IBM

BP05: IBM DB2 Migration Toolkit for iSeries

Migrating from Oracle to DB2 UDB for iSeries

Jarek Miszczyk, PartnerWorld for Developers, Rochester
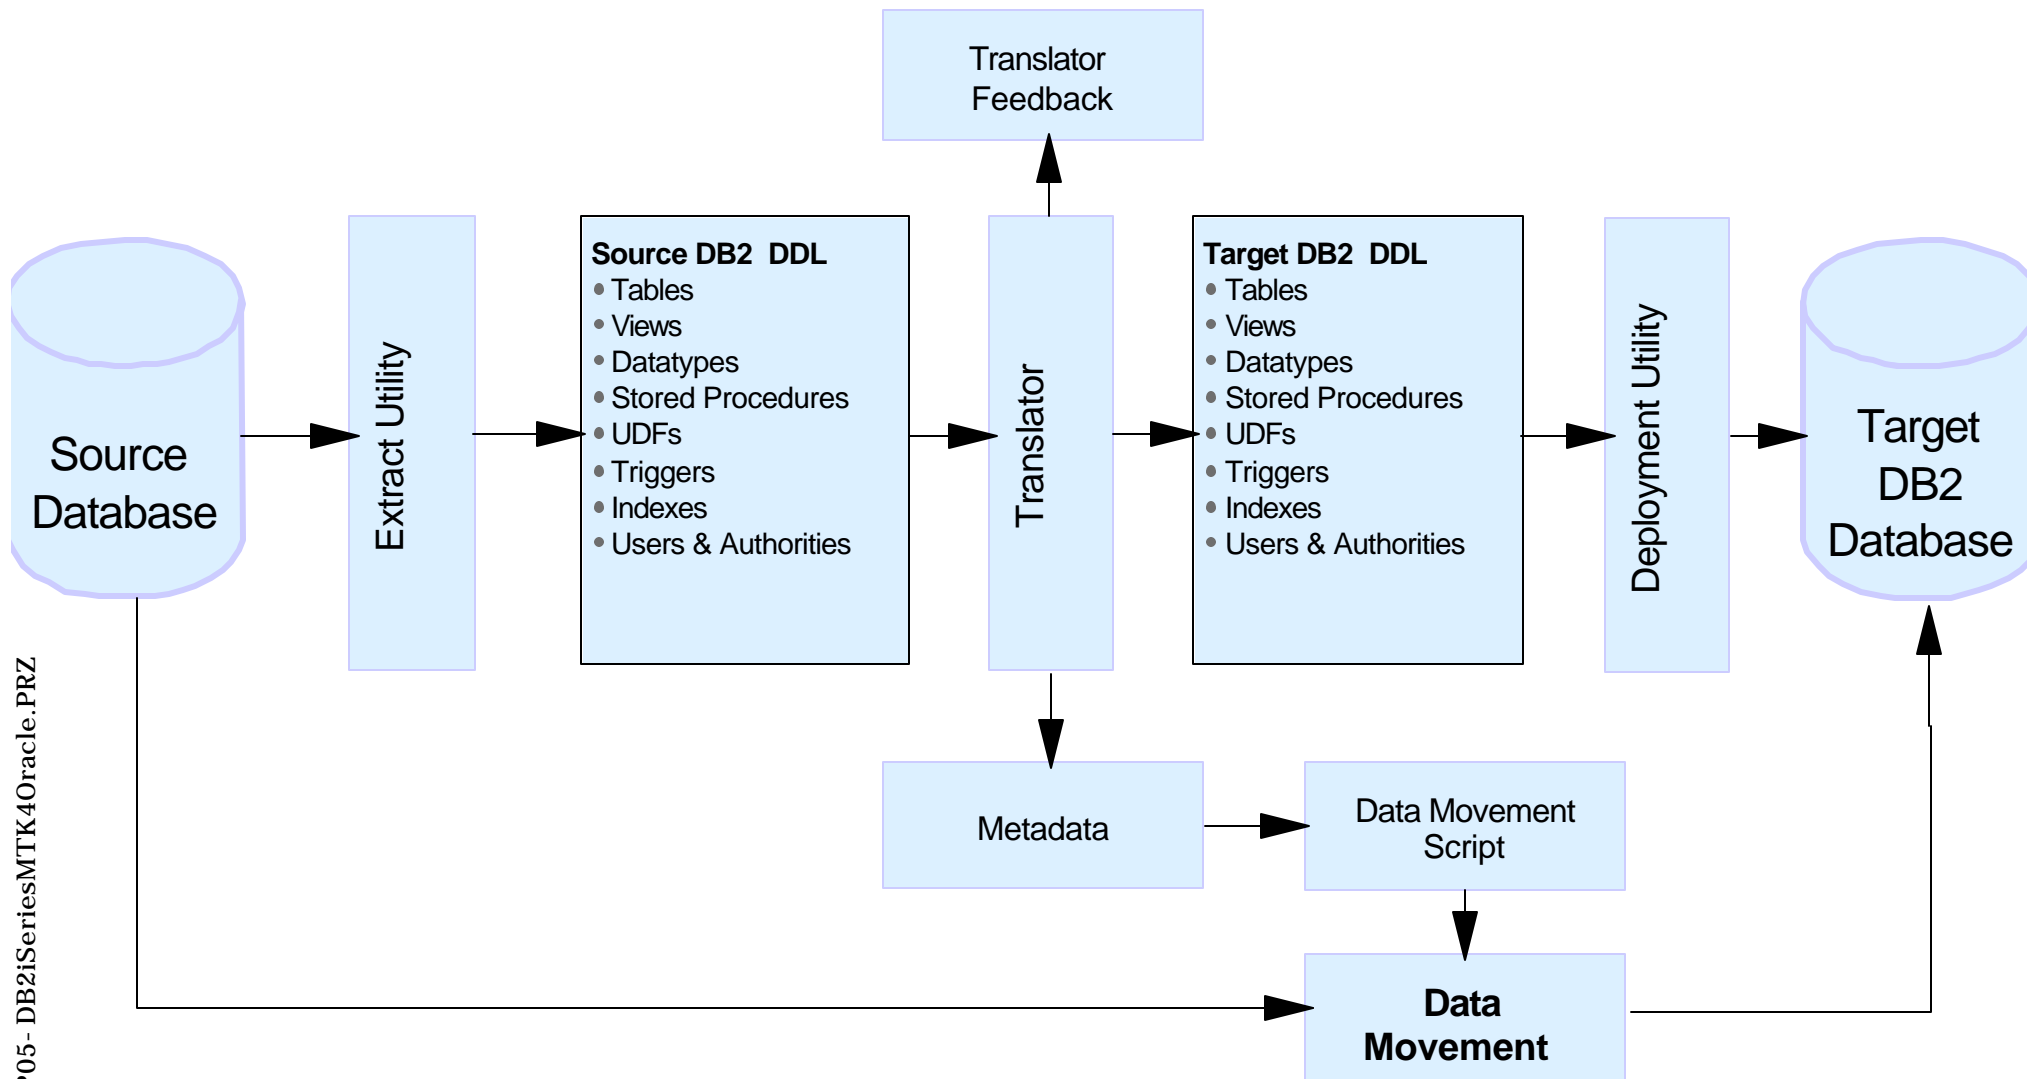
F03BP05- DB2iSeriesMTK4Oracle.PRZ

IBM DB2 Migration Toolkit for iSeries

- Provides single development environment for migration from Oracle 8
  - ▶ supports DB2 UDB for iSeries V5R2 and DB2 UDB EE/EEE V7 and V8
  - ▶ requires the latest DB2 UDB for iSeries database FixPak SF99502
  - ▶ additionally SI06748 and SI06675 PTFs must also be loaded
- Runs on Windows NT or Windows 2000
  - ▶ requires an ODBC or JDBC connection to source and target databases
  - ▶ available in English only
  - ▶ Linux, AIX, and Solaris in the future
- Provides the functions to graphically create, build, and deploy the migrated solution
- Provides migration reports and logs

The Architecture of IBM DB2 Migration Toolkit

Translator
Feedback

Source Database

Extract Utility

Source DB2 DDL
- Tables
- Views
- Datatypes
- Stored Procedures
- UDFs
- Triggers
- Indexes
- Users & Authorities

Translator

Target DB2 DDL
- Tables
- Views
- Datatypes
- Stored Procedures
- UDFs
- Triggers
- Indexes
- Users & Authorities

Deployment Utility

Target DB2 Database

Metadata

Data Movement Script

Data Movement
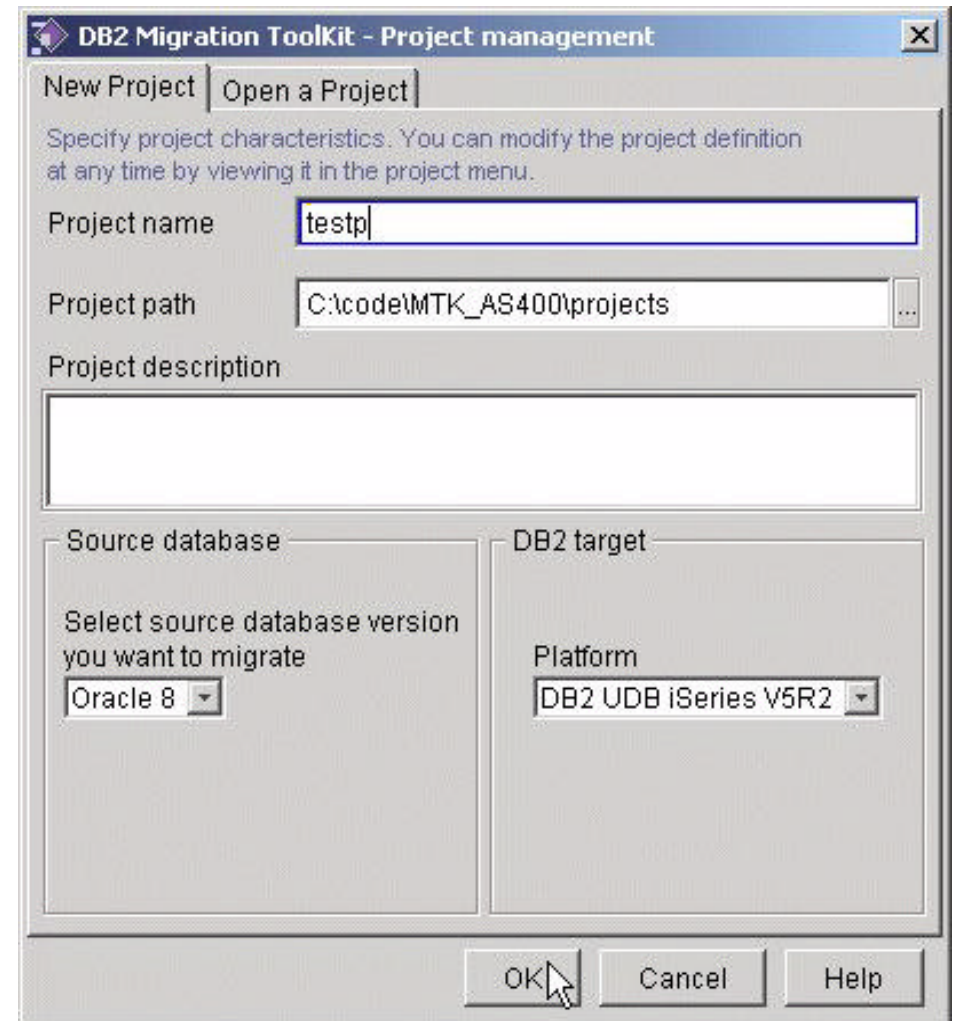
The DB2 MTK Basics

- MTK manages work by projects
- Migration process implemented in a sequence of steps
  - ► Import/extract metadata
    - – The toolkit converts the source metadata into DDL scripts
  - ► Convert metadata
    - – SQL Translator converts the source DDL scripts to DB2 constructs
    - – metadata for each source and target database object is also generated
  - ► Refine metadata
    - – helps rectify the inconsistencies between source and target scripts
  - ► Generate data transfer scripts
    - – Data can be exported from Oracle and loaded into DB2
    - – On the iSeries, MTK utilizes the enhanced CPYFRMIMPF command
    - – CPYFRMIMPF supports LOB and identity columns
  - ► Deploy to DB2
    - – deployment scripts executed on target iSeries

What gets converted?

- MTK converts a range of Oracle constructs
  - ► SQL DDL (create table, create index, create view, alter table...)
  - ► SQL DML Statements (insert, update, delete,...)
  - ► SQL Queries (select)
  - ► Triggers
  - ► Procedures
  - ► Functions
- Oracle objects not supported yet
  - ► Applications (OCI, embedded SQL)
  - ► Replication
  - ► OLAP specific features
  - ► Objects relational features & collections
  - ► Catalog/system tables
  - ► Statement specific to system administration
- For a list of supported features refer to "Summary of features"
 - ► can be found in Converter Reference of the online help

Setting up a new project

- Project name
- Project path on the local workstation
- Source Oracle database
  - ► currently Oracle 8 is supported
- Target DB2 database
  - ► DB2 UDB for iSeries V5R2
  - ► DB2 UDB EE/EEE V7 and V8

DB2 Migration ToolKit – Project management

New Project | Open a Project

Specify project characteristics. You can modify the project definition at any time by viewing it in the project menu.

Project name: `testp`

Project path: `C:\code\MTK_AS400\projects`

Project description:

Source database

Select source database version you want to migrate

`Oracle 8`

DB2 target

Platform

`DB2 UDB iSeries V5R2`

OK   Cancel   Help

Specifying the Source

- Two possible sources of Oracle metadata
  - ► directly extracted from Oracle database
  - ► imported from existing source file
- Direct extraction requires a connection to Ora
  - ► ODBC
    - – DSN entry needs to exist prior to conversion
  - ► JDBC
    - – classes111.zip required in the CLASSPATH global variable
      - • usually located in [ORACLE_HOME]\jdbc\lib
- Run stats on Oracle recommended prior to extraction
  - ► run statistics on SYS.DEPENDENCY$, SYS.OBJ$, SYS.USER$
  - ► can be accomplished through Oracle's admin tools
    - – DBA Studio
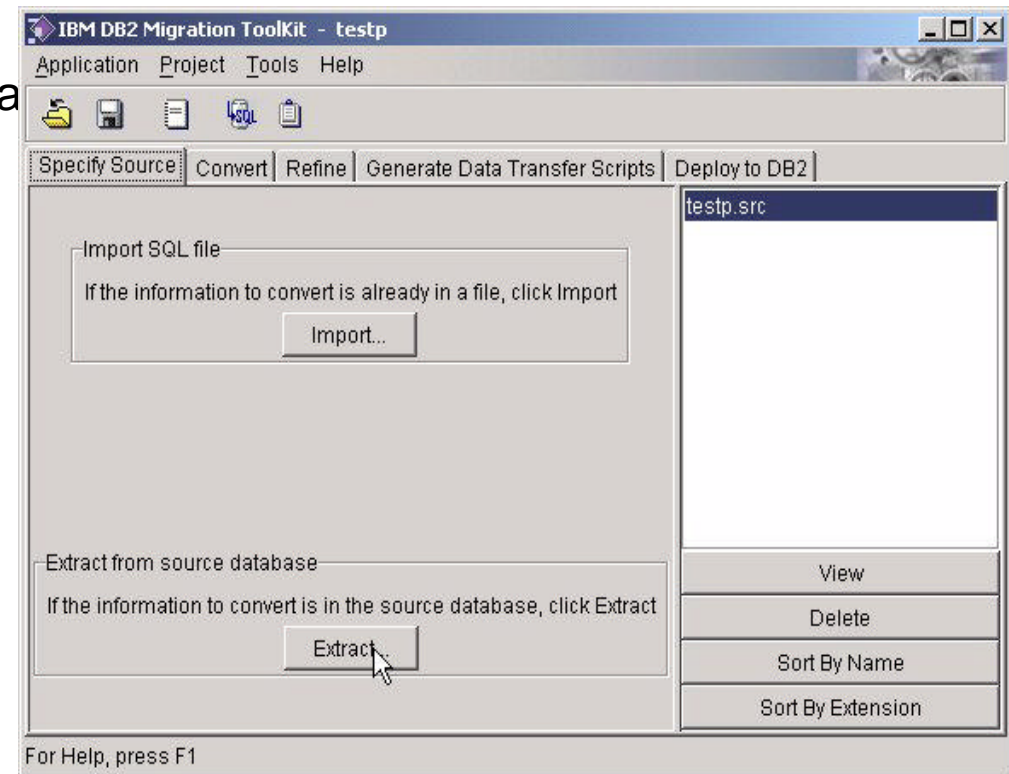
Extracting the source metadata

- The extraction is performed by the Extract utility
  - ► retrieves the metadata from Oracle's catalogs
  - ► generates native (Oracle) SQL statement for all selected objects
  - ► writes the SQL statements into a flat file on the local workstation
  - ► the flat file used as input to the Translator utility
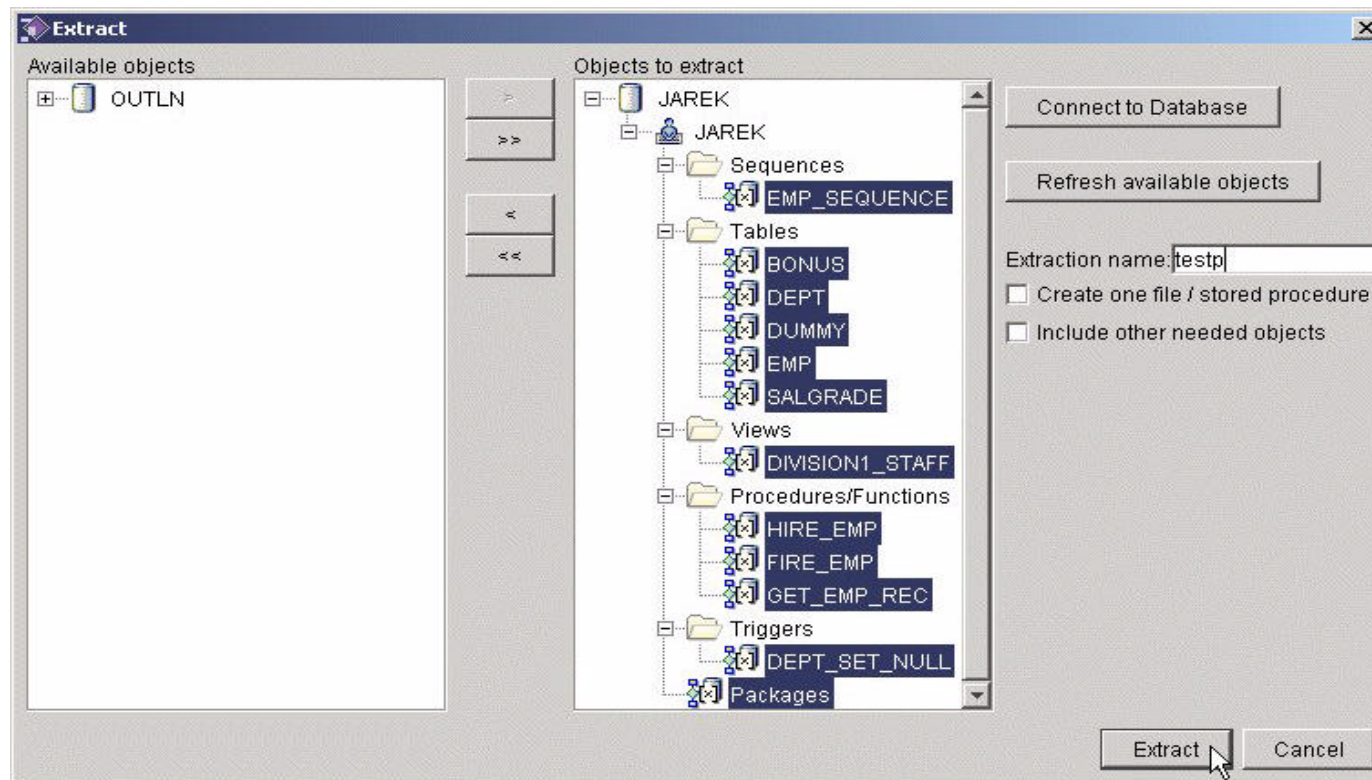
Converting to DB2

- The conversion process can be fine-tuned
  - ► Global Type Mapping function
    - – provides a list of default type mappings
    - – some mappings can be changed
    - – the current MTK version doesn't allow VARCHAR2 to CHAR remapping
      - this remapping can be accomplished in the Refine step
      - recommended for performance reasons
  - ► Advanced Options function
    - – controls the format of the output script file
      - insert DROP statements before CREATE
      - copy source as comments
      - do not copy full source for procedures
      - ...
- The translator is invoked by selecting Convert on the Convert dialog

Refining the conversion

- Once the conversion is accomplished, MTK switches to Refine dialog
- Refine process is critical for the quality of the migration project
  - ► requires fairly high level of expertise in both Oracle and DB2 UDB for iSeries
- Refine - Convert is an iterative process
  - ► Oracle source files are manually "tweaked" and then reconverted
  - ► The process repeats until the results are satisfactory or further improvement not possible

Refining cont.

- Analyze the translator feedback in the Messages dialog
- The messages are sorted by category
  - ► Input Error, Translator Information, Translator Limitation, Translator Omission
  - ► The order of message category reflects the relative importance
  - ► The upper levels require modification
  - ► The lower levels informational and require little or no action

Refine example - Changing Oracle source

- Oracle's %ROWTYPE and %TYPE not supported by the translator as input parameters of stored procedure or function
  - These constructs are supported in the body of procedure or function
- The affected Oracle procedure or function needs to be rewritten
  - %ROWTYPE converted to list of variables

Original code extracted from Oracle:

```
CREATE OR REPLACE PROCEDURE   Get_emp_rec (Emp_number IN Emp.Empno%TYPE,  Emp_ret OUT Emp%ROWTYPE) IS
BEGIN
SELECT Empno, Ename, Job, Mgr, Hiredate, Sal, Comm, Deptno INTO Emp_ret
FROM Emp WHERE Empno = Emp_number;
END;
```

Manually modified version

```
CREATE OR REPLACE PROCEDURE Get_emp_rec (Emp_number IN NUMBER,o_Empno OUT
NUMBER,o_Ename OUT VARCHAR2, o_Job OUT VARCHAR2, o_Mgr OUT NUMBER, oHiredate
OUT DATE, o_Sal OUT NUMBER, o_Comm OUT NUMBER,o_Deptno OUT NUMBER) IS
BEGIN
SELECT Empno, Ename, Job, Mgr, Hiredate, Sal, Comm, Deptno
 INTO o_Empno, o_Ename, o_Job, o_Mgr, o_Hiredate, o_Sal, o_Comm, o_Deptno
FROM Emp WHERE Empno = Emp_number;
END;
```

Refine example - Changing DB2 target script

- Oracle's SEQUENCE not supported on DB2 UDB for iSeries
  - ► need to remove the CREATE SEQUENCE from the source to eliminate the Translator Omission error message
- DB2 identity column can be used in place of sequence
 - ► very similar functionality

DB2 Script modifications
```
CREATE TABLE EMP(
    EMPNO INTEGER NOT NULL GENERATED ALWAYS AS IDENTITY,
    ENAME CHAR(10),
    JOB CHAR(9),
    MGR INTEGER,
    HIREDATE TIMESTAMP,
    SAL DECIMAL(7,2),
    COMM DECIMAL(7,2),
    DEPTNO INTEGER)


INSERT INTO EMP ( ENAME,JOB,MGR1,HIREDATE,SAL,COMM,DEPTNO )
 VALUES(NAME,JOB1,ORA8.ROUND(MGR1),HIREDATE1,SAL1,COMM1,ORA8.ROUND(DEPTNO1));


SET NEW_EMPNO = (SELECT identity_val_local() FROM SYSIBM.SYSDUMMY1);
```

Note: The VARCHAR columns were manually changed to CHAR for better performance

Generate Data Transfer Scripts

- MTK can create the scripts on local workstation or target iSeries server
  - ► on iSeries script is stored in IFS
- Field Definition Files (FDF) are generated at this step
  - ► FDF defines the format of the data import file used by CPYFRMIMPF
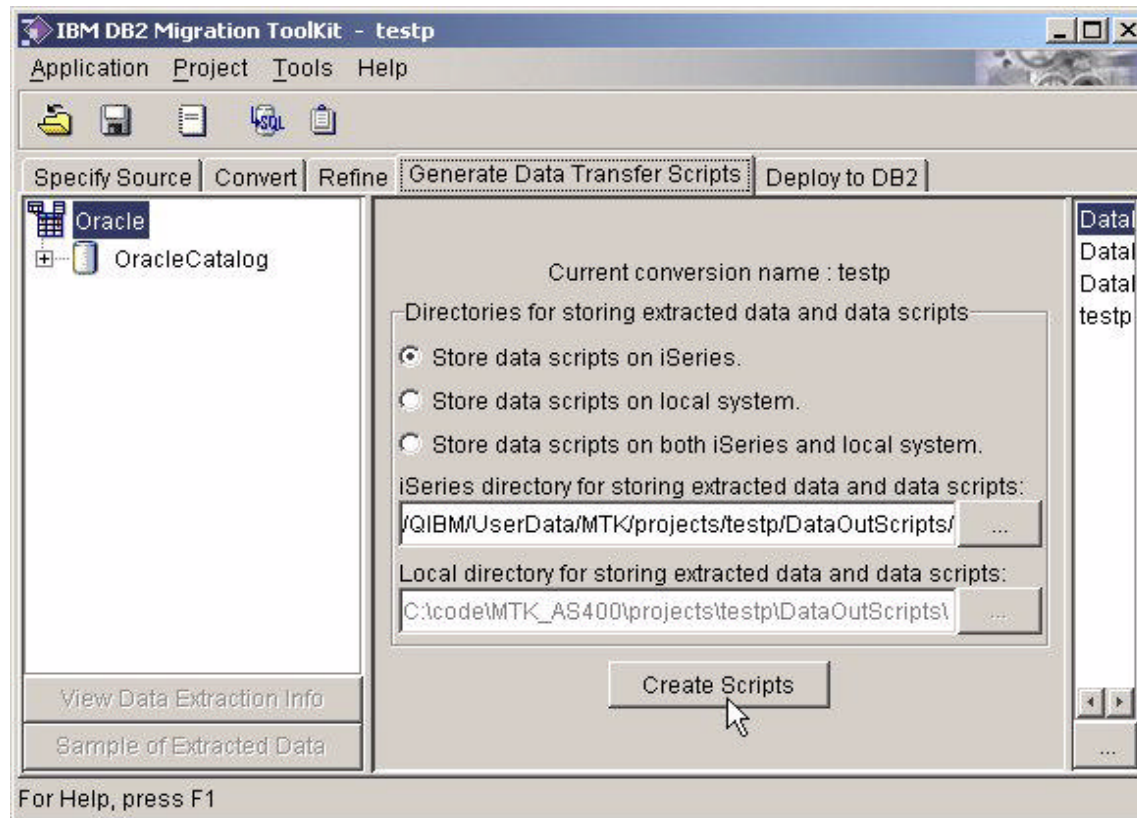
Deploy to DB2

- Several scenarios supported by the deployment utility
  - ► Deploy only the DB2 metadata
  - ► Load only the data
  - ► Combine the metadata deployment with data transfer
- Deploy to DB2 dialog contains summary of the conversion process
  - ► Error messages should not be ignored
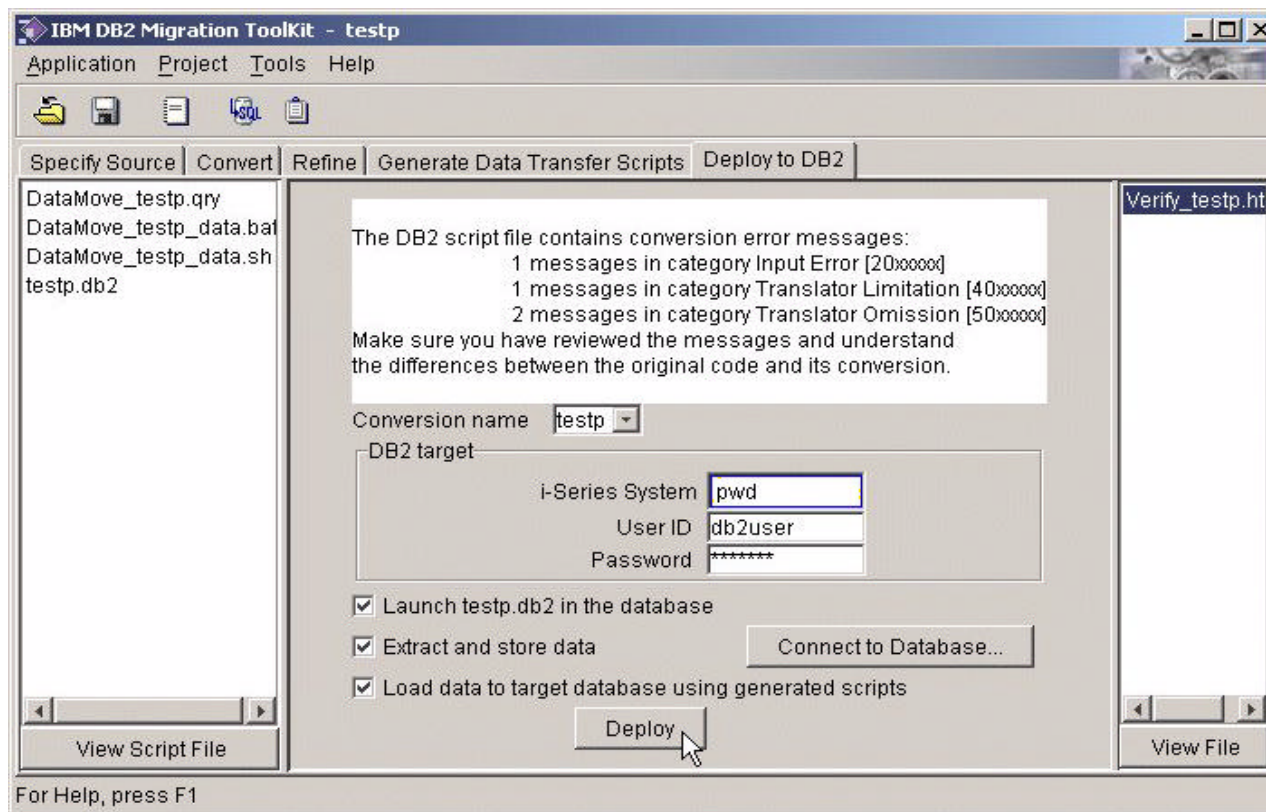
Deployment cont.

- Several tasks are performed by MTK during the deployment
  - ► The data is extracted from Oracle and stored on the iSeries in the project's directory
  - ► The user-defined functions (UDFs) are created on the iSeries
    - – Provided to emulate Oracle functions that do not exist on DB2
    - – Reside in the ORA8 schema
  - ► The DB2 script is executed
    - – metadata gets created on the iSeries
  - ► The source data is loaded into the DB2 tables
  - ► The integrity of the database is checked
  - ► The deployment process is verified

MTK Deployment Reports

- Summary of the migration process for an open project.
  - ► Consists of several reports:
    - – An overview table with Conversion Date, Project, Name...
    - – total number of objects (by type) converted from source to DB2
    - – total number of errors that occurred during conversion by:
      - number of files
      - number of objects
      - number of error messages
  - ► details for each object in error, including the location of the error and the associated error message.
  - ► Table size report (for DB2 tablespace creation)

MTK Deployment Reports example

- The Verify log

Post-Deployment Activities

- The migration process does NOT stop after the deployment
  - ► need to validate that source and target constructs are functionally equivalent
  - ► SQL performance tuning is required
    - set of indexes migrated from Oracle is usually not adequate for DB2
    - Oracle and DB2 cost-based optimizers are different
      - Optimizer hints not supported on DB2
      - QAQQINI query option file maybe used in some cases
    - need to understand the performance characteristics of SQL/PSM routines and triggers on DB2
  - ► DB2 training necessary for application maintenance and SQL tuning
    - The SQL Performance Tuning workshop (S6140) highly recommended

Additional Information

- The MTK can be downloaded free of charge at:
  http://www-919.ibm.com/servers/eserver/iseries/developer/db2/oraclemtk.html
- 'Cut Your Migration Time from Oracle to DB2 UDB for iSeries' article
  http://www.mcpressonline.com/mc
  - ► Select Database from the Feature Sections
- DB2 UDB for iSeries Porting Guide - Oracle to IBM   iSeries white paper
  http://www-919.ibm.com/servers/eserver/iseries/developer/db2/guides.html
- ITSO Redbooks
  - ► Developing Cross-Platform DB2 Stored Procedures (SG24-5485)
  - ► Stored Procedures and Triggers on DB2 UDB for iSeries (SG24-6503)
- Redbooks are available at:
  www.redbooks.ibm.com
- Feedback
  - ► send your questions directly to the development team:
  db2mtki@us.ibm.com
- SQL Performance Tuning workshop
  http://www-1.ibm.com/servers/eserver/iseries/service/igs/db2performance.html