definition if it does not already exist.

If you wish to use the Statistics held on the remote LPAR when explaining SQL on this LPAR, then the statistics must have been migrated separately using the EZ-Stats menu options.

If you have used XMIT to transmit the exported workload files (see option 9.4) you can use option 9.6 to RECEIVE the transmitted files before importing the workload using this option.

Select option 9.5 to display the Import Workload panel as shown in the following figure:-

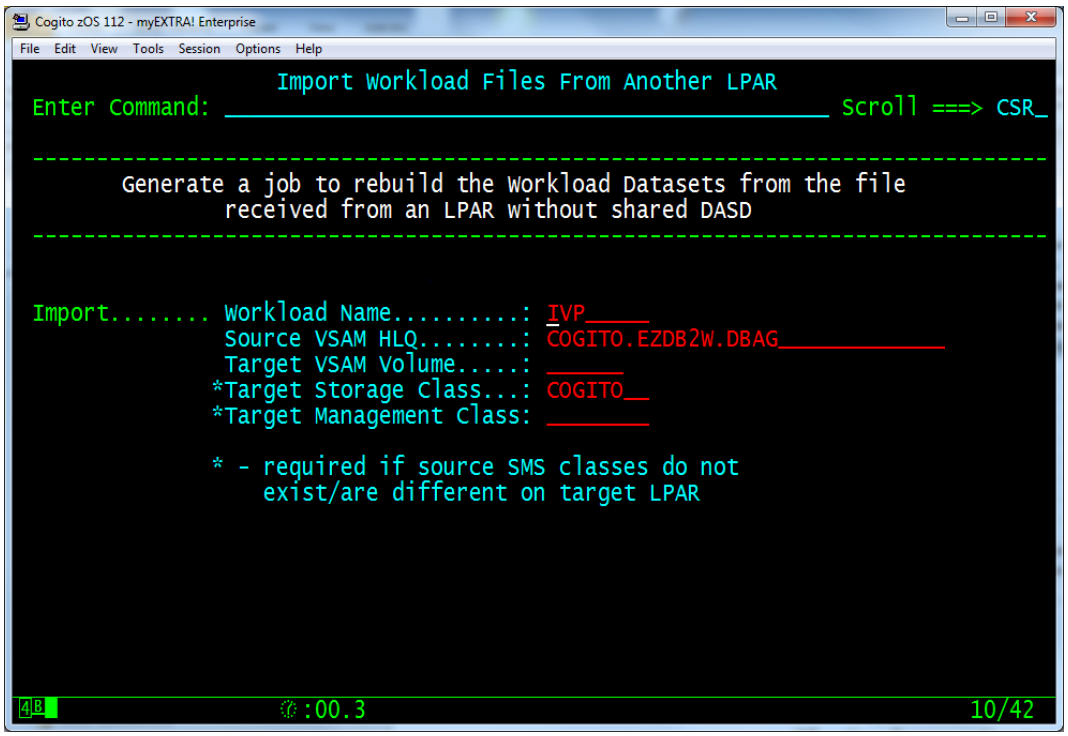

Figure 22 Import Workload

The following parameters may be specified on this panel.

Import

Workload Name	Set this field to the name of the Workload to be rebuilt from the binary file received from the remote LPAR. The binary file received, must reside in the 'ezdbpfx.INSTALL.UPLOAD' dataset, as a member with name of the workload.
Source VSAM HLQ	Set this field to the high level qualifier of the original VSAM files on the remote source LPAR. This is required by IDCAMS when running IMPORT for the packaged VSAM files.
***Target Storage Class**	If the exported workload VSAM files were originally allocated with Storageclass or Managementclass set, then the IDCAMS IMPORT assumes that the same classes are to be used on the LPAR running the IMPORT. If they do not exist and you do not set override values, then the IMPORT will fail. Set the optional STORAGECLASS and/or MANAGEMENTCLASS overrides for the VSAM IMPORT to use.
***Target Management Class**	If the exported workload VSAM files were originally allocated with Storageclass or Managementclass set, then the IDCAMS IMPORT assumes that the same classes are to be used on the LPAR running the IMPORT. If they do not exist and you do not set override values, then the IMPORT will fail. Set the optional STORAGECLASS and/or MANAGEMENTCLASS overrides for the VSAM IMPORT to use.

X Press <enter> to display the generated job stream to Import the Workload.
X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the Housekeeping and Controls menu.

9.6 TSO RECEIVE of files from another LPAR

Use this panel to receive datasets queued for your userid. This allows you to receive datasets/files that have been XMIT'd to you.

Each file queued is processed separately and you must specify the filename to receive the file into.

When sending workloads between LPARS where there is no shared DASD, the workload is exported and sent as a single file. This file must be received into the 'ezdbprfx.INSTALL.UPLOAD' dataset. The member name is the same as the workload name.

Select option 9.6 to display the TSO Receive panel as shown in the following figure:-

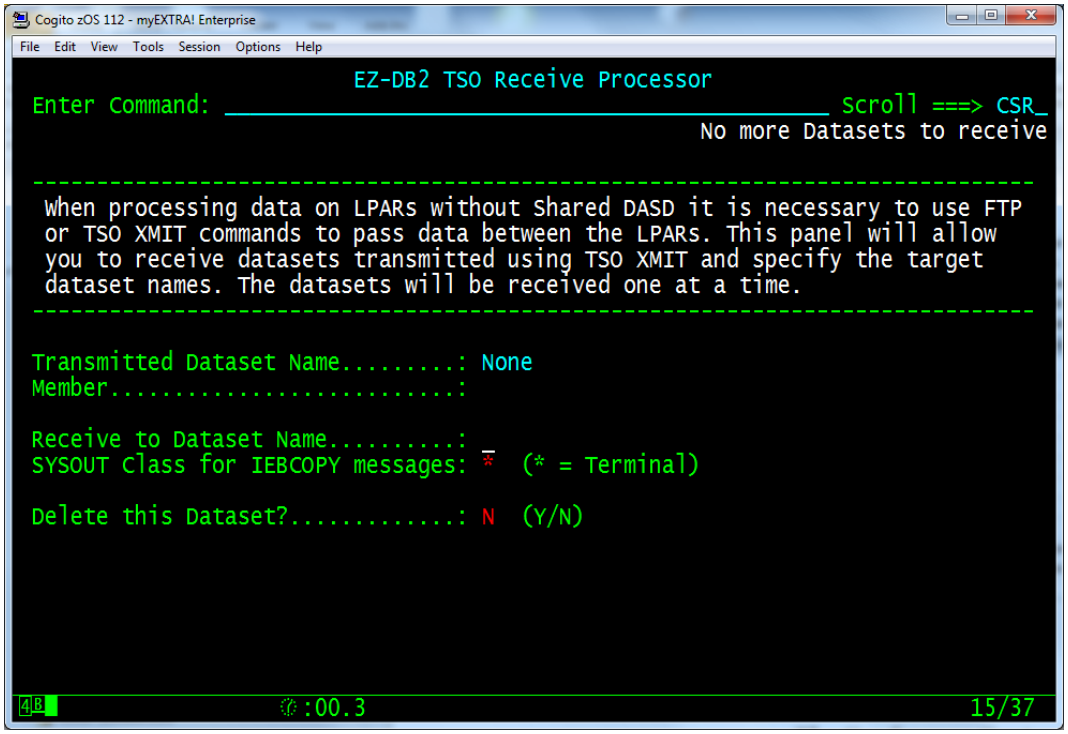

Figure 23 TSO receive files from another LPAR

The following parameters may be specified on this panel.

Transmitted Dataset name The name of the transmitted dataset to be received.

Member The name of the transmitted member to be received.

Receive To Dataset name The name of the dataset to receive the transmitted files. Note: when receiving workloads exported from another LPAR, the file must be received into 'ezdbpfx.INSTALL.UPLOAD'.

SYSOUT Class for IEBCOPY messages	Specify the SYSOUT Class to be used for the IEBCOPY messages. The default value is '*' for terminal display.
Delete this Dataset?	Specify Y to delete the transmitted dataset after it has been successfully received. The default value is N.

X Press **<enter>** to display the generated job stream to Receive the Workload.
X Type SUBMIT in the command line to submit the job for execution.
X Press the **<PF3>** key to return to the Housekeeping and Controls menu.

9.7 Import Plan Table data from a remote system using DRDA.

Select additional function option 9.7 to display the new panel as shown in the following figure:-

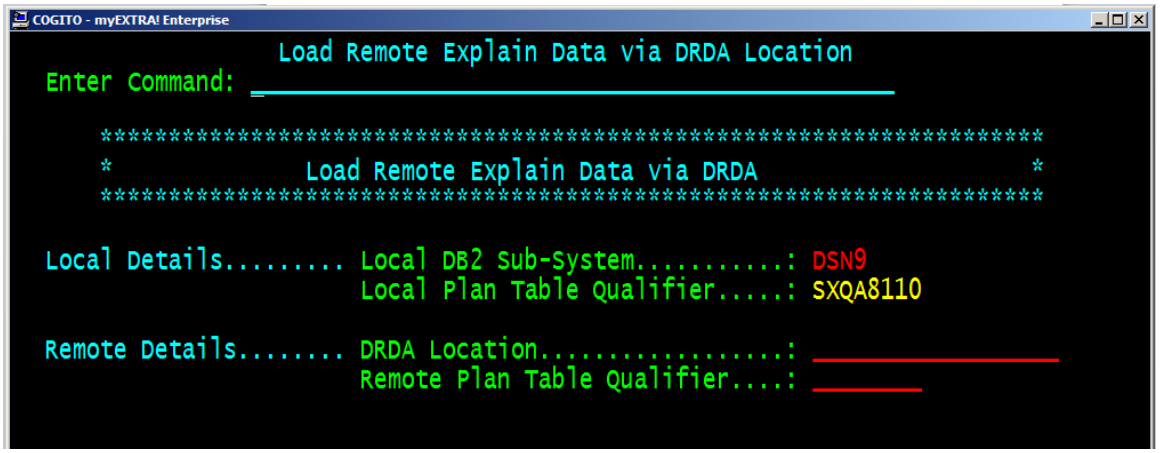

Figure 24 Load Remote Explain Data using DRDA

You can load a local set of plan tables, directly from a remote set of plan tables, using DRDA.

The generated JCL uses the DB2 LOAD from INCURSOR utility. The cursor accesses the remote tables via DRDA.

This option can only be used if you have DRDA defined correctly between the two systems. The specified LOCATION must already be defined in the local communications Database - refer to the IBM DB2 documentation on defining DRDA locations.

The generated JCL assumes the plan table format, is for the local DB2 version. It also assumes the following tables are available on both subsystems:

PLAN_TABLE
DSN_STATEMNT_TABLE
DSN_DETCOST_TABLE
DSN_FILTER_TABLE
DSN_STRUCT_TABLE
DSN_QUERY_TABLE

Modify the generated JCL if not all tables are available, or for a different DB2 version.

NOTE: The default generated JCL will cause the local table contents to be replaced. Modify the generated JCL if you need to retain the existing plan table data.
The following is a sample BIND to the remote location:

```
DSN SYSTEM(DB9G) RETRY(0) TEST(0)  <-- DB9G is the local subsystem
BIND PACKAGE (DBAGLU1.DSNUTIL) -   <-- DBAGLU1 is the DRDA location
   COPY(DSNUTIL.DSNUGSQL) -
   CURRENTDATA(NO) ISOLATION(CS) -
   ACTION(REPLACE) OPTIONS(COMMAND)
```

Local Details

Local DB2 Subsystem	The local DB2 subsystem ID.
Local Plan Table Qualifier	The PLAN Table qualifier on the local subsystem you are copying to.
Remote Details	
DRDA Location	The DRDA location of the remote system you are copying from.
Remote Plan Table Qualifier	The PLAN table qualifier on the remote system you are copying from.

Miscellaneous

9.8 Workload Job Run Status

Select option 9.8 to display the Run Status Display as shown in the following figure:-

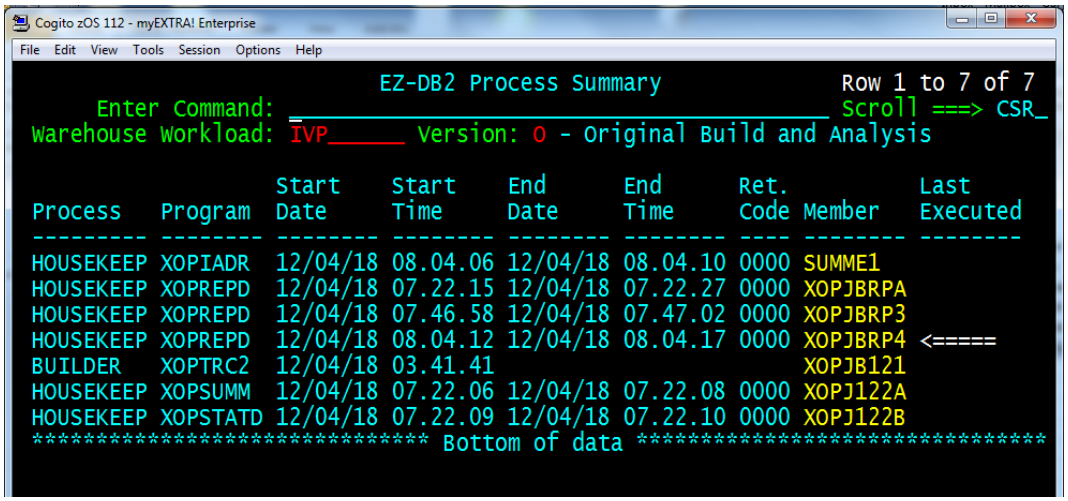

Figure 25 EZ-DB2 Run Status

This panel is used to identify the steps that have been run for the current Workload and the start and end date and time and completion code.

The panel also highlights the last step executed so that a user will know where to restart from.

The user may also select the LOGFILE member for any particular process by cursor selecting the Process member name and hitting the <PF4> drill function key.

9.9 Display DSNZPARM parameters

Select option 9.9 to display the DSNZPARM macro parameters for a given DB2 subsystem. The following panel is displayed:-

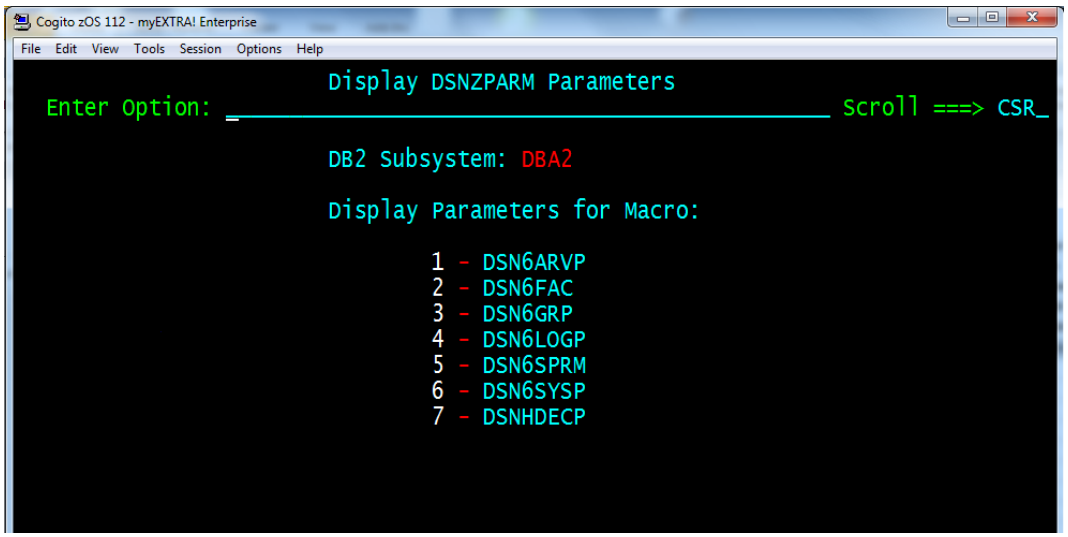

Figure 26 Display DSNZPARM macro parameters

DB2 Subsystem Enter the DB2 subsystem ID as applicable for the DB2 system to be displayed.

Enter the option number to display the DB2 parameters as required. For example, enter option 5 to display the DSN6SPRM parameters as shown in the following panel:-

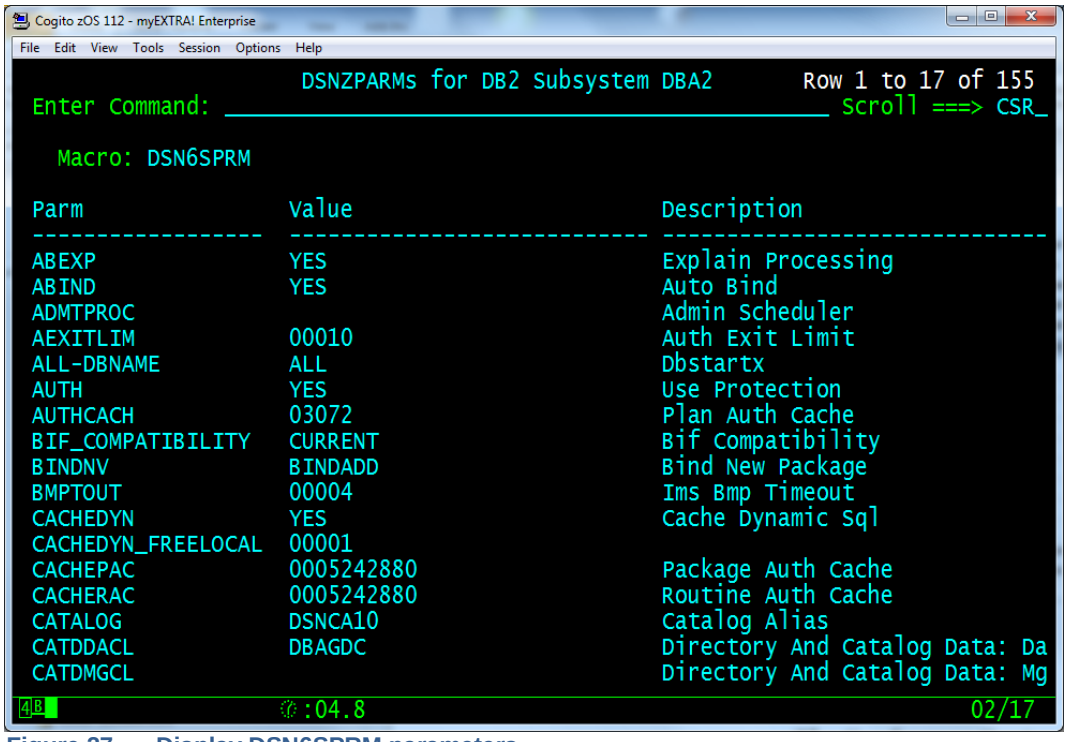

Figure 27 Display DSN6SPRM parameters

Refer to the IBM documentation - DB2 Installation Reference Guide for further information.

9.10 PFKeys and Commands

Select option 9.10 to display the EZ-D2 PFkey Functions panel as shown in the following figure:-

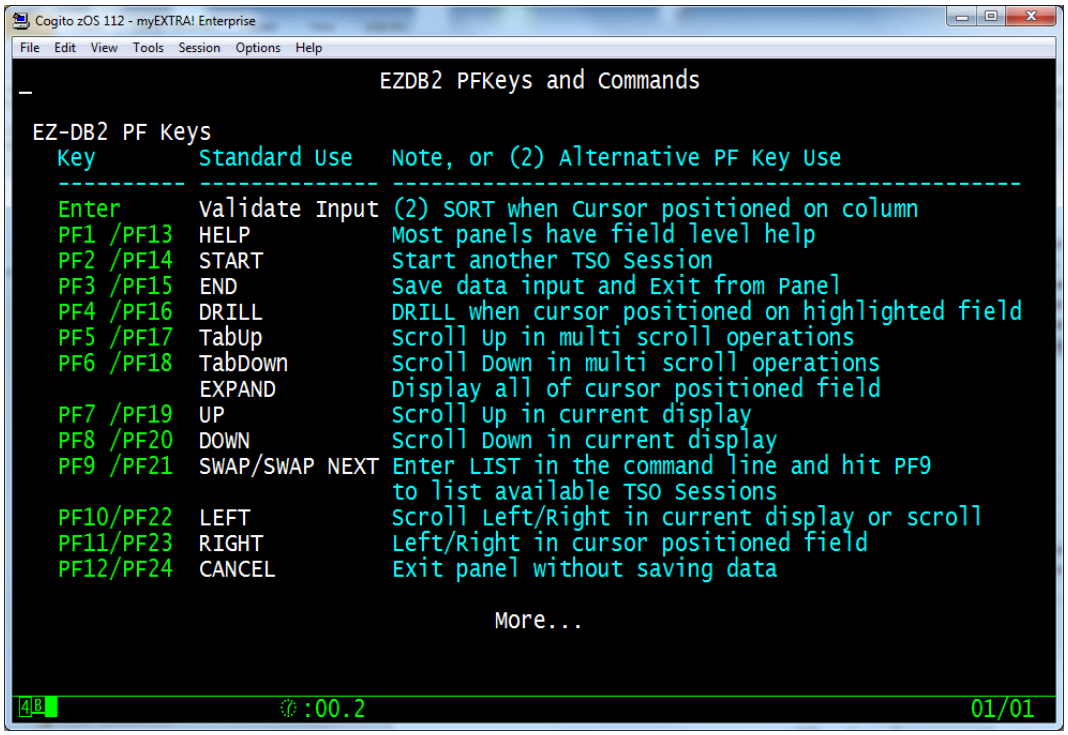

Figure 28 Display EZ-DB2 PFKey Functions

Press enter to display the EZ-DB2 commands shown in the following figure:-

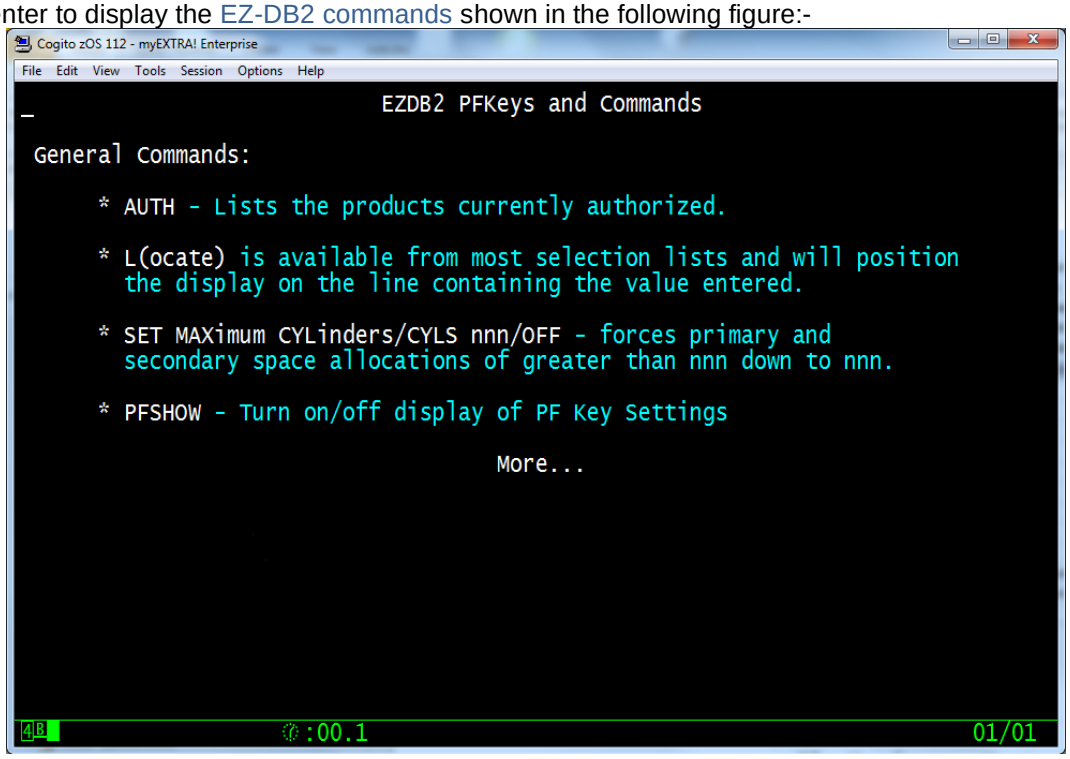

Figure 29 Display EZ-DB2 commands

Press enter to display the EZ-DB2 commands (2) shown in the following figure:-

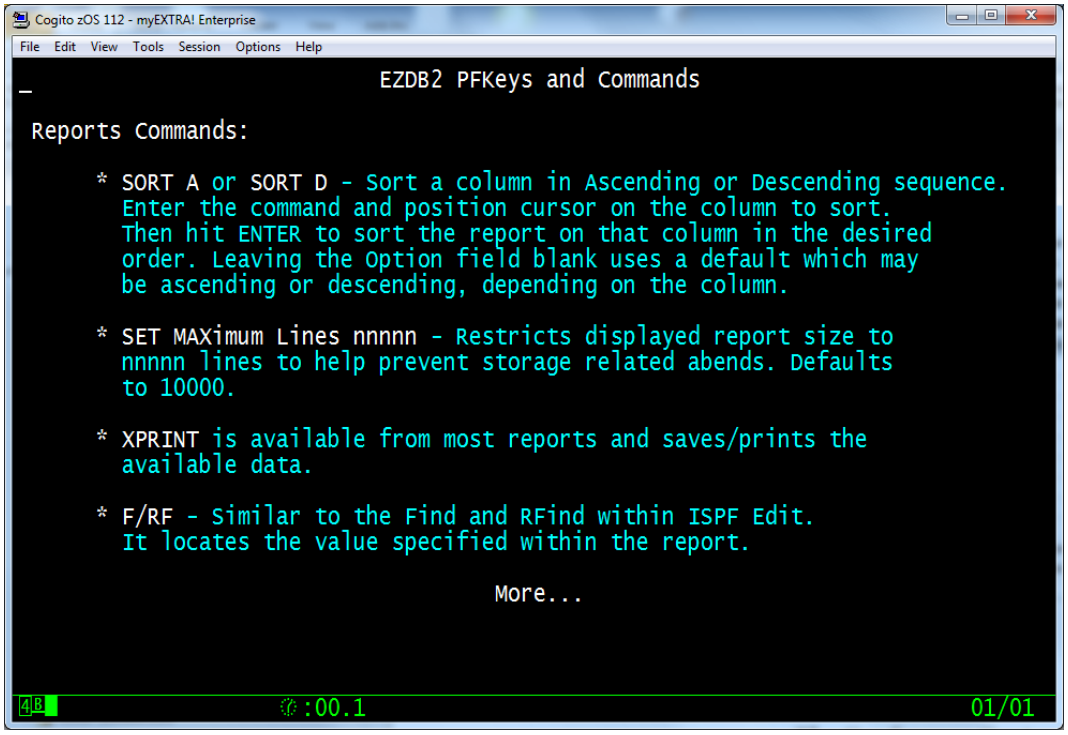

Figure 30 Display EZ-DB2 commands (2)

Press enter to display the EZ-DB2 commands (3) shown in the following figure:-

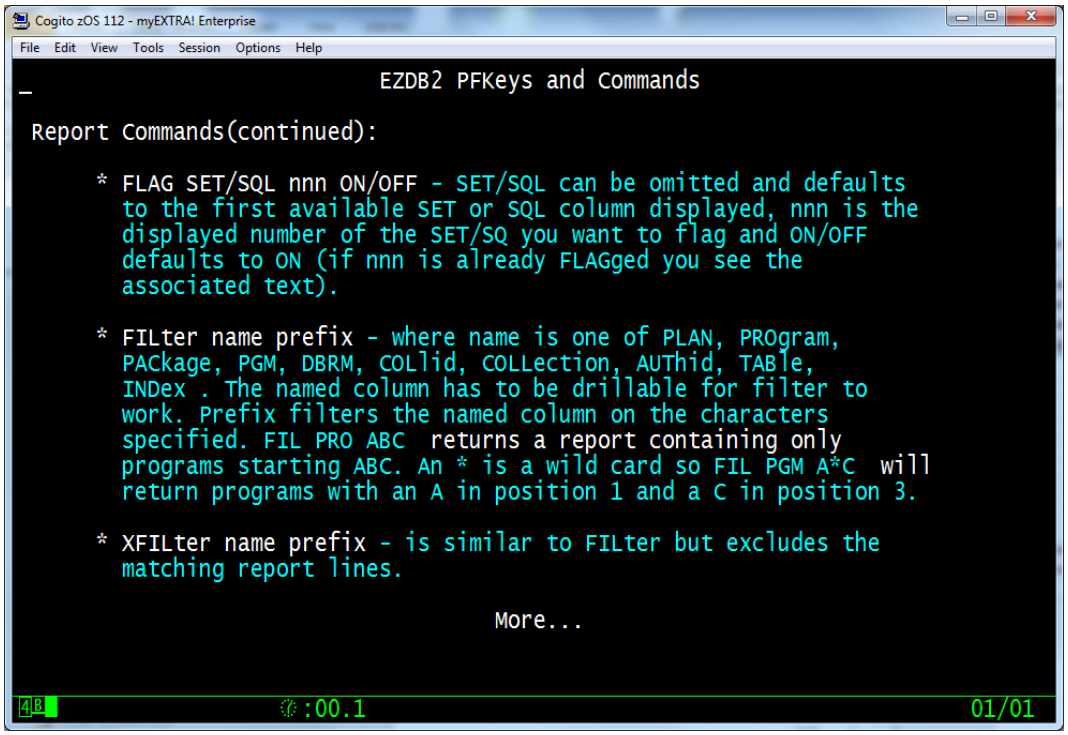

Figure 31 Display EZ-DB2 commands (3)

Press enter to display the EZ-DB2 commands (4) shown in the following figure:-

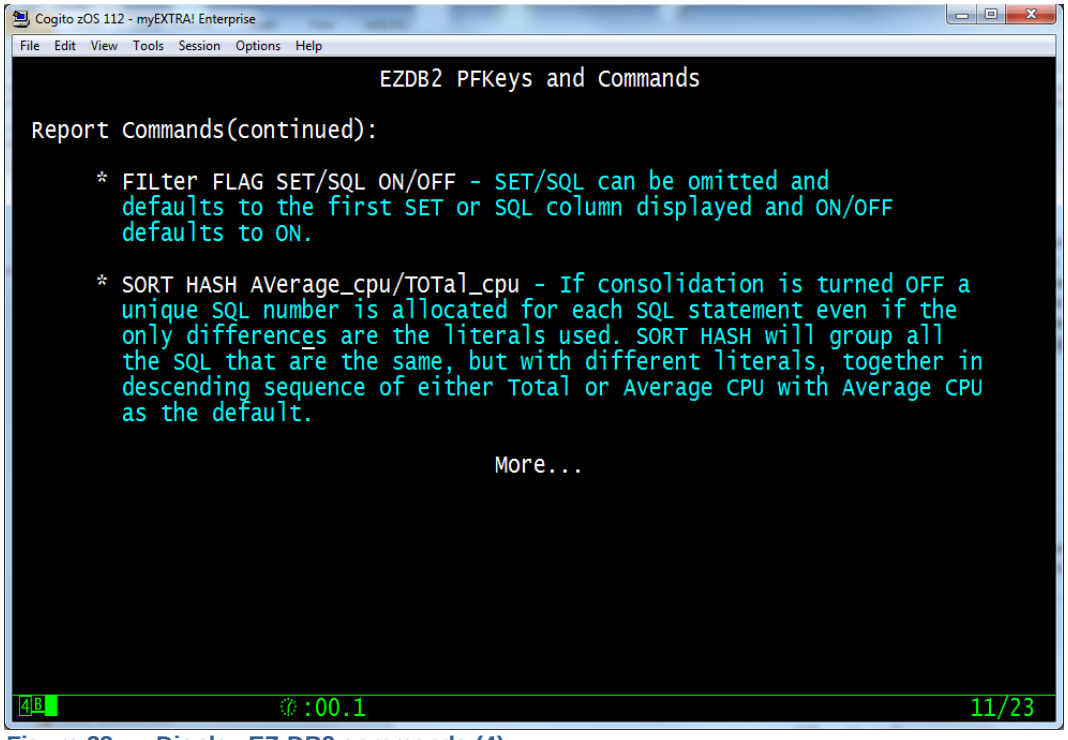

Figure 32 Display EZ-DB2 commands (4)

Press enter to display the EZ-DB2 commands (5) shown in the following figure:-

Figure 33 Display EZ-DB2 commands (5)

Press enter to display the EZ-DB2 Commands(6) shown in the following figure:-

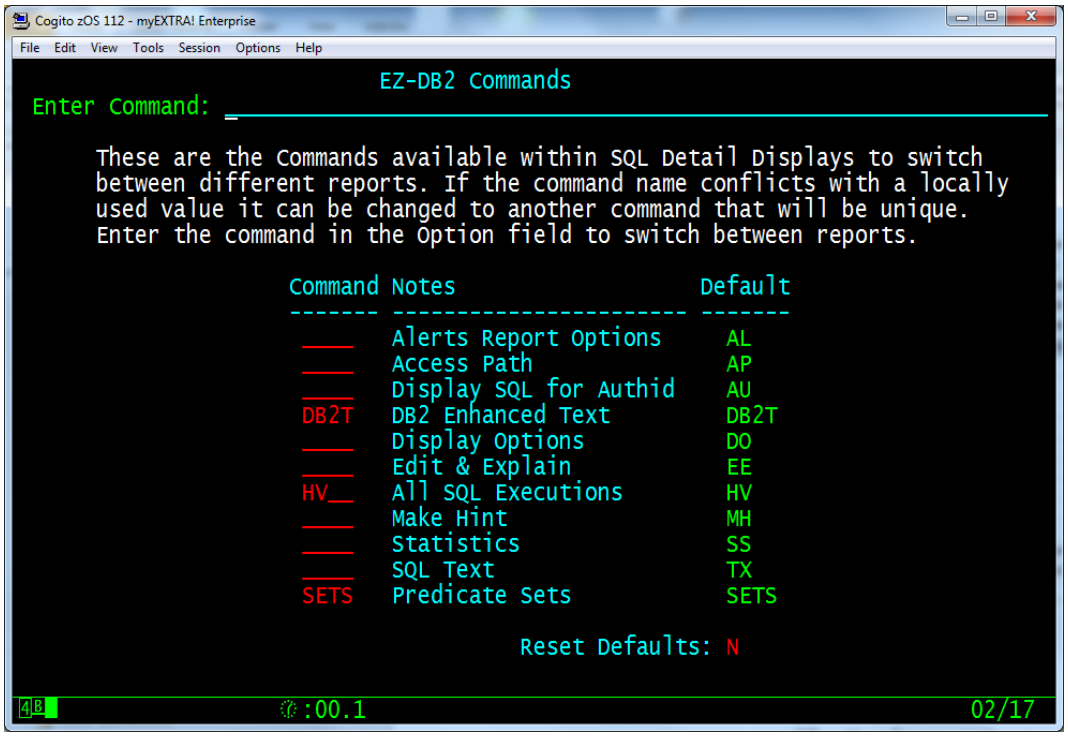

Figure 34 Display EZ-DB2 commands (6)

This panel allows you to change the **AP** command for **Access Path** to an alternative name of your choice. Specify the alternative command name in the **Command** field adjacent to the required command.

You can reset the commands back to their default settings by specifying **Y** in the **Reset Defaults** field at the bottom of the display.

9.11 List APARs Applied

Select option 9.11 to display the Applied APAR List panel as shown in the following figure:-

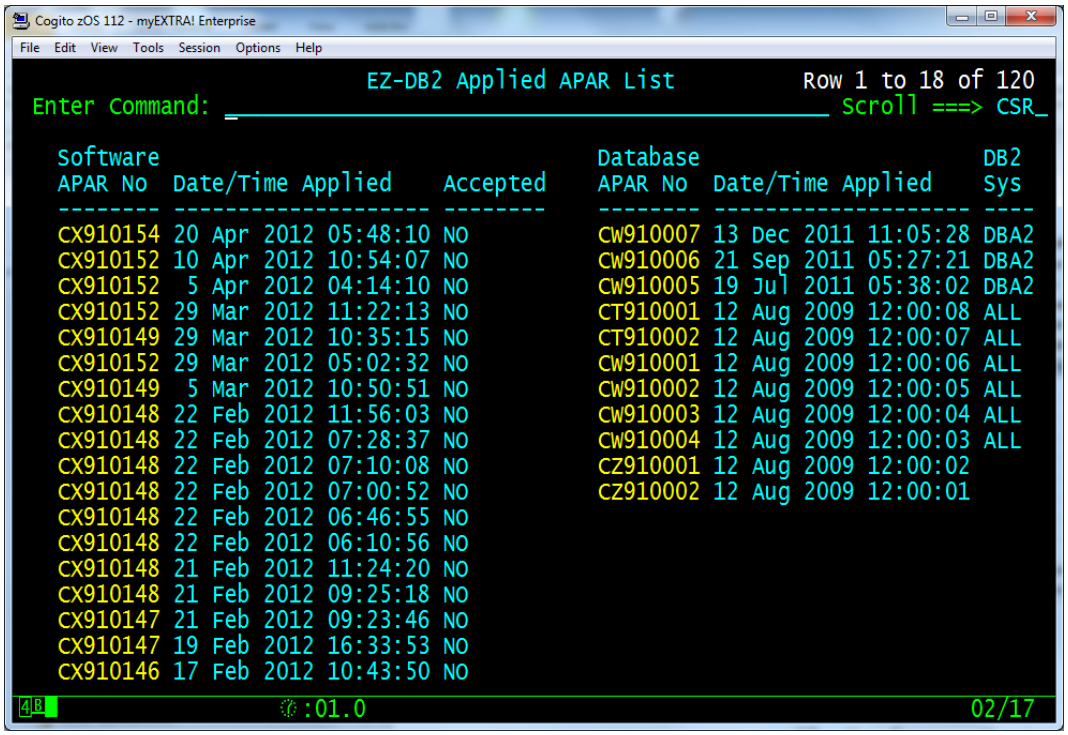

Figure 35 Display Applied APARs

The APARs are displayed in 2 columns, the left hand column shows Software APARs, the right hand column shows DDL APARs and one-off APARs. The following information is displayed on this panel for each column:-

APAR No — The APAR number. The 2 digit prefix indicates the type of APAR

CX	-	Software APAR
CT	-	DDL APAR for Trend Analysis DB
CW	-	DDL APAR for Warehouse DB
CZ	-	One-Off APAR

Applied by — The User ID who applied the APAR.

DRILL on the APAR number to display this information.

Date — The date the APAR was applied

Time — The time the APAR was applied

Accepted — Indicates if the APAR has been accepted (CX apars only)

DB2 Sys — The DB2 subsystem ID where the APAR is applied.

Refer to the Installation Reference Guide for further information.

9.12 FTP a file to EZ-DB2 Support

Select option 9.12 to display the FTP a file to EZ-DB2 Support panel as shown in the following figure:-

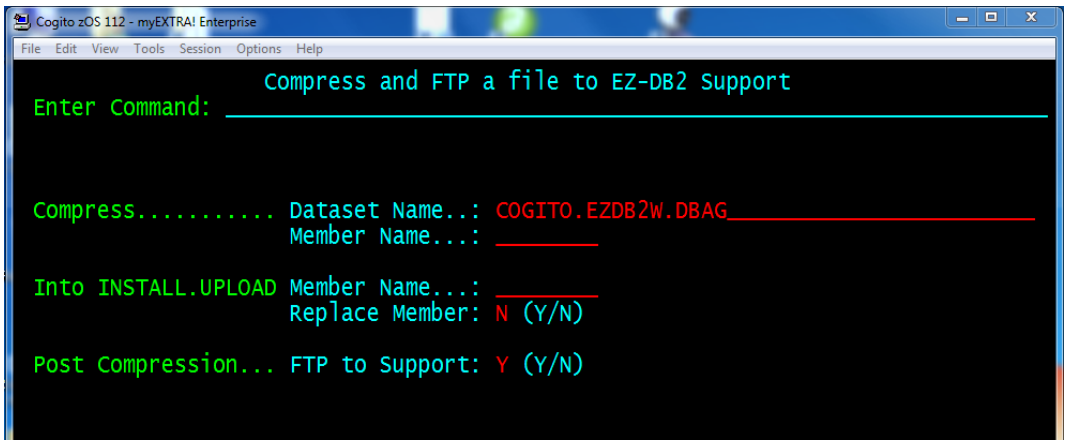

Figure 36 FTP a file to EZ-DB2 Support

On occasions EZ-DB2 support may request you to provide further information with respect to an issue you may be having. EZ-DB2 support may request that you FTP certain files to them. You can use this panel to compress and FTP the requested information to EZ-DB2 Support.

The following information is displayed on this panel:-

Compress

Dataset Name — Enter the name of the file to be compressed and transmitted or the name of a PDS which contains the file.

Member Name — If the file name specified is a PDS then you must specify the name of the member to be compressed and transmitted.

Into INSTALL.UPLOAD

Member Name — Specify a member name to use in your EZ-DB2 installation INSTALL.UPLOAD library into which the selected file will be compressed before transmission.

Replace Member — Enter Y to over write an existing member in the INSTALL.UPLOAD library. Enter N to avoid over writing an existing member.

Post Compression

FTP to Support — Enter Y if you wish to FTP the file to Support.

Once you have entered the information press <**Enter**> to generate a JCL stream which will compress the file and then FTP the compressed file to the EZ-DB2 Support server.

Whenever you send a file to EZ-DB2 Support please also send an email to support@cogito.co.uk to notify EZ-DB2 support what information has been sent.

9.13 FTP a file FROM EZ-DB2 support and Decompress

Select option 9.13**Error! Reference source not found.** to display the FTP file FROM EZ-DB2 support panel as shown in the following figure:-

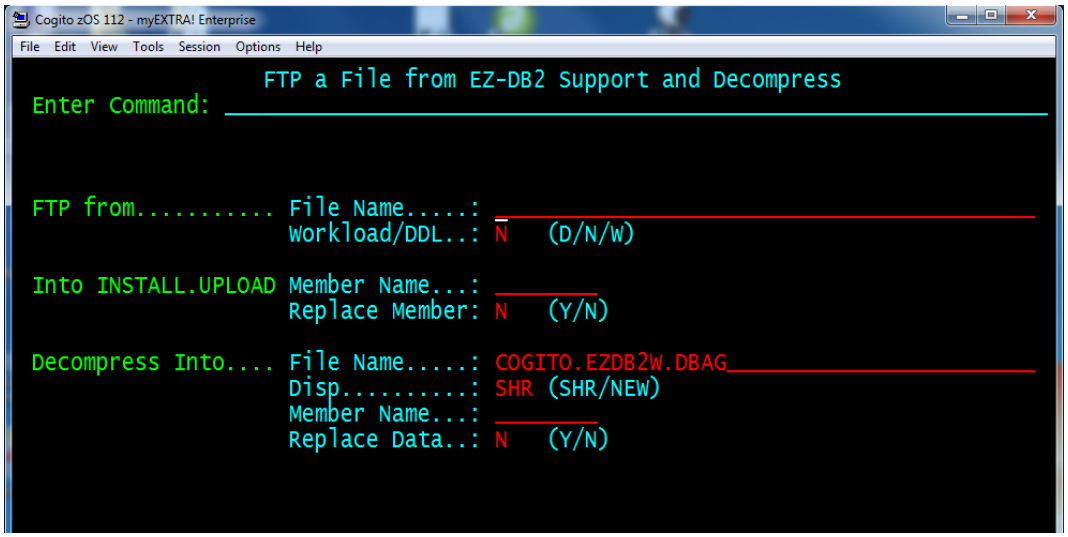

Figure 37 FTP file FROM EZ-DB2 support

On occasion you may be asked to RECEIVE a file from EZ-DB2 support. The file will be sent as a compressed file. Use this panel to receive the file and decompress into a dataset of your choosing,

The following information is displayed on the panel:-

FTP from

File Name Specify the name of the file as supplied by EZ-DB2 support.

Workload/DDL Specify whether the file is a Workload file or DDL.

D – DDL
W – Workload

Member Name Specify a member name to use in your EZ-DB2 installation INSTALL.UPLOAD library into which the selected file will be received.

Replace Member Enter Y to over write an existing member in the INSTALL.UPLOAD library. Enter N to avoid over writing an existing member.

Decompress Into

File Name Enter the name of the file or PDS to be de-compressed into

. **Disp** Specify the disposition of the file (SHR or NEW).

Member Name If the file name specified is a PDS then you must specify the name of the member.

Replace Data	Enter Y to overwrite an existing member in the specified dataset. . Enter N to avoid over writing an existing member.

9.14 FTP Database DDL to EZ-DB2 support

Select option 9.14 to display the FTP Database DDL to EZ-DB2 support panel as shown in the following figure:-

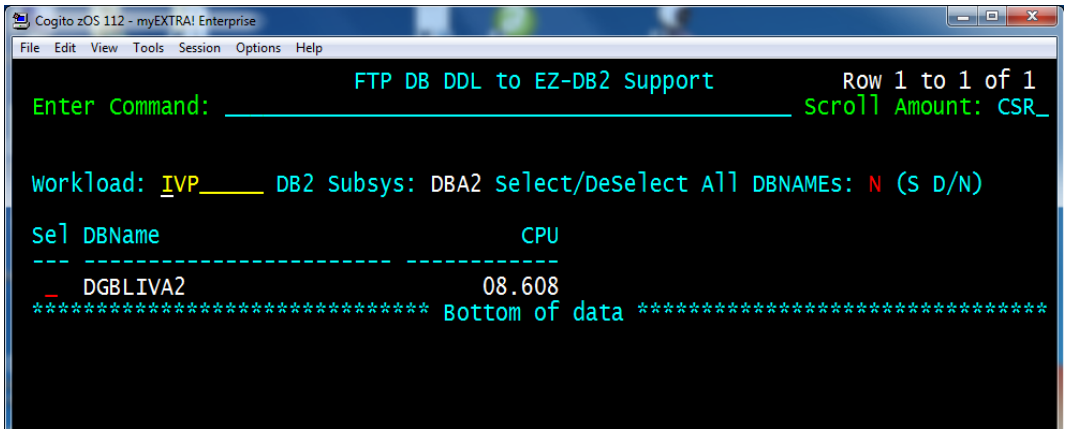

Figure 38 FTP Database DDL to EZ-DB2 support

This option is used to generate the DDL for selected databases and FTP the DDL to EZ-DB2 support.

Workload — Specify the name of the applicable EZ-DB2 workload. All of the database names for the selected database will be displayed.

DB2 Subsys — Specify the DB2 subsystem ID.

Select/Deselect All DBNAMES — You may select or deselect all of the displayed database names.

S – Select all DBNames
D – DeSelect all DBNames

See all Sel parameter below to select/deselect individual databases.

Sel — Enter S to select a particular database.

See also Select/Deselect All Dbnames above.

Dbname — The name of the database.

CPU — The total CPU cost for the database.

Once you have entered the information press <**Enter**> to generate a JCL stream which will generate the DDL for the selected database and then FTP the compressed file to the EZ-DB2 Support server.

Whenever you send a file to EZ-DB2 Support please also send an email to support@cogito.co.uk to notify EZ-DB2 support what information has been sent.

Appendix