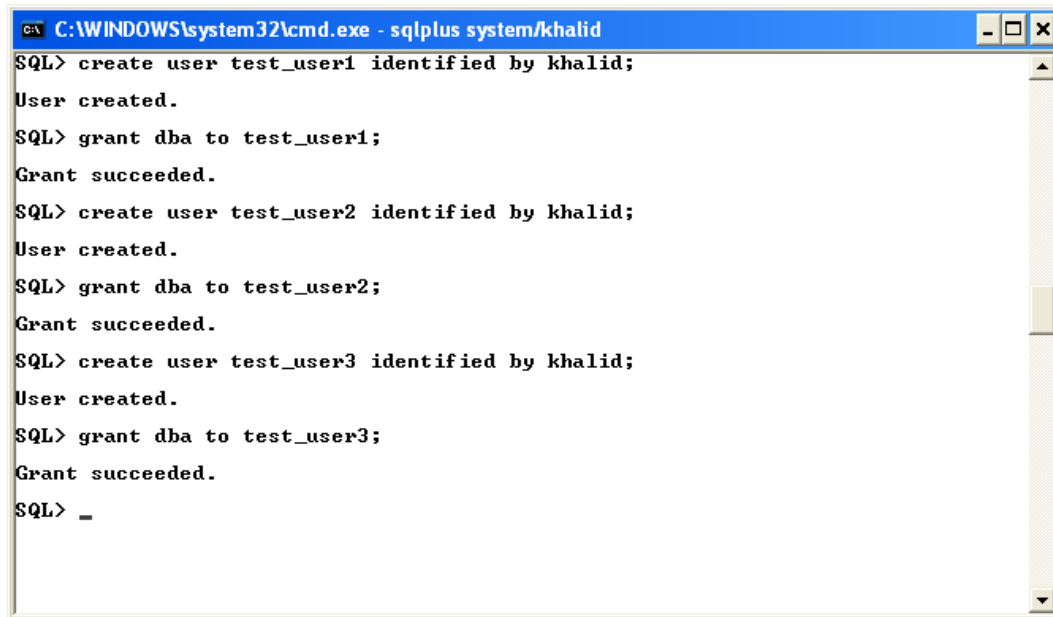


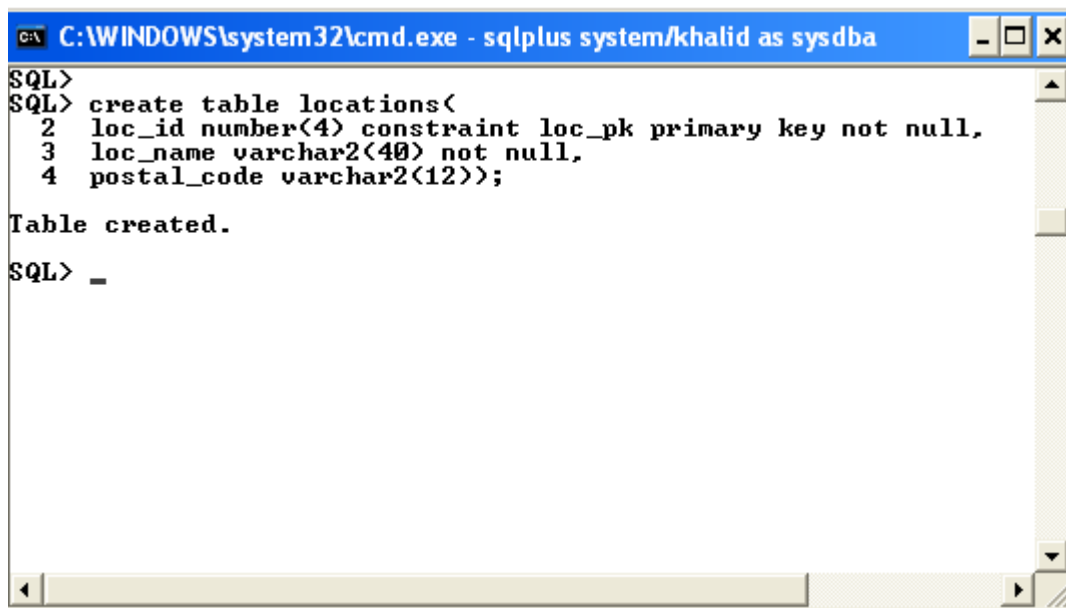
APPENDIX A

Oracle Database Creation (Oracle 10g)



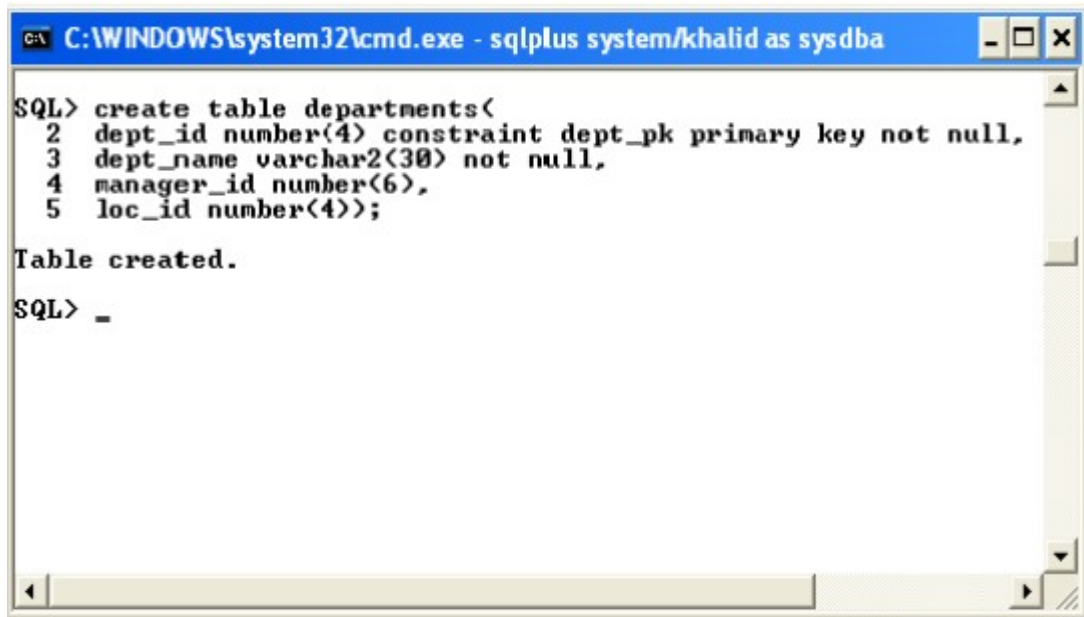
```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid
SQL> create user test_user1 identified by khalid;
User created.
SQL> grant dba to test_user1;
Grant succeeded.
SQL> create user test_user2 identified by khalid;
User created.
SQL> grant dba to test_user2;
Grant succeeded.
SQL> create user test_user3 identified by khalid;
User created.
SQL> grant dba to test_user3;
Grant succeeded.
SQL> _
```

Figure 1 – Creation of the users



```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid as sysdba
SQL>
SQL> create table locations(
  2 loc_id number(4) constraint loc_pk primary key not null,
  3 loc_name varchar2(40) not null,
  4 postal_code varchar2(12));
Table created.
SQL> _
```

Figure 2 – 'Locations' table creation



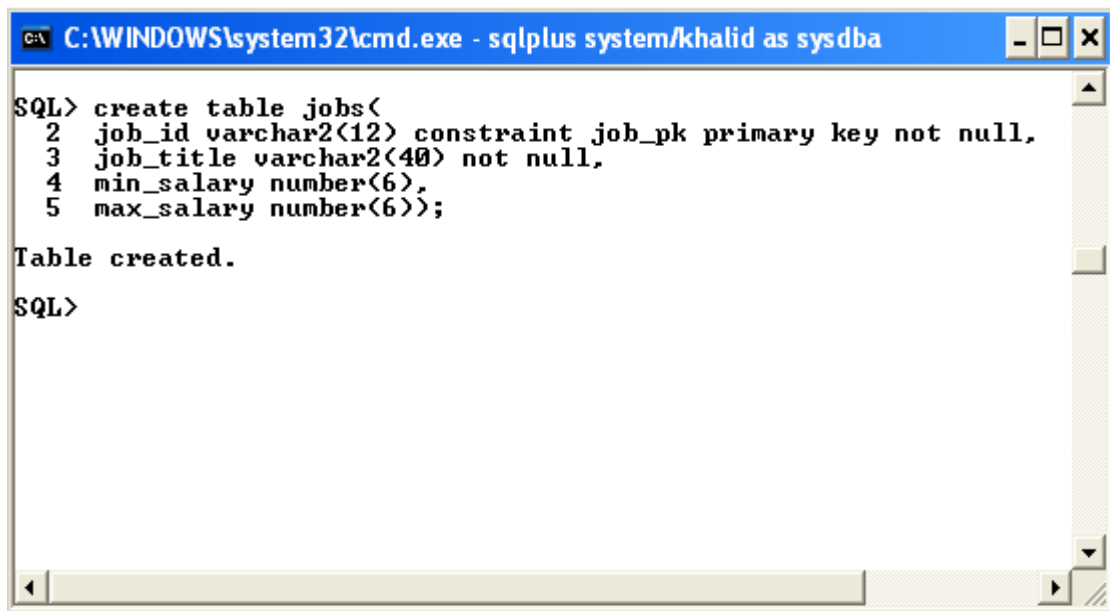
```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid as sysdba

SQL> create table departments(
  2  dept_id number(4) constraint dept_pk primary key not null,
  3  dept_name varchar2(30) not null,
  4  manager_id number(6),
  5  loc_id number(4));

Table created.

SQL> _
```

Figure 3 – 'Departments' table creation



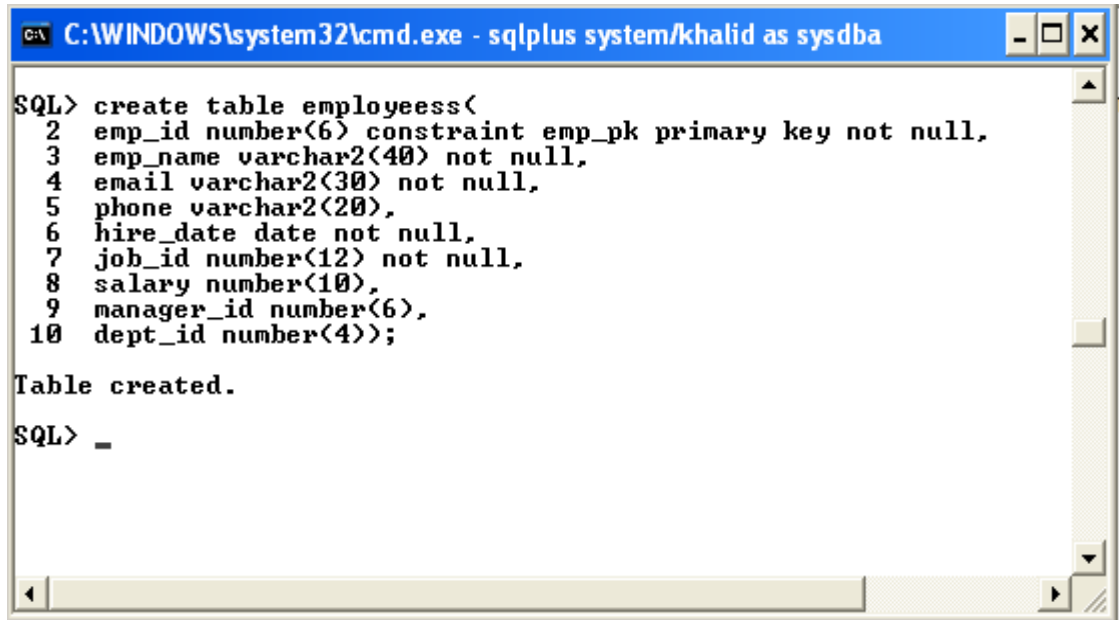
```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid as sysdba

SQL> create table jobs(
  2  job_id varchar2(12) constraint job_pk primary key not null,
  3  job_title varchar2(40) not null,
  4  min_salary number(6),
  5  max_salary number(6));

Table created.

SQL>
```

Figure 4 – 'Jobs' table creation



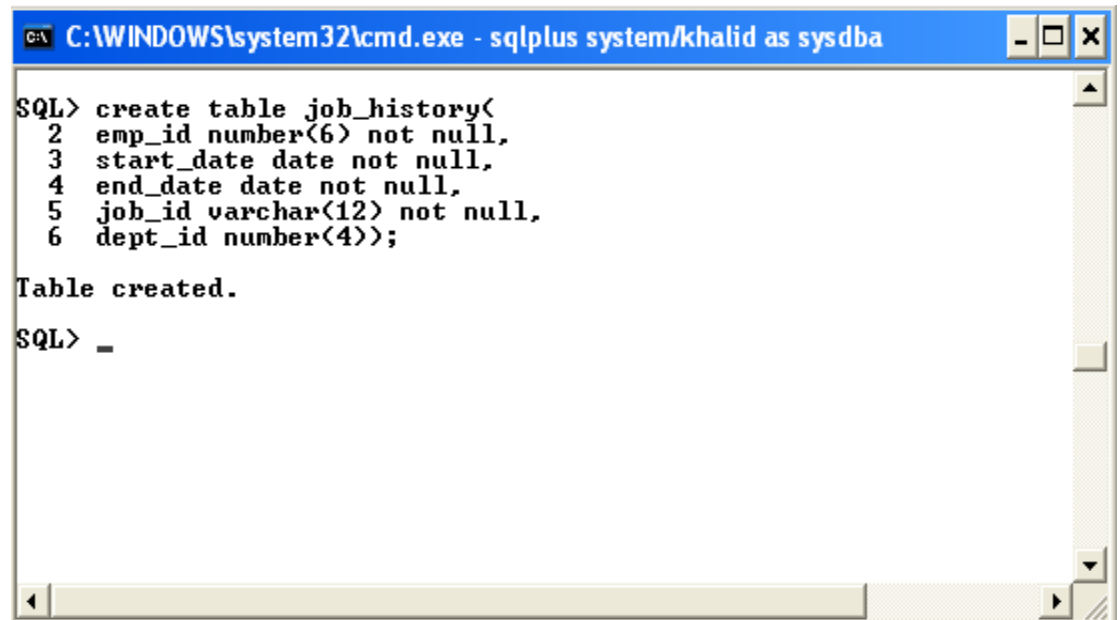
```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid as sysdba

SQL> create table employees(
  2 emp_id number(6) constraint emp_pk primary key not null,
  3 emp_name varchar2(40) not null,
  4 email varchar2(30) not null,
  5 phone varchar2(20),
  6 hire_date date not null,
  7 job_id number(12) not null,
  8 salary number(10),
  9 manager_id number(6),
 10 dept_id number(4));

Table created.

SQL> _
```

Figure 5 – ‘Employees’ table creation



```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid as sysdba

SQL> create table job_history(
  2 emp_id number(6) not null,
  3 start_date date not null,
  4 end_date date not null,
  5 job_id varchar(12) not null,
  6 dept_id number(4));

Table created.

SQL> _
```

Figure 6 – ‘Job_history’ table creation

```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid

SQL> conn test_user1/khalid;
Connected.
SQL>
SQL> create table u1t1(c number(10));
Table created.
SQL> create table u1t2(c number(10), d varchar2(10));
Table created.
SQL> _
```

Figure 7 – Test tables' creation into test user1

```
C:\WINDOWS\system32\cmd.exe - sqlplus system/khalid

SQL> conn test_user2/khalid;
Connected.
SQL>
SQL> create table u2t1(c number(10));
Table created.
SQL> create table u2t2(c number(10), d varchar2(10));
Table created.
SQL> _
```

Figure 8 – Test tables' creation into test user2

APPENDIX B

The Migration

1 About EnterpriseDB

EnterpriseDB is the leading provider of enterprise class products and services based on PostgreSQL, the world's most advanced open source database. The company's Postgres Plus products are ideally suited for transaction-intensive applications requiring superior performance, massive scalability, and compatibility with proprietary database products. PostgresPlus also provides an economical open source alternative or complement to proprietary databases without sacrificing features or quality.

EnterpriseDB has offices in North America, Europe, and Asia. The company was founded in 2004 and is headquartered in Edison, N.J. For more information, please call +1-732-331-1300 or visit <http://www.enterprisedb.com>.

2 The Migration from Oracle to Postgres

2.1 Installing Migration Studio

Migration Studio is distributed with PostgresPlus Advanced Server. You can download the latest version of PostgresPlus Advanced Server at:

<http://www.enterprisedb.com/products/download.do>.

Choose the version that is compatible with your system, and click the Download link. Complete the registration information, and click Submit to start the download. Extract the installation folder from the downloaded archive.

- If you are installing into Windows, navigate to the installer, and double click the PostgresPlus Advanced Server icon to start the installation wizard.

- If you are installing into Linux, open a command window, and navigate into the pgplus-advsvr-linux-x86_830106 directory. Give yourself super user privileges, and then start the installation wizard with the command \$sh pgplus-advsvr-linux-x86_830106.bin.

Follow the on-screen directions, choosing the appropriate responses, and clicking Next to progress through the installation. When the Features dialog opens, make sure the box next to the Developer/Client Tools option is checked to include the Migration Studio Graphical tool in the installation (shown in Figure 4.1). Click Next to continue.

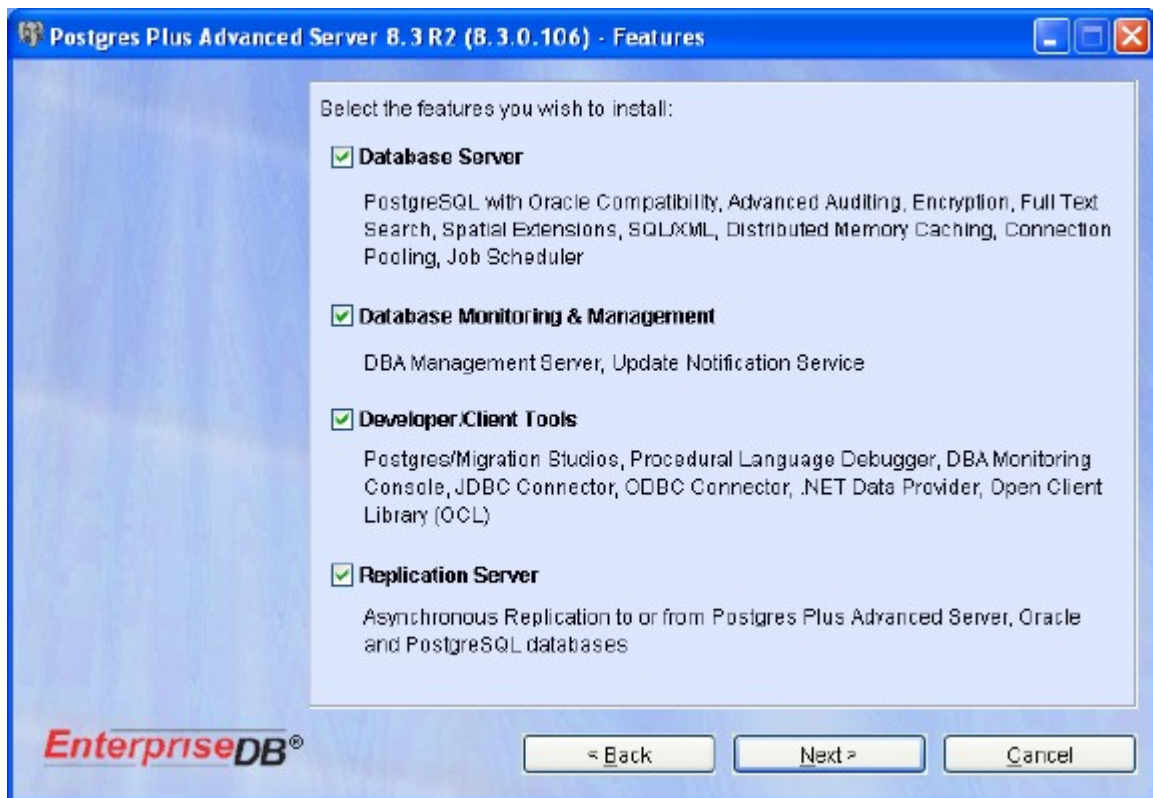


Figure 1 - Advanced Server Installation Dialog

Continue through the dialogs until the installation completes. Postgres Plus Advanced Server should now appear on the start menu.

You must install a JDBC driver before migrating from an Oracle or MySQL server. You can find a link to the appropriate driver in the Third Party Drivers section of the EnterpriseDB downloads page. Download the driver file, and move it into the \jre\lib\ext\ directory under the Postgres Plus Advanced Server home directory.

JDBC drivers for Sybase and SQL Server are distributed and installed with Advanced Server.

See figures 2 and 3 below.

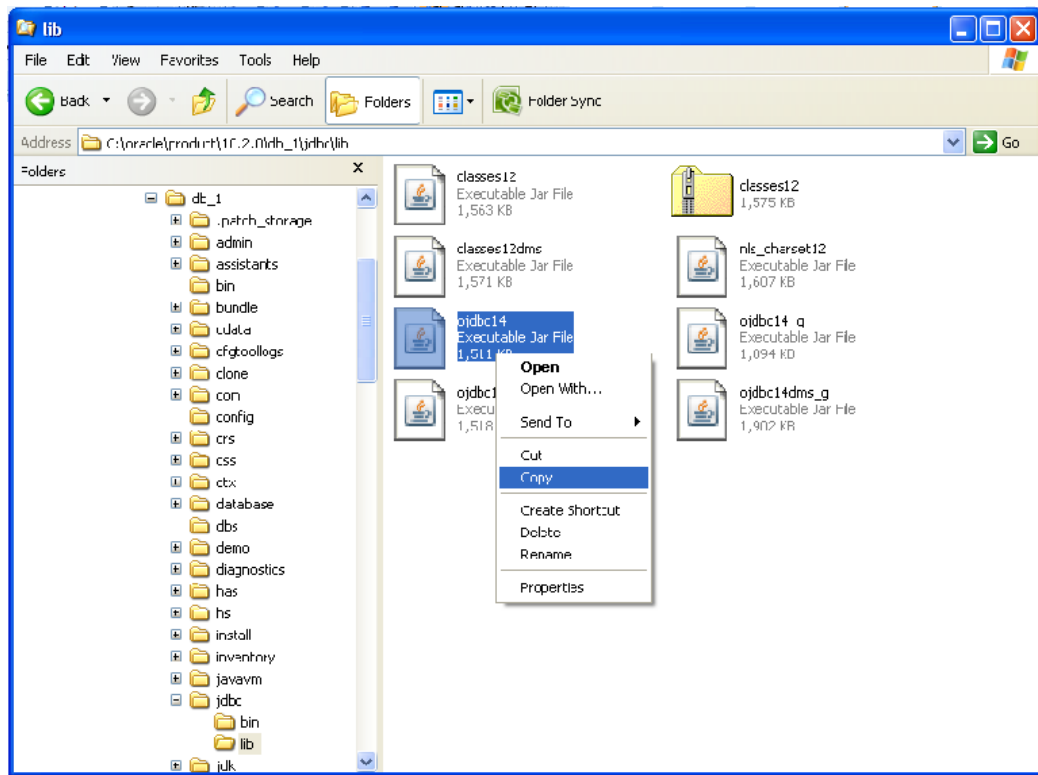


Figure 2 – Copying Oracle JDBC Driver from Oracle Home

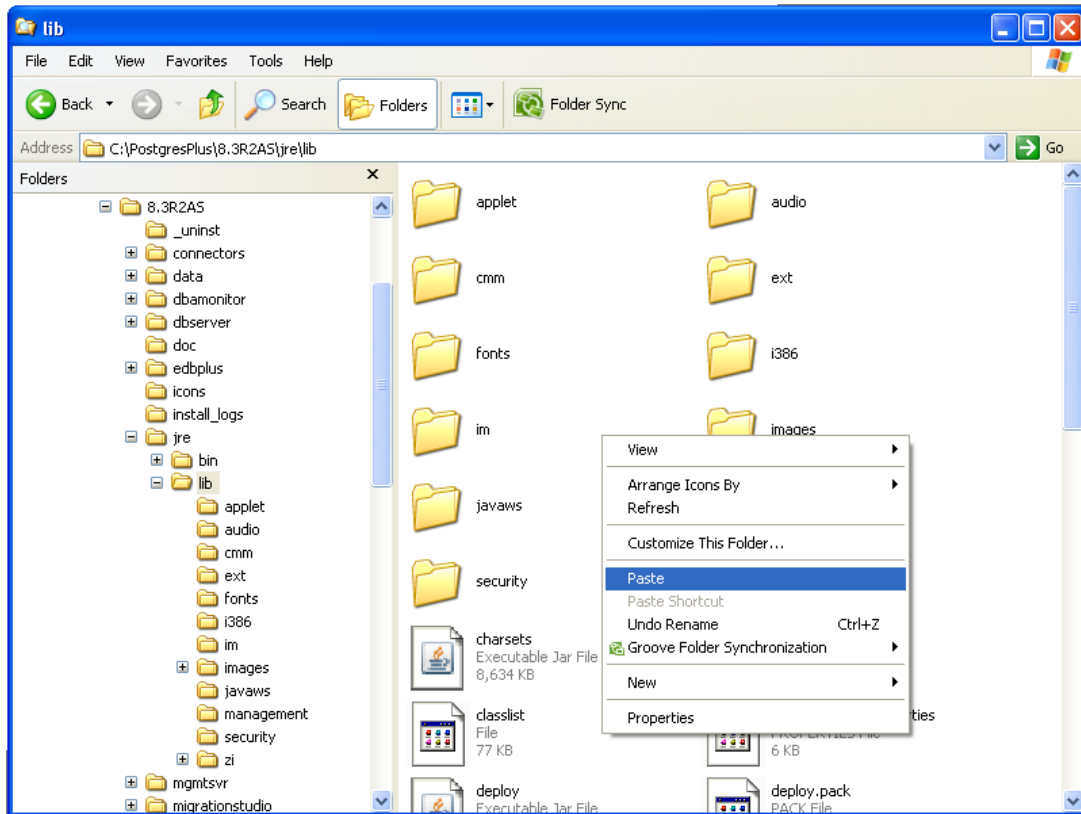


Figure 3 - Locating Oracle JDBC Driver into jre Directory of PostgresPlus

To run Migration Studio, open the start menu and select PostgresPlus Advanced Server 8.3R2; choose Migration Studio from the pull-aside menu to open the EnterpriseDB Migration Studio 8.3 window, as shown in Figure 2.

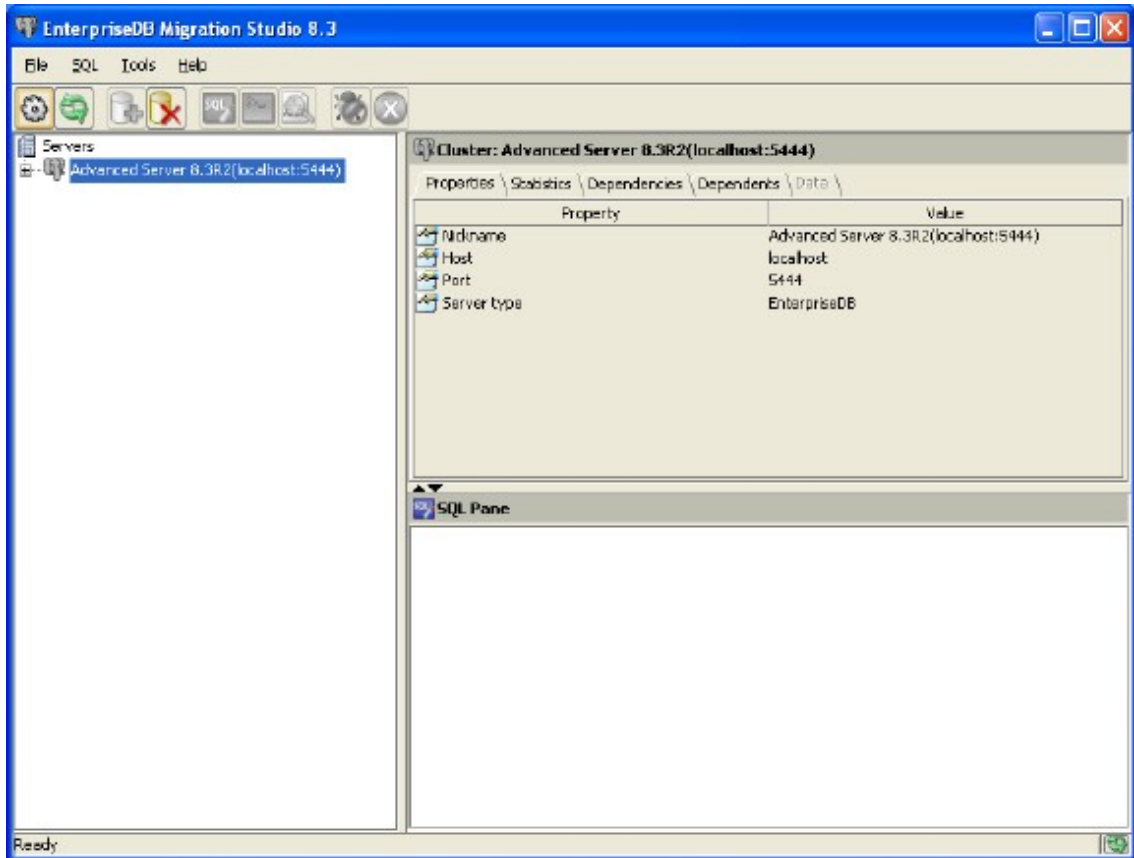


Figure 4 - EnterpriseDB Migration Studio 8.3

Alternatively, you can start Migration Studio by opening a Command window and navigating to the migration studio directory under the current installation of PostgresPlus. Enter the command: runMigrationStudio.bat to open the Migration Studio window.

2.2 Connecting an Application to Advanced Server

In the case of a Java application, change the JDBC driver name (`Class.forName`) and JDBC URL.

A Java application running on Oracle might have the following connection properties:

```
Class.forName("oracle.jdbc.driver.OracleDriver");
Connection con =
```

```
DriverManager.getConnection("jdbc:oracle:thin:@localhost:1521:xe",
                            "user",
                            "password")
```

Change the connection string to connect to an instance of Advanced Server:

```
Class.forName("com.edb.Driver")
Connection con =
    DriverManager.getConnection("jdbc:edb://localhost:5444/edb",
                                "user",
                                "password");
```

1. Use an ODBC data source administrator to create a data source that defines the connection properties for the Advanced Server database.

Most Linux and Windows systems include graphical tools that allow you to create and edit ODBC data sources. After installing ODBC, check the Administrative Tools menu for a link to the ODBC Data Source Administrator. Click the Add button to start the Create New Data Source wizard; complete the dialogs to define the Advanced Server data source.

2. Change the application to use the new data source.

The application will contain a call to `SQLConnect` (or possibly `SQLDriverConnect`); edit the invocation to change the data source name. In the following example, the data source is "OracleDSN":

```
result = SQLConnect(conHandle,           // Connection
                    handle (returned)
                    "OracleDSN", SQL_NTS, // Data source
                    name
                    username, SQL_NTS,   // User name
                    password, SQL_NTS);  // Password
```

To connect to an instance of Advanced Server defined in a data source named "AdvServerDSN", change the data source name:

```
result = SQLConnect(conHandle,           // Connection
                    handle (returned)
                    "AdvServerDSN", SQL_NTS, // Data source
                    name
                    username, SQL_NTS,   // User name
                    password, SQL_NTS);  // Password
```

2.3 Servers Creation

These screen shots present the steps could be walked through to create an Oracle server. The same steps could be followed to create an EnterpriseDB server.

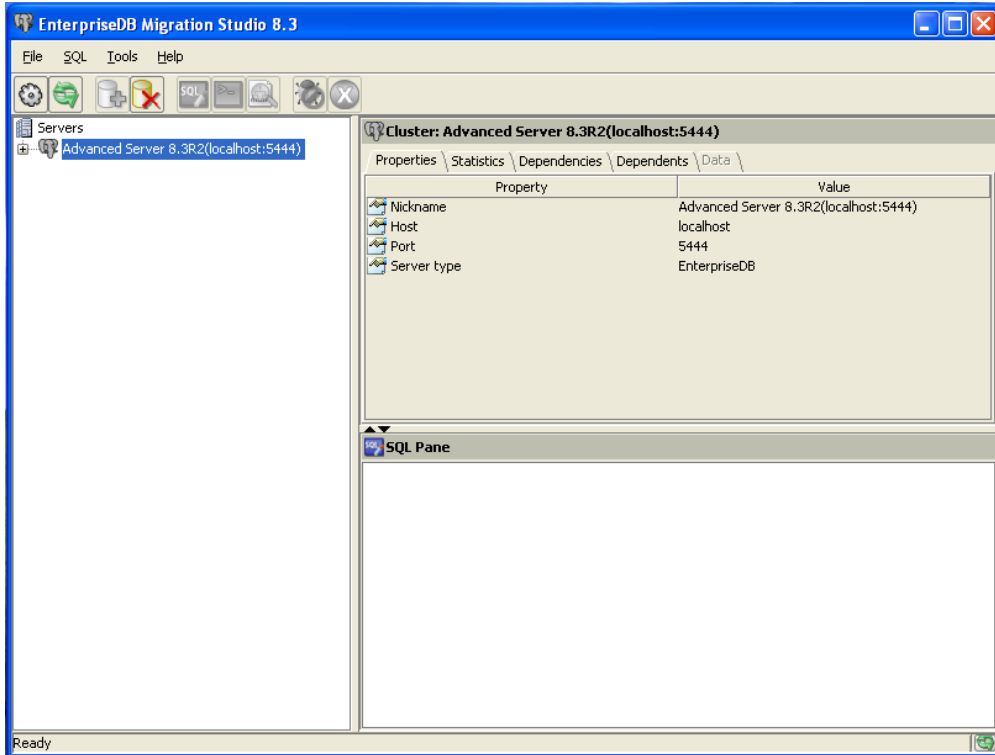


Figure 5 – EnterpriseDB migration studio 8.3 main menu

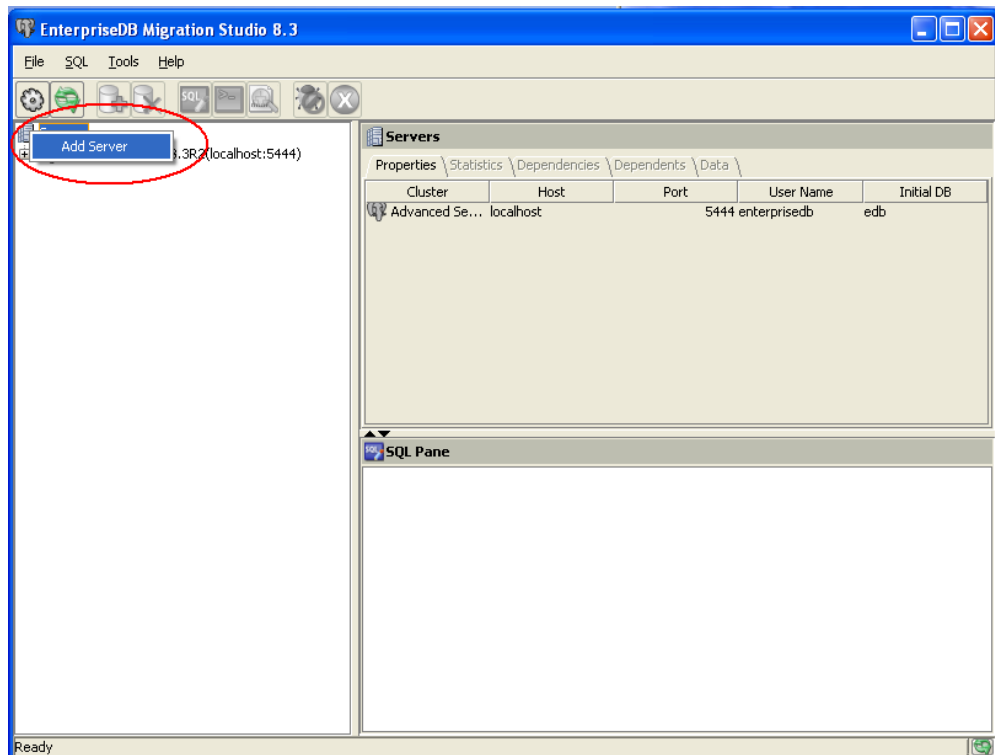


Figure 6 – EnterpriseDB migration studio 8.3, add new server

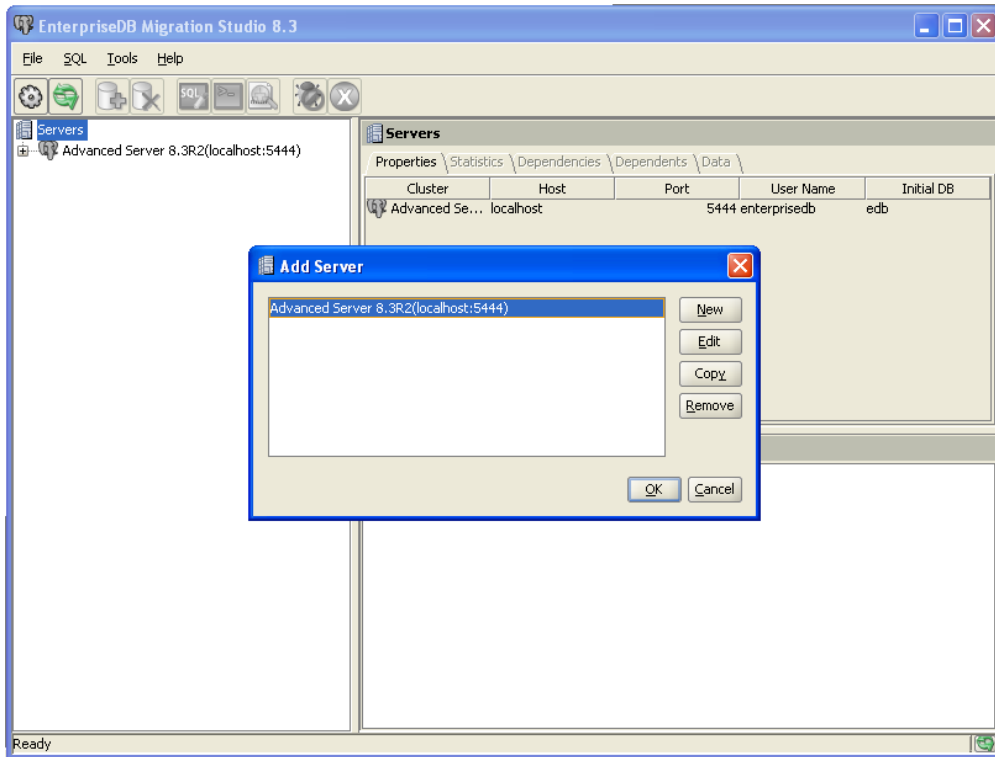


Figure 7 – EnterpriseDB migration studio 8.3, add new server dialog box

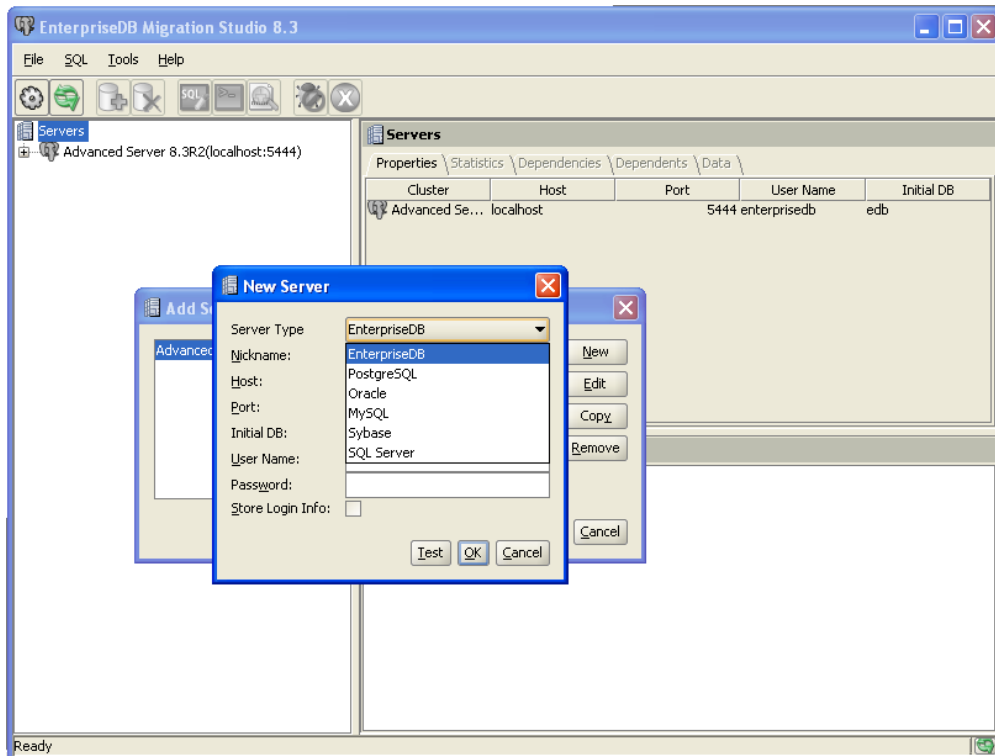


Figure 8 – EnterpriseDB migration studio 8.3, select the server's type

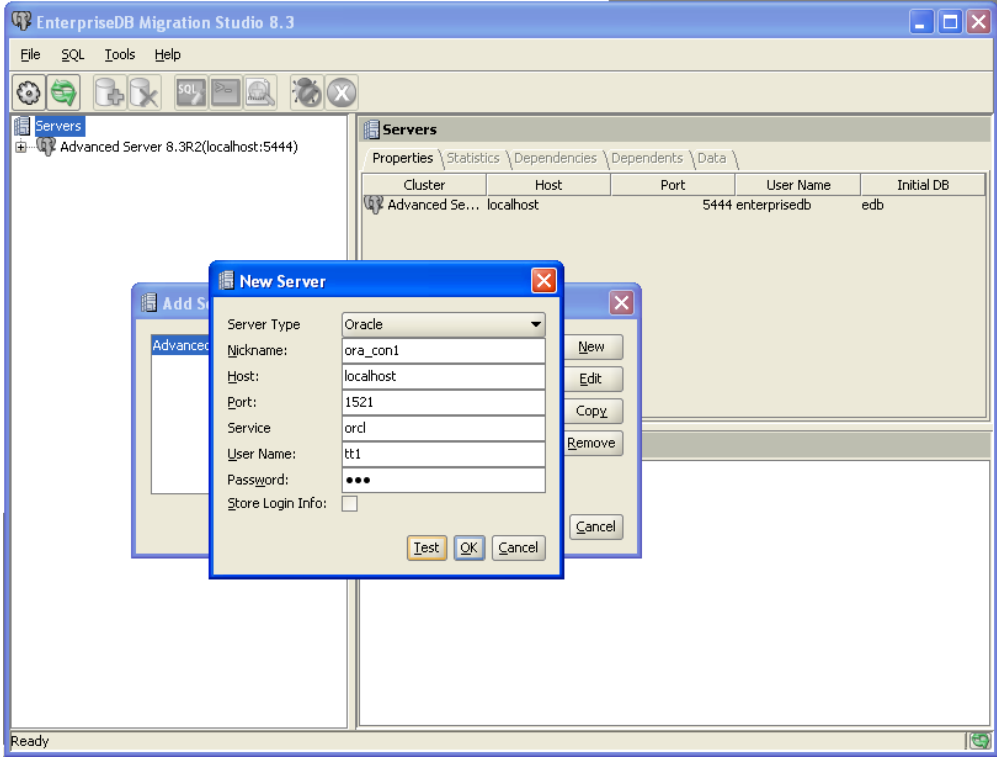


Figure 9 – EnterpriseDB migration studio 8.3 set the server information

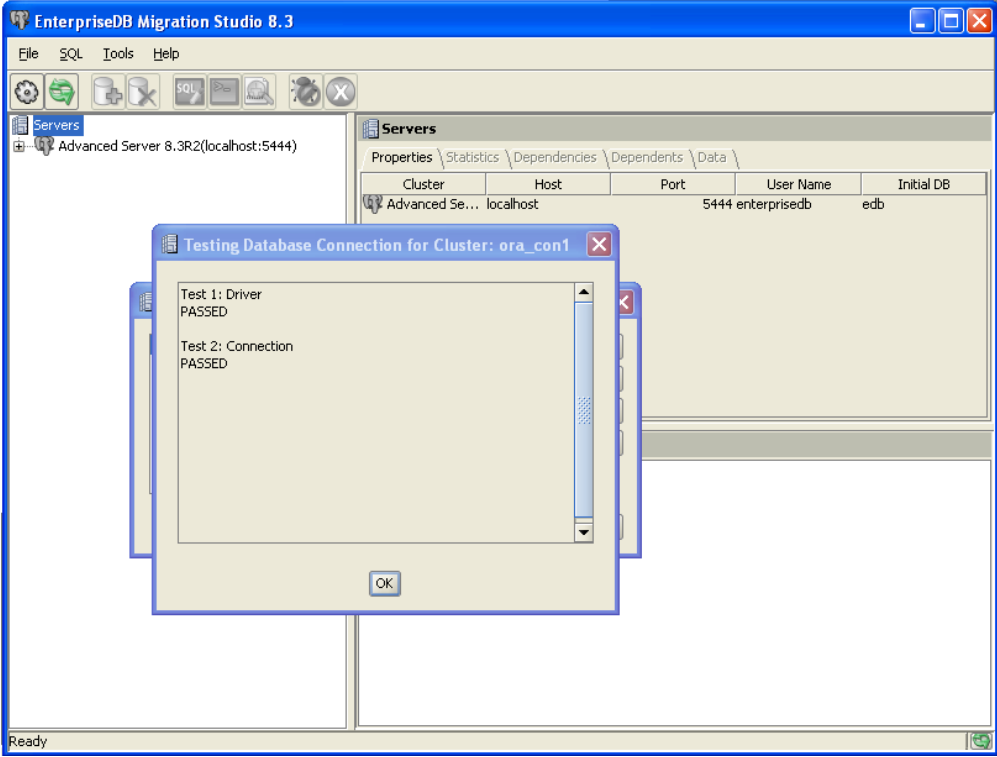


Figure 10 – EnterpriseDB migration studio 8.3, test the created server

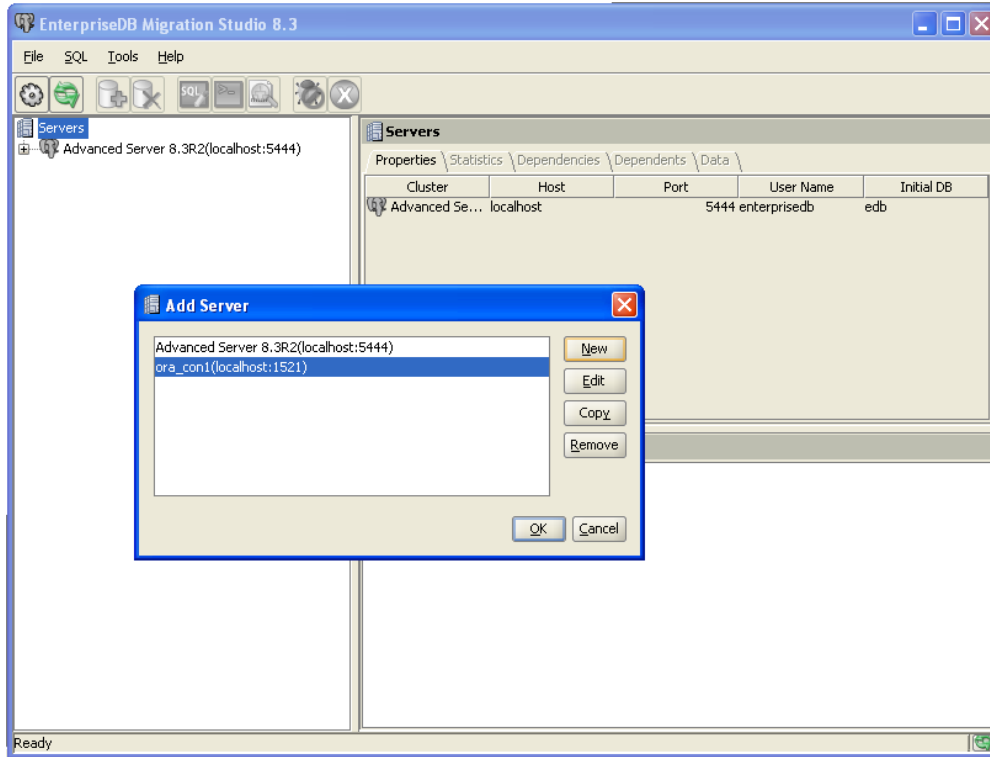


Figure 11 – EnterpriseDB migration studio 8.3, new server is ready to be added

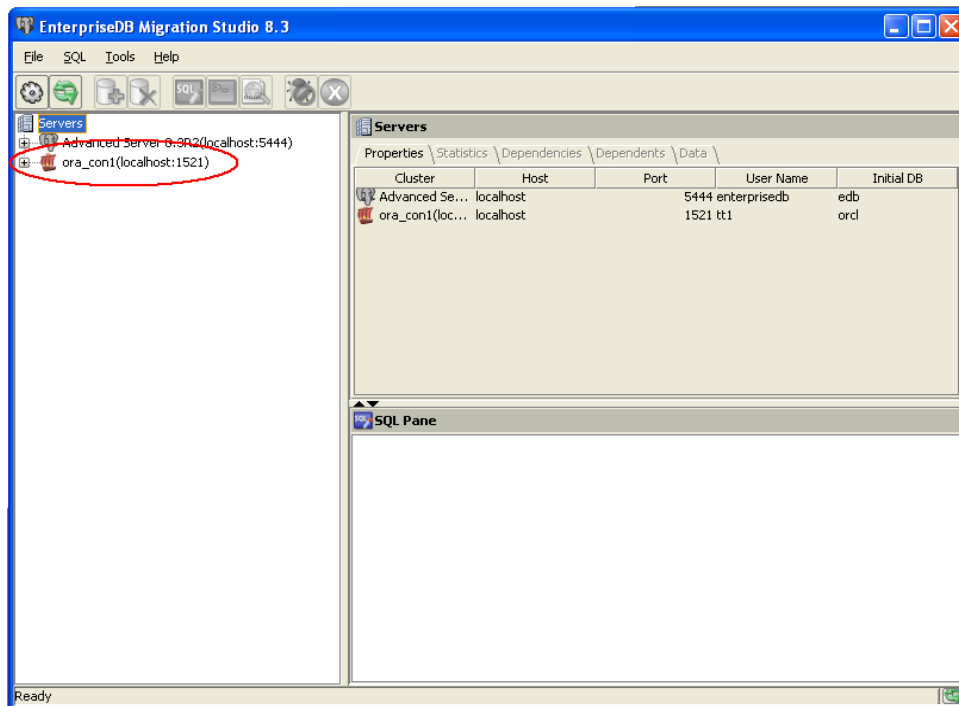


Figure 12 – EnterpriseDB migration studio 8.3, new server has been added

2.4 Online Migration

The screenshots below show how the online migration steps are accomplished.

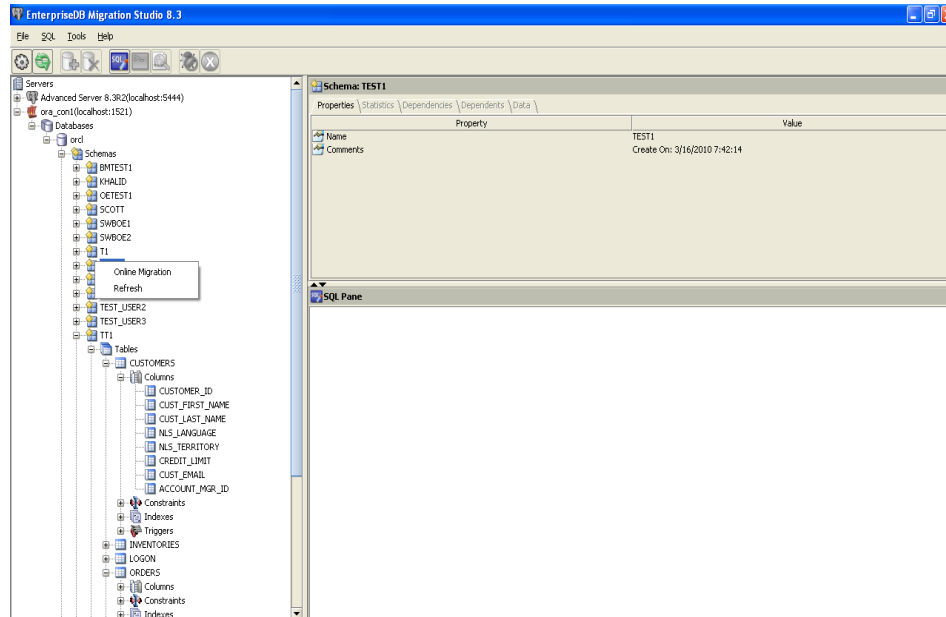


Figure 13 – Right click an object then click online migration

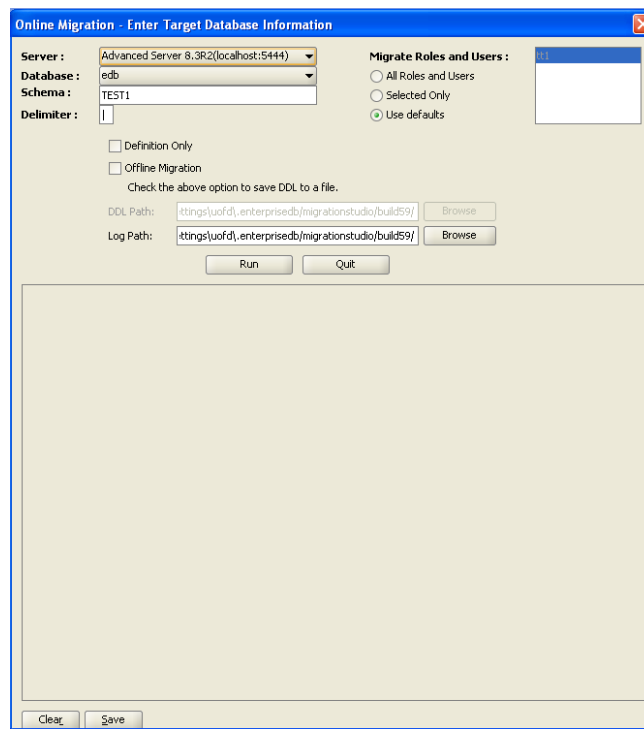


Figure 14 – Online migration dialog box

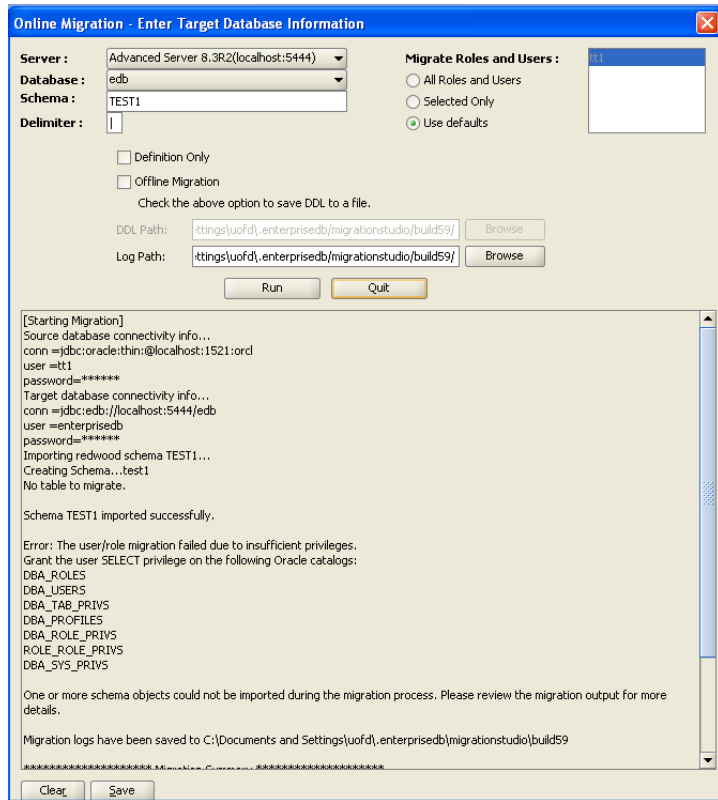


Figure 15 – After running online migration (whole T1 user)

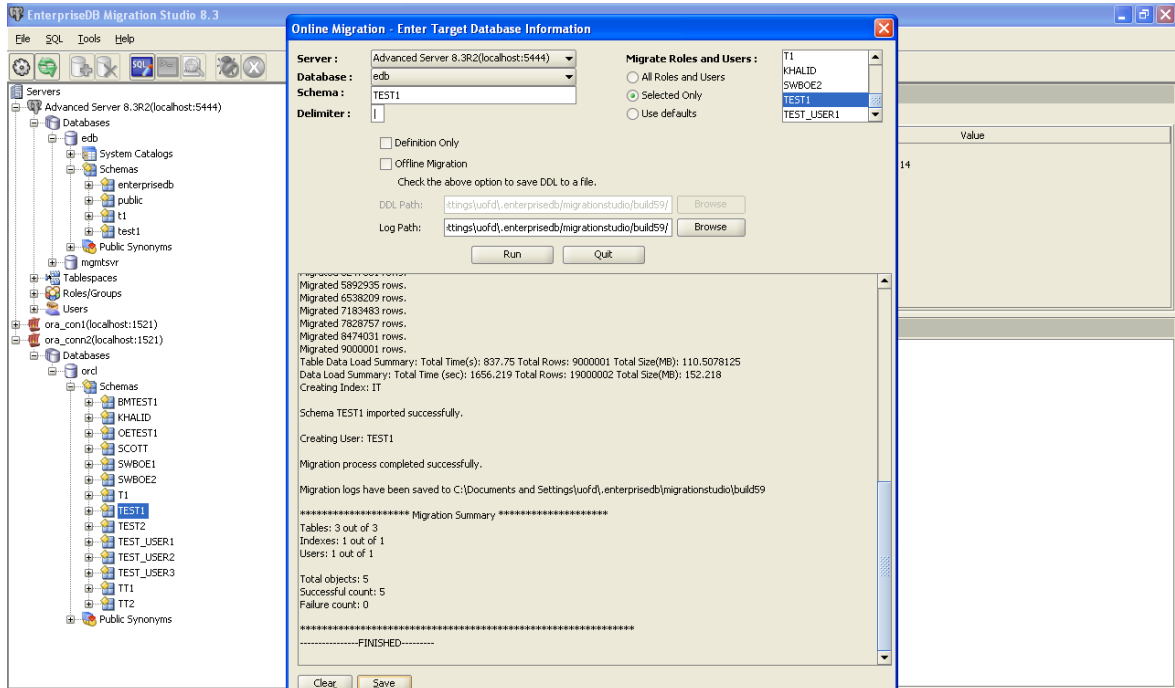


Figure 16 –User T1 has been successfully migrated

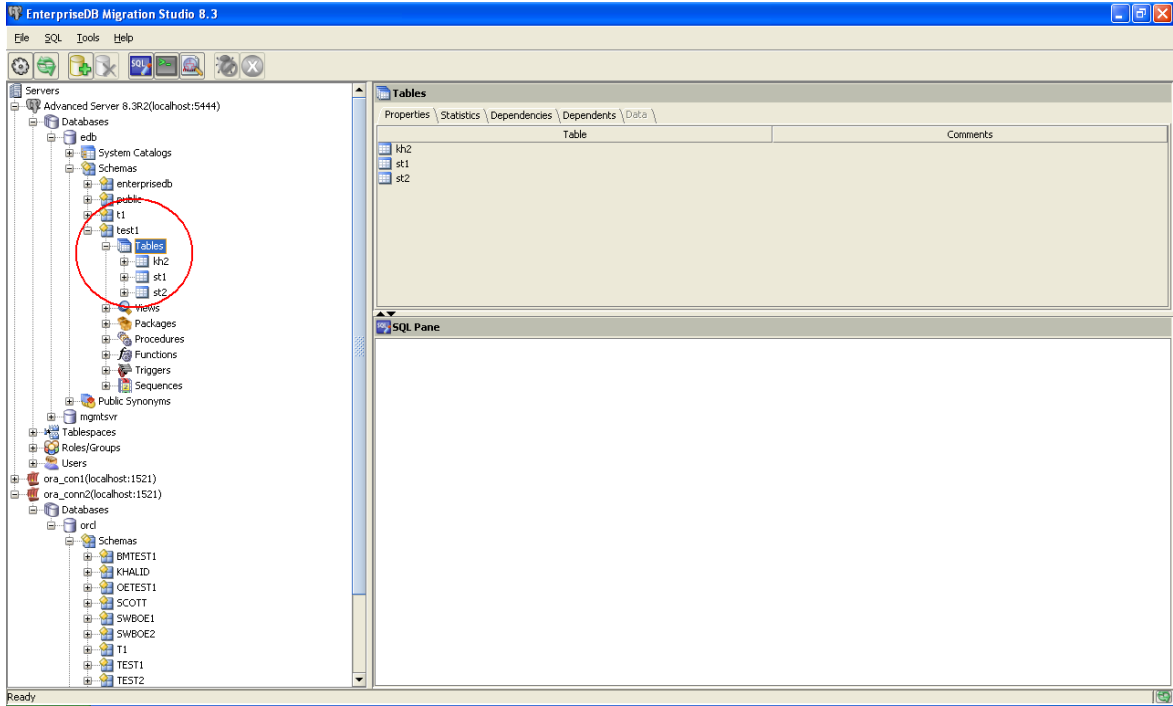


Figure 17 – The objects of user T1 could be appeared into Postgres Advanced Server

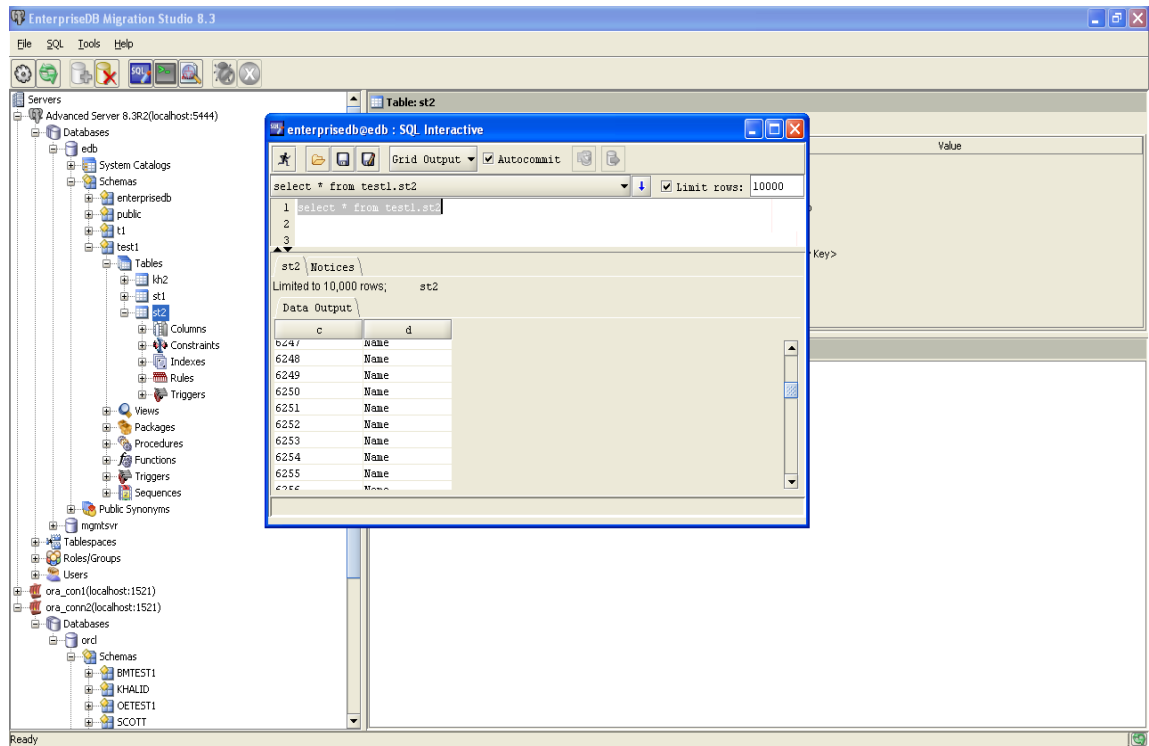


Figure 18 – Manipulation of user T1 in Postgres Advanced Server

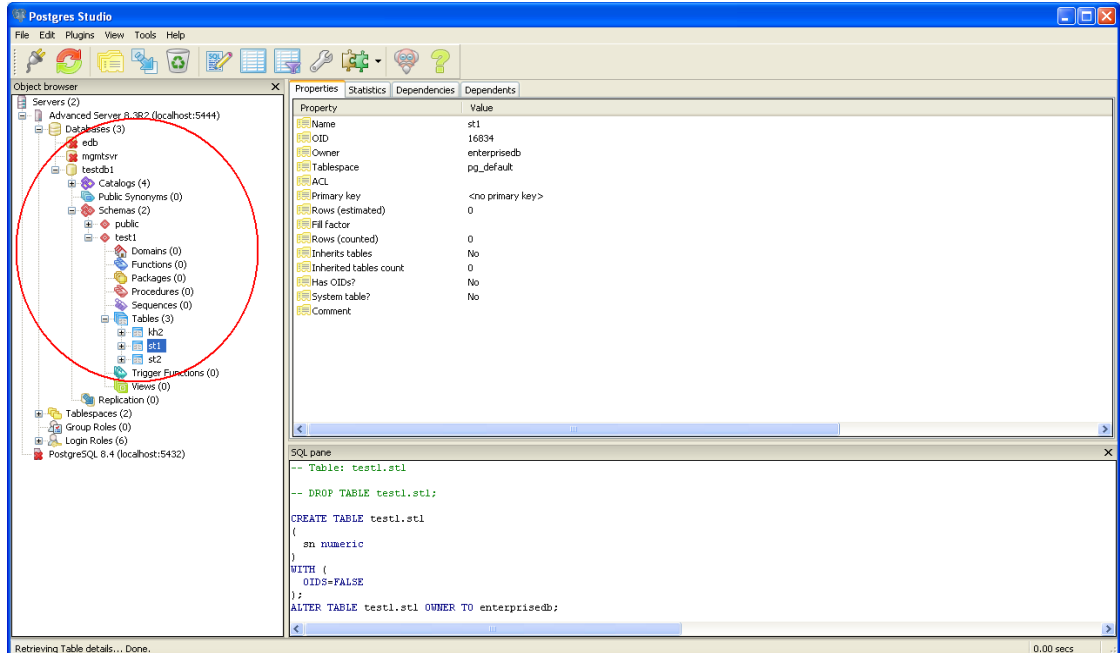


Figure 19 –Migrated user T1 in Postgres environment

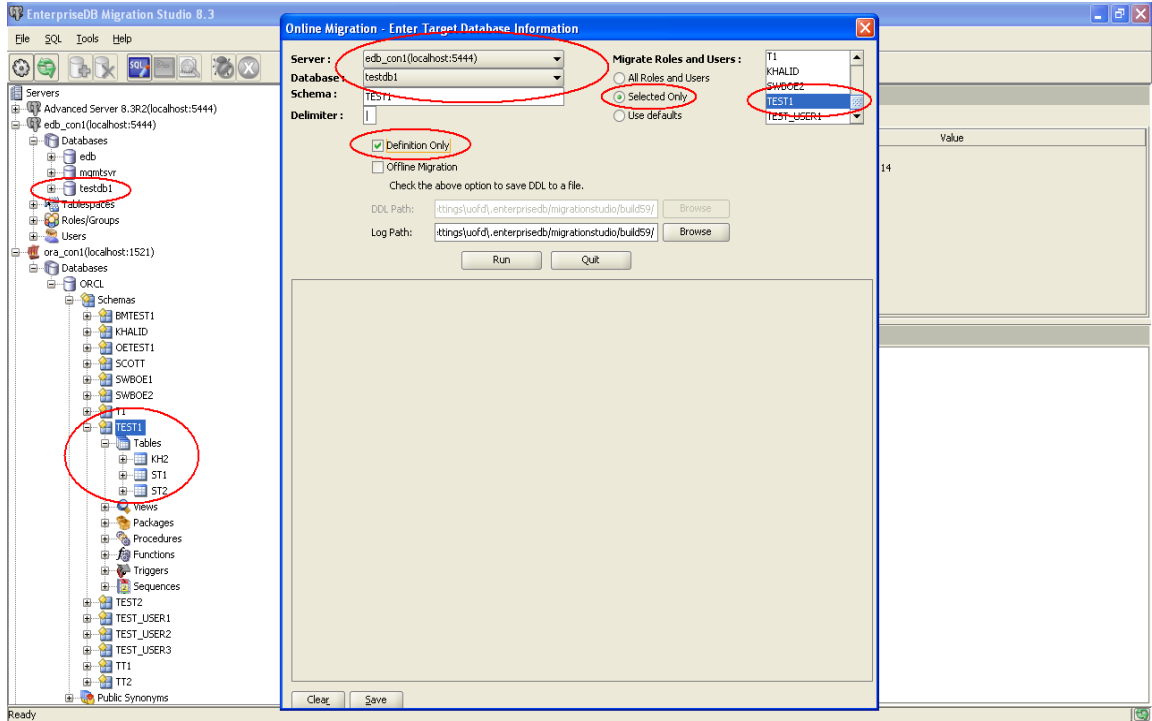


Figure 20 – Online migration, migrate definition only

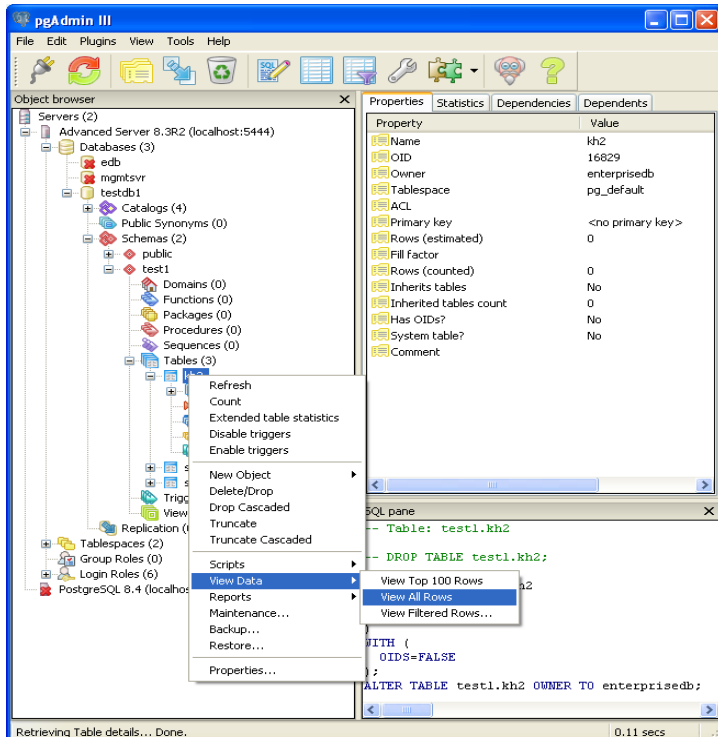


Figure 21 – Manipulation of migrated user in pgAdmin (view data)

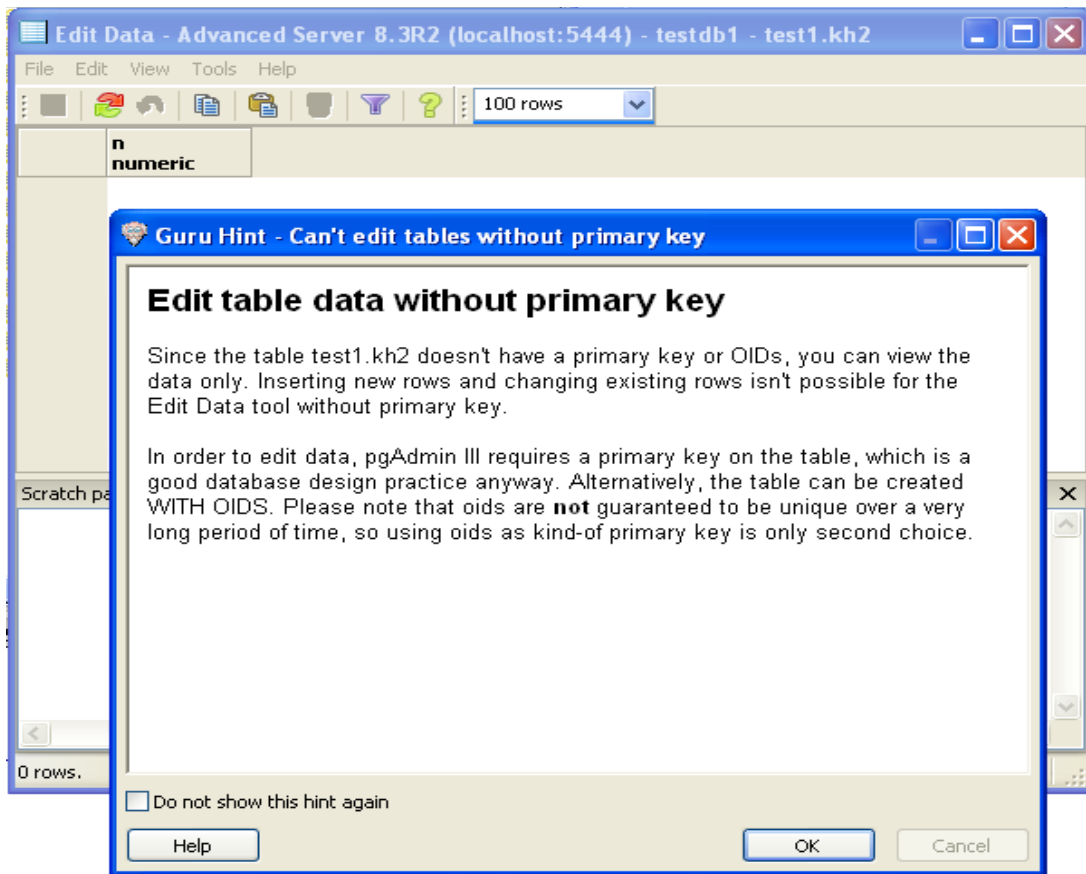


Figure 22 – Inability to view the data of migrated user, primary key recommendation

2.5 Some Conflicted Features in Postgres

ASM is the ability of the database to act as its own volume management system. Postgres doesn't have this and relies on the volume management system of the hardware for this functionality.

Table compression can be accomplished by putting those tables in a tablespace that resides on a compressed file system. This is supported by Postgres.

Advanced Queuing, although not built into the database can be and has been worked around using external messaging systems such as ActiveMQ, TIBCO® or MQ Series so those are options for you if you use AQ.

External tables will need to be loaded into staging tables in the database.

Using EnterpriseDB's EDB*Loader tool to load the data quickly is the recommended approach.

Materialized views exist in the form of summary tables that can be maintained by triggers, similar to the Oracle implementation but the setup is manual and not Oracle compatible. Automatic query rewrite is not currently supported so your application will need to be made aware of the summary table's existence.

Analytic Functions like lag, lead and dense_rank are not supported in the 8.3x releases of PostgresPlus Advanced Server. These have been introduced in the PostgreSQL 8.4 release so they will be available in the 8.4 release of Postgres Plus Advanced Server.

If errors are encountered during migration, EnterpriseDB can be quickly notified by emailing them at migrations@enterprisedb.com. They are continuously enhancing their compatibility and input from customers is invaluable in helping us prioritize our development plans and helping customers troubleshoot new issues.

APPENDIX C

Oracle AWR

1 AWR Steps

This section illustrates the steps necessary for [Creating Snapshots](#), [Dropping Snapshots](#), and Modifying Snapshot Settings.

1.1 Creating Snapshots

Snapshots can manually create with the `CREATE_SNAPSHOT` procedure to capture statistics at times different than those of the automatically generated snapshots. For example:

```
BEGIN
    DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT ();
END;
/
```

In this example, a snapshot for the instance is created immediately with the flush level specified to the default flush level of `TYPICAL`. You can view this snapshot in the `DBA_HIST_SNAPSHOT` view.

Steps in the sample database:

1. Connect to Oracle as system administrator and grant DBA to the two test users

```
CONN sys/oracle AS SYSDBA
```

```
GRANT DBA TO user_test1;
```

```
GRANT DBA TO user_test2;
```

2. Extract number of rows from the tested table, primarily

```
SELECT NUM_ROWS FROM DBA_TABLES WHERE TABLE_NAME LIKE  
'U1T1';
```

3. Let the analyzer to analyze the table, by issuing the statement

```
ANALYZE TABLE TEST_USER1.U1T1 ESTIMATE STATISTICS;
```

```
ANALYZE TABLE TEST_USER1.U1T1 ESTIMATE STATISTICS FOR  
ALL COLUMNS;
```

Or better use when table is small

```
ANALYZE TABLE TEST_USER1.U1T2 COMPUTE STATISTICS;
```

```
ANALYZE TABLE TEST_USER1.U1T2 COMPUTE STATISTICS FOR  
ALL COLUMNS;
```

4. Again, extract the number of records by issuing the following statement

```
SELECT NUM_ROWS FROM DBA_TABLES WHERE TABLE_NAME LIKE  
'U1T1';
```

5. Perform high insert activity

```
CONN AS USER_TEST2
```

```
CREATE TABLE U2T1(C NUMBER);
```

```
BEGIN
```

```
FOR I IN 1..10000000 LOOP
```

```
INSERT INTO U2T1 SELECT SID FROM V$MYSTAT WHERE  
ROWNUM<2;
```

```
COMMIT;
```

```
DELETE U2T1 WHERE I IN (SELECT SID FROM V$MYSTAT  
WHERE ROWNUM<2);
```

```

COMMIT;

END LOOP;

END;

/

OR:

CREATE TABLE U2T2 (N NUMBER);

BEGIN

FOR I IN 1..10000000 LOOP

INSERT INTO U2T2 SELECT SID FROM V$MYSTAT WHERE
ROWNUM<2;

COMMIT;

END LOOP;

END;

/

EXIT;

CONN AS USER_TEST1

CREATE TABLE U1T3 (I NUMBER);

DROP TABLE U1T1 CASCADE CONSTRAINTS;

CREATE TABLE U1T2(C NUMBER (10), D VARCHAR2(20));

CREATE INDEX IT ON U1T2(C);

INSERT INTO U1T2 VALUES (1);

COMMIT;

BEGIN

FOR I IN 1..900000 LOOP

```

```
INSERT INTO U1T2 VALUES (I, 'INITIAL');  
  
END LOOP;  
  
END;  
  
/
```

6. Simulate high CPU

```
CONN TEST_USER2  
  
SELECT * FROM V$SESSION A, V$SESSION B, V$SESSION C;
```

1.2 AWR Steps and Results

1. Creating Snapshots:

```
CONN SYS/KHALID AS SYSDBA  
  
BEGIN  
  
    DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT ();  
  
END;  
  
/
```

2. Dropping Snapshots:

```
SELECT DBID FROM V$DATABASE;
```

Then use the DBID in the next step

```
BEGIN  
  
DBMS_WORKLOAD_REPOSITORY.DROP_SNAPSHOT_RANGE  
(LOW_SNAP_ID => 899, HIGH_SNAP_ID => 901, DBID =>  
3492516897);  
  
END;
```


/

3. Running Workload Repository Reports Using SQL Scripts:

You can view AWR reports by running one of the SQL scripts such as the `awrrpt.sql` SQL script generates an HTML or text report that displays statistics for a range of snapshot Ids.

Then choose the `SQL_ID` after that run the following script:

```
@?/RDBMS/ADMIN/AWRSQRPT.SQL
```

Parse Calls:

```
SET LINESIZE 100  
SET PAGESIZE 100  
SELECT * FROM  
(SELECT SUBSTR (SQL_TEXT, 1, 40) SQL,  
PARSE_CALLS, EXECUTIONS, HASH_VALUE, ADDRESS  
FROM V$SQLAREA  
WHERE PARSE_CALLS > 1000  
ORDER BY PARSE_CALLS DESC)  
WHERE ROWNUM <= 10;
```

Sharable Memory:

```
SELECT * FROM  
(SELECT SUBSTR (SQL_TEXT, 1, 40) SQL,  
SHARABLE_MEM, EXECUTIONS, HASH_VALUE, ADDRESS  
FROM V$SQLAREA  
WHERE SHARABLE_MEM > 1048576
```

```
ORDER BY SHARABLE_MEM DESC)
```

```
WHERE ROWNUM <= 10;
```

Version Count:

```
SELECT * FROM
```

```
(SELECT SUBSTR (SQL_TEXT, 1, 40) SQL, VERSION_COUNT,  
EXECUTIONS, HASH_VALUE, ADDRESS FROM V$SQLAREA
```

```
WHERE VERSION_COUNT > 20
```

```
ORDER BY VERSION_COUNT DESC)
```

```
WHERE ROWNUM <= 10;
```

Buffer Gets:

```
SELECT * FROM
```

```
(SELECT SUBSTR (SQL_TEXT, 1, 40) SQL,
```

```
BUFFER_GETS, EXECUTIONS, BUFFER_GETS/EXECUTIONS  
"GETS/EXEC",
```

```
HASH_VALUE, ADDRESS
```

```
FROM V$SQLAREA
```

```
WHERE BUFFER_GETS > 10000
```

```
ORDER BY BUFFER_GETS DESC)
```

```
WHERE ROWNUM <= 10;
```

1.3 Dropping Snapshots

You can drop a range of snapshots using the `DROP_SNAPSHOT_RANGE` procedure. To view a list of the snapshot IDs along with database IDs, check the `DBA_HIST_SNAPSHOT` view. For example, you can drop the following range of snapshots:

```
BEGIN
```

```
DBMS_WORKLOAD_REPOSITORY.DROP_SNAPSHOT_RANGE (LOW_SNAP_ID
=>22, HIGH_SNAP_ID => 32, DBID => 3310949047);

END;

/
```

In the above example, the range of snapshot Ids to drop is specified from 22 to 32. The optional database identifier is 3310949047. If you do not specify a value for DBID, the local database identifier is used as the default value.

ASH that belongs to the time period specified by the snapshot range is also purged when the DROP_SNAPSHOT_RANGE procedure is called.

1.4 AWR Reports

You can view the AWR reports with Oracle Enterprise Manager or by running SQL scripts. To run an AWR report, a user must be granted the DBA role.

The reports are divided into multiple sections. The HTML report includes links that can be used to navigate quickly between sections. The content of the report contains the workload profile of the system for the selected range of snapshots.

Note:

If you run a report on a database that does not have any workload activity during the specified range of snapshots, calculated percentages for some report statistics can be less than 0 or greater than 100. This result simply means that there is no meaningful value for the statistic.

1.5 Running AWR Using SQL Scripts

AWR reports can view by running the following SQL script:

- The `awrrpt.sql` SQL script generates an HTML or text report that displays statistics for a range of snapshot Ids.

Running the awrrpt.sql report as follow:

To generate an HTML or text report for a range of snapshot Ids, run the `awrrpt.sql` script at the SQL prompt:

```
@$ORACLE_HOME/RDBMS/ADMIN/AWRRPT.SQL
```

First, you need to specify whether you want an HTML or a text report.

```
ENTER VALUE FOR REPORT_TYPE: TEXT
```

Specify the number of days for which you want to list snapshot Ids.

```
ENTER VALUE FOR NUM_DAYS: 2
```

After the list displays, you are prompted for the beginning and ending snapshot Id for the workload repository report.

```
ENTER VALUE FOR BEGIN_SNAP: 150
```

```
ENTER VALUE FOR END_SNAP: 160
```

Next, accept the default report name or enter a report name. The default name is accepted in the following example:

```
ENTER VALUE FOR REPORT_NAME:
```

```
USING THE REPORT NAME AWRRPT_1_150_160
```

The workload repository report is generated.

2 Managing Snapshot Data with APIs

While the primary interface for managing the AWR is the Oracle Enterprise Manager Database Control, monitoring functions can be managed with procedures in the `DBMS_WORKLOAD_REPOSITORY` package.

Snapshots are automatically generated for an Oracle database; however, you can use `DBMS_WORKLOAD_REPOSITORY` procedures to manually create, drop, and modify the snapshots that are used by automatic database diagnostic monitoring. Snapshots are sets of historical data for specific time periods that are used for performance comparisons. [4]

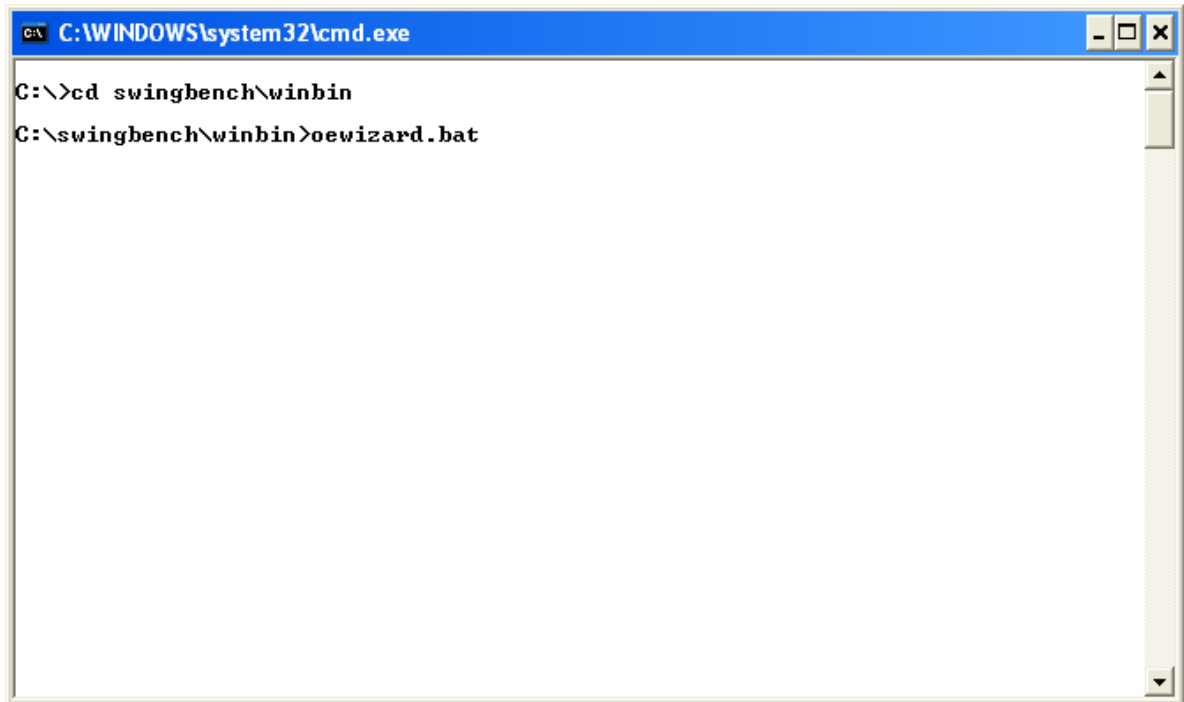
To invoke these procedures, a user must be granted the DBA role.

The steps necessary for [Creating Snapshots](#), dropping Snapshots, and Modifying Snapshot Settings are placed in appendix C.

APPENDIX D

Oracle Benchmark Test (Swingbench)

The following screen-shots illustrate how swingbench works.



```
C:\WINDOWS\system32\cmd.exe
C:\>cd swingbench\winbin
C:\swingbench\winbin>oewizard.bat
```

Figure 1 – Enter swingbench directory and open oewizard

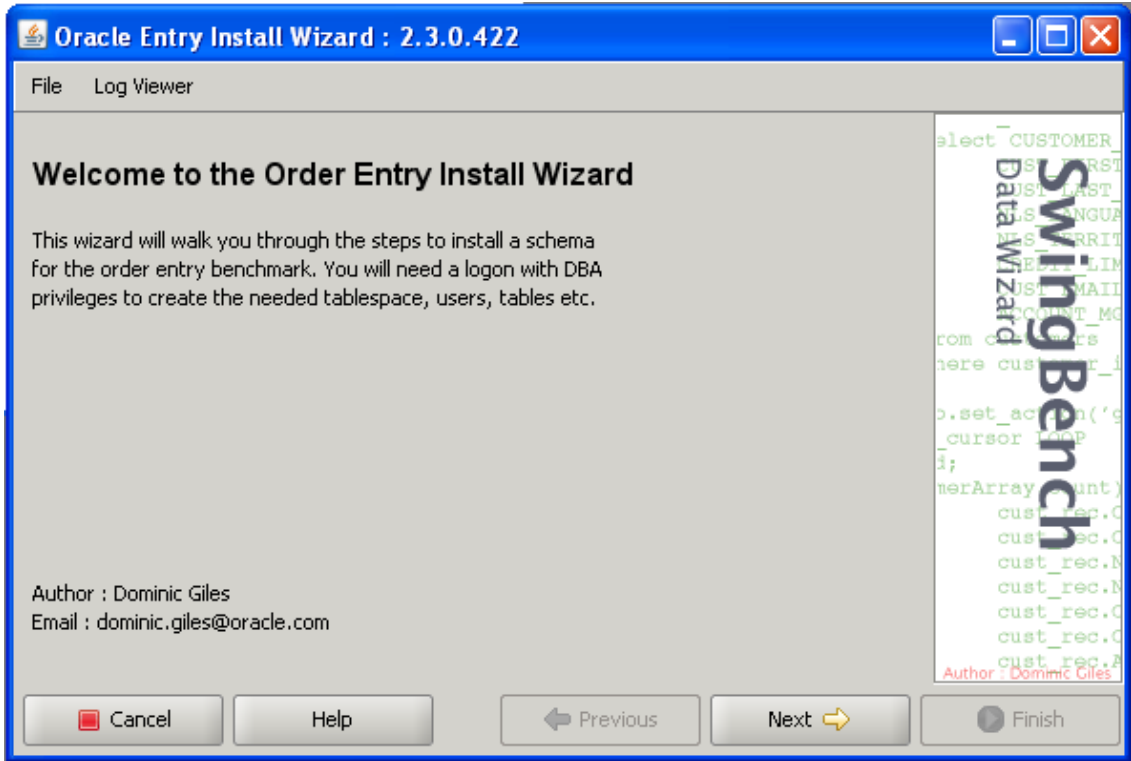


Figure 2 – oewizard, welcome screen




Figure 3 – Oewizard, task selection

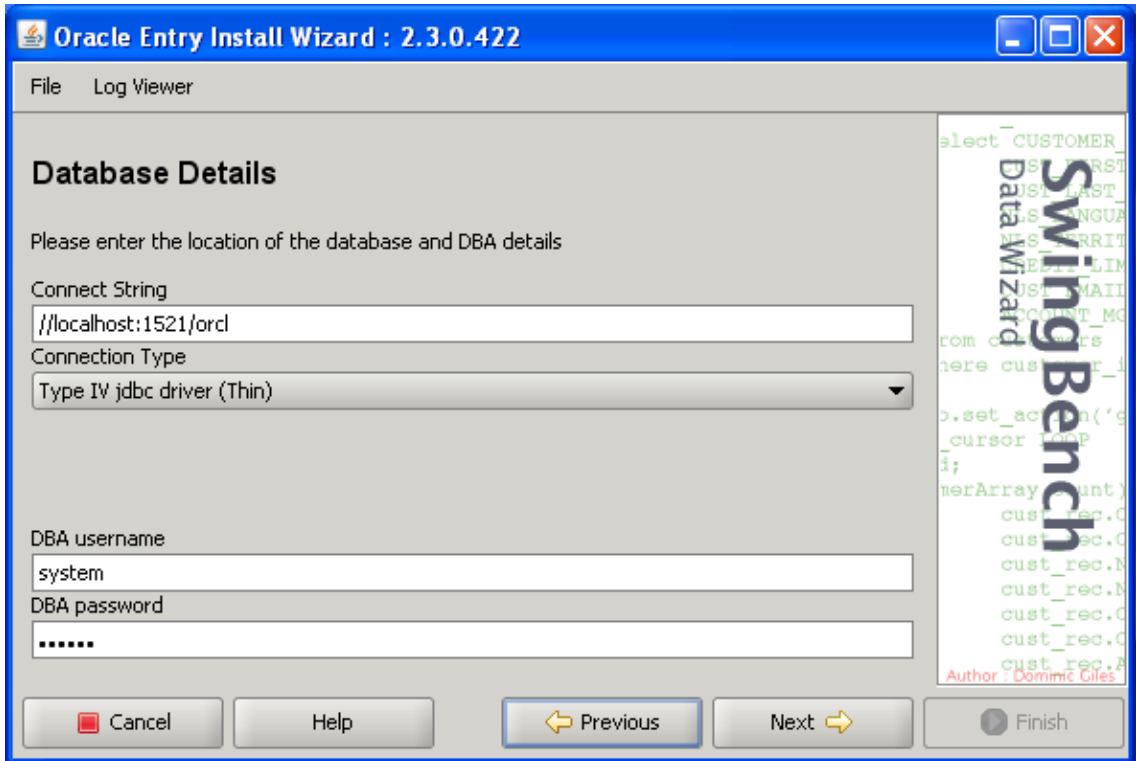


Figure 4 – Oewizard, DBA details

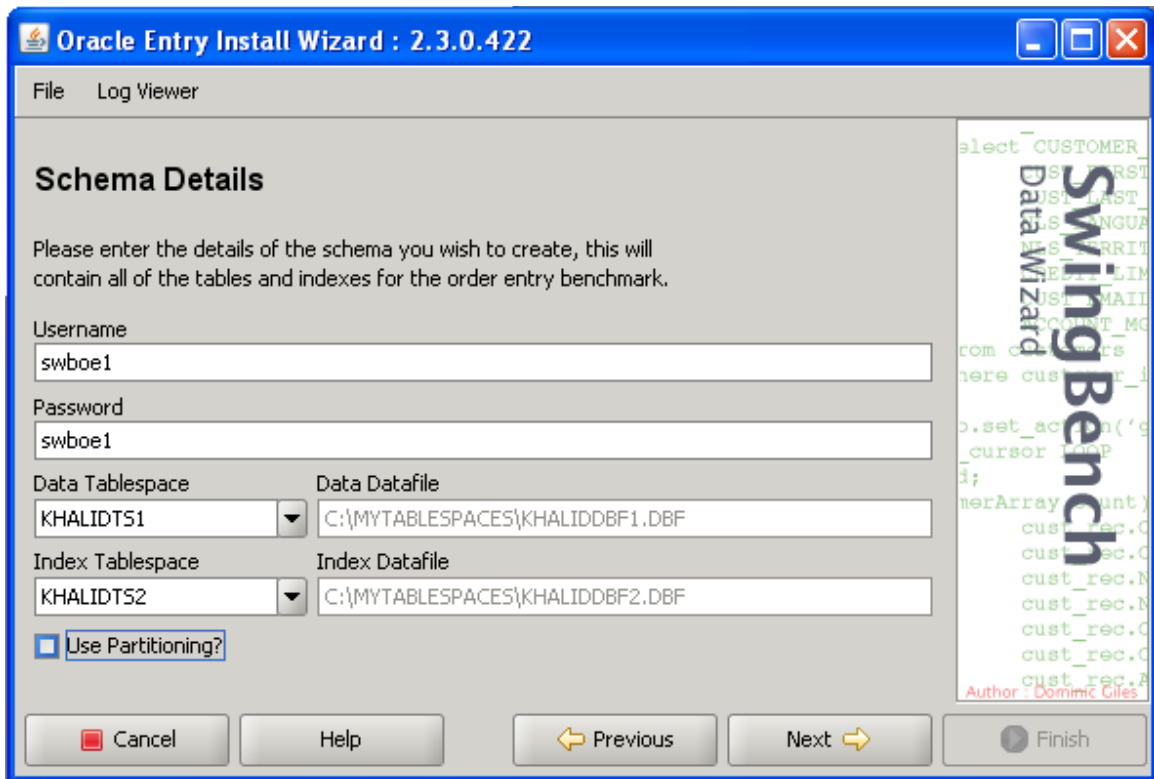


Figure 5 – Oewizard, details of schema to be created




Figure 6 – Oewizard, sizing details

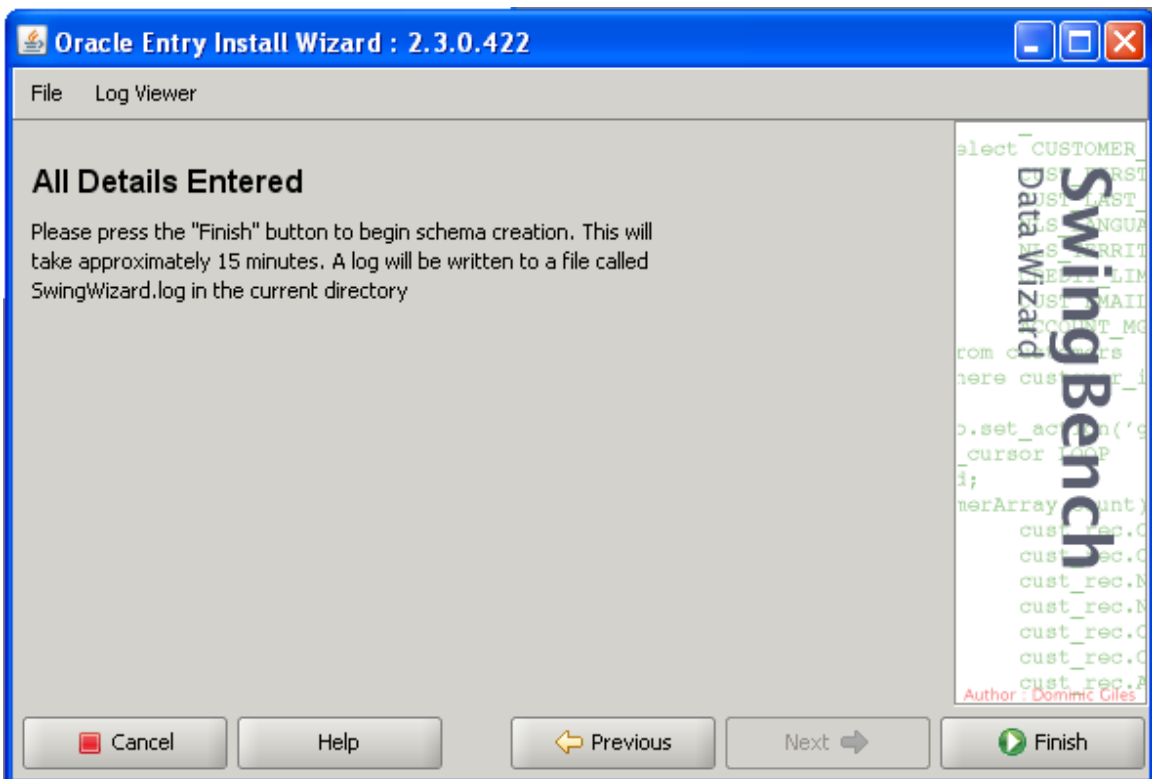


Figure 7 – Oewizard, ready to create schema

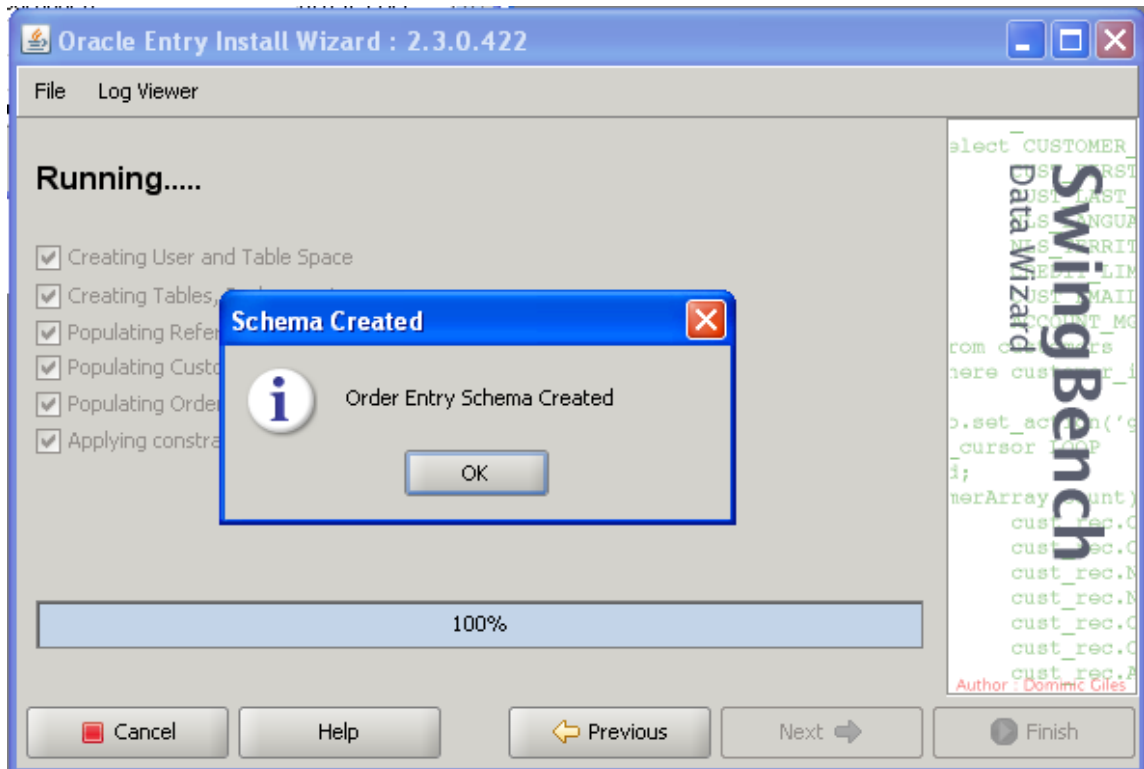


Figure 8 – Oewizard, oe schema created




Figure 9 – Oewizard, oe schema created




Figure 10 – Oewizard, saving configuration

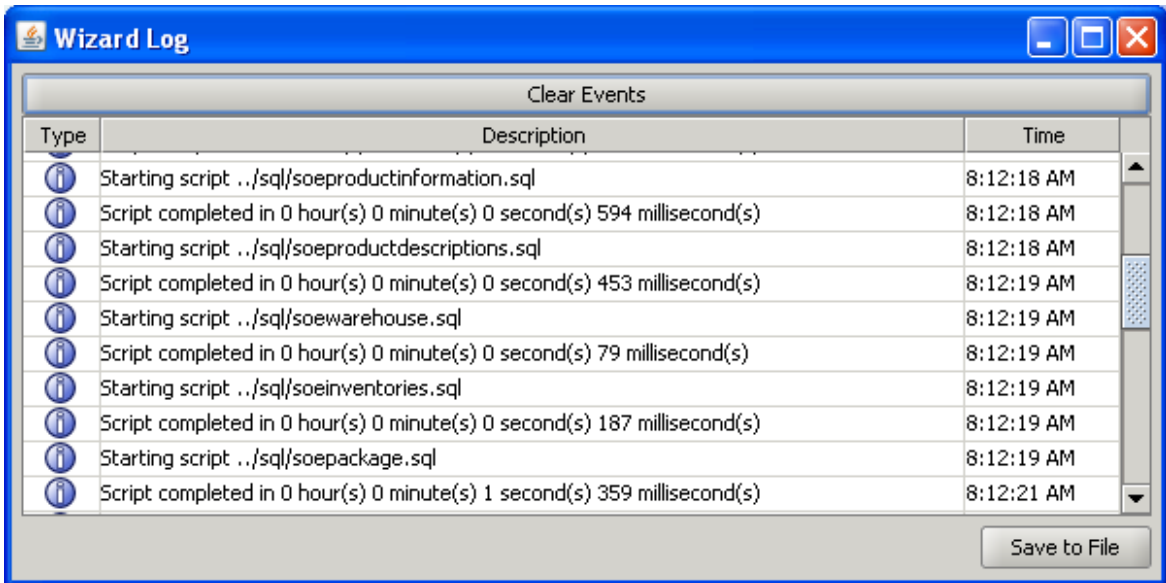


Figure 11 – Oewizard, events log




Figure 12 – Oewizard, exit




Figure 13 – Start swingbench

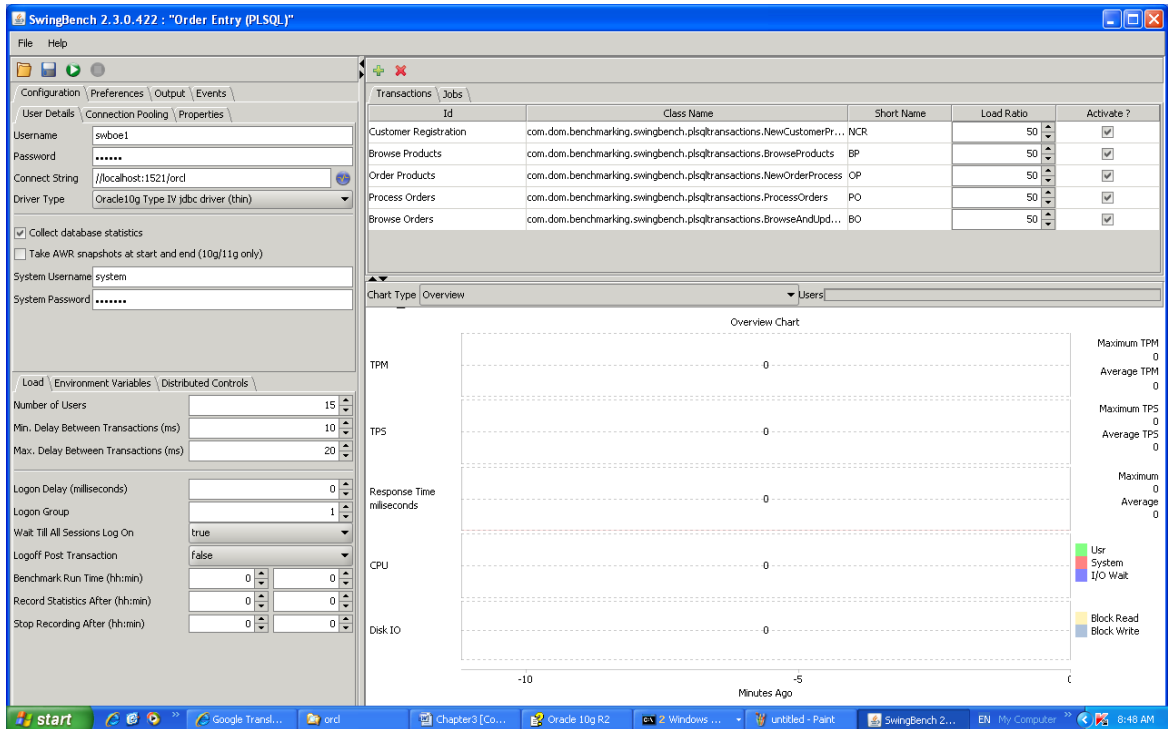


Figure 14 – Swingbench, main menu

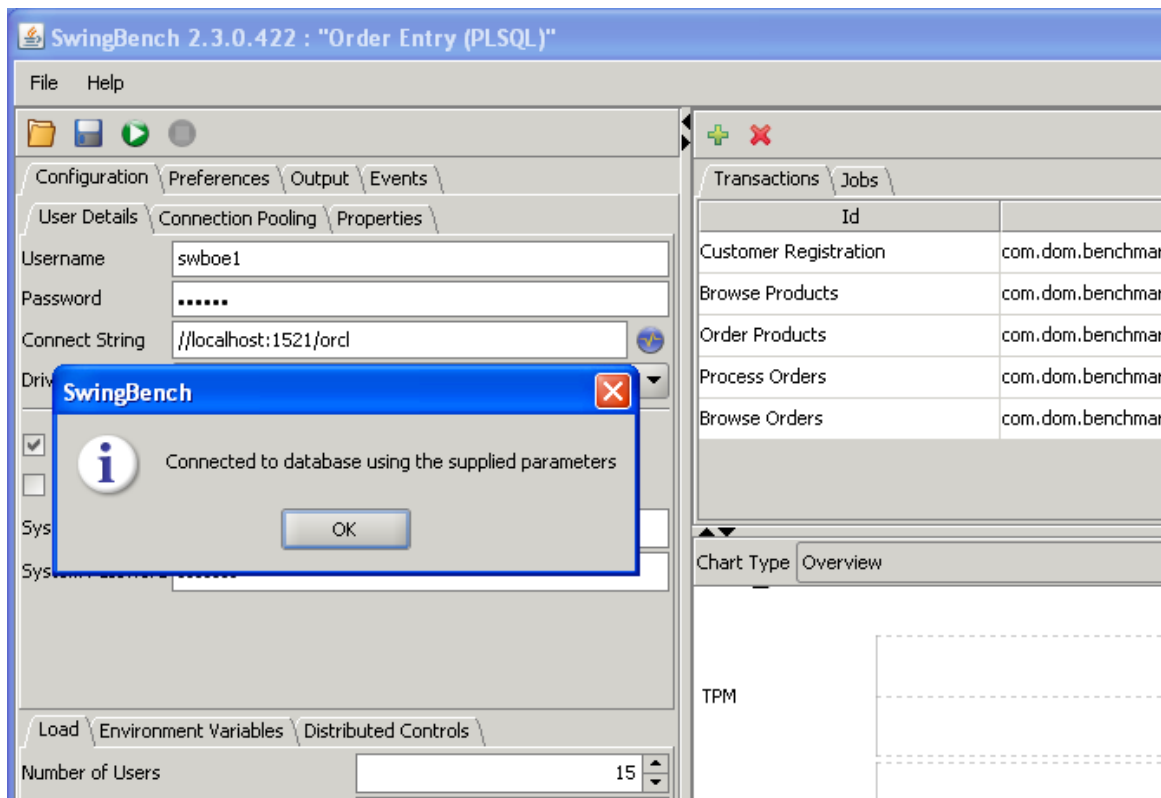


Figure 15 – Swingbench, testing the connection to database