



IBM Software Group

IBM Software Group

Taking RI Relationships from Passive to Active

Ernie Mancill – IBM Sr Certified Consulting IT Specialist

Barry Davis – Product Specialist Rocket Software



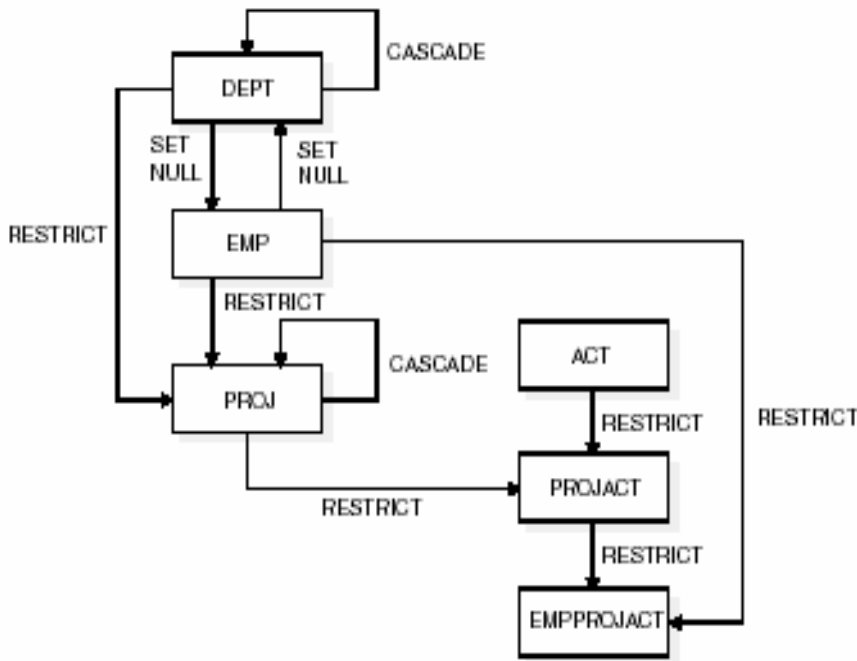
ON DEMAND BUSINESS™

Session Agenda

- RI Terminology and Overview
- DB2 Grouper Overview, Architecture, and Terminology
- DB2 Grouper – Product Usage
- DB2 Recovery Expert
- Wrap-up and Questions



Sample DB2 Table - Relationships



For example in the DB2 sample table shown to the left:

EMP is a Dependent table (child) of DEPT. The DDL (references clause) will look something like:

“Primary Key (EMPNO)

Foreign Key RED (WORKDEPT)

References DSN8510.DEPT on Delete Set Null”

Example of relationships among tables, arrows point from parents to dependents

DB2 Enforced Relationships - Terminology

- **Parent Table** - A table with a primary key. In the sample application, the employee table serves that purpose for employees.
- **Parent Key** - The table should have one column or a set of columns that provides a unique identifier for the rows of the table. This column (or set of columns) is called the parent key of the table. Also known as a Primary Key.
- **Unique Index** - To ensure that the parent key does not contain duplicate values, you must create a unique index on the column or columns that constitute the parent key.



DB2 Enforced Relationships – Terminology cont.

- Foreign Key - can refer to either a unique or a primary key of the parent table. The column names you specify identify the columns of the parent key.
- Child Table – A table which references the Parent table with the Foreign Key clause.
- Referential Constraint – Composed of three elements:
 - ▶ Constraint name
 - ▶ Columns comprising the foreign key
 - ▶ References clause (which includes the constraint rule)



DB2 V8 – Informational Referential Constraints

- An informational referential constraint is a referential constraint that is not enforced by DB2 during normal operations
 - ▶ DB2 ignores informational referential constraints during insert, update, and delete operations.
 - ▶ Some utilities ignore these constraints; other utilities recognize them. For example, CHECK DATA and LOAD ignore these constraints. QUIESCE TABLESPACESET recognizes these constraints by quiescing all table spaces related to the specified table space.
- Informational referential constraints are often useful, especially in a data warehouse environment, for several reasons:
 - ▶ To avoid the overhead of enforcement by DB2.
 - ▶ To allow more queries to qualify for automatic query rewrite.
- To define an informational referential constraint, use the NOT ENFORCED option of the referential constraint definition in a CREATE TABLE or ALTER TABLE statement.
- Note: You should use this type of referential constraint only when an application process verifies the data in a referential integrity relationship.


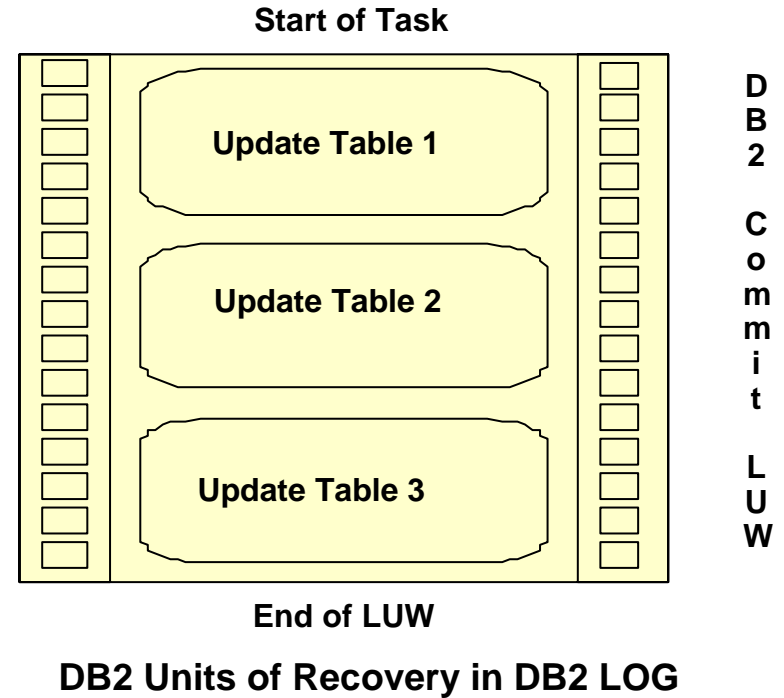
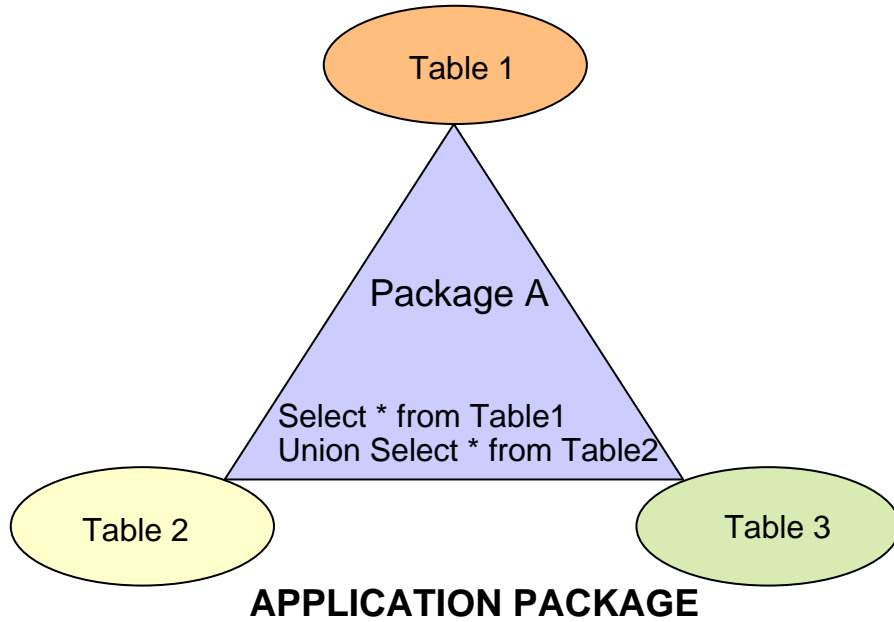


Application Enforcement

- Non-enforced referential constraint when an application process verifies the data in a referential integrity relationship. For example, when inserting a row in a dependent table, the application will verify that a foreign key exists as a primary or unique key in the parent table.
- Advantages
 - ▶ Performance by avoiding enforcement overhead in DB2
 - ▶ Simplistic data model not requiring referential integrity
- Disadvantages
 - ▶ Strict adherence required by all applications updating tables involved these relationships (everyone must play by the same rules).
 - ▶ Hard to keep track (over time) of the relationships and enforcement embedded in the application (until now).



DB2 Relationships - Others



Where to look for relationship information

- DB2 Catalog (V7 for enforced RI, V8 for enforced and non-enforced RI)
 - ▶ Can be discovered using catalog navigation tools such as DB2 Administration Tool
- OEM DB2 Data Modeling Tools (Rational, Bachman, ERWIN, etc.)
 - ▶ Dependent on how the model is maintained
- Application Programs
 - ▶ Hard to decipher and only useful if application program adheres to relationship model.
- Application Software Vendor documentation
 - ▶ Not always updated and some designs are proprietary
- Application design documentation
 - ▶ Value of information directly related to effort spent to keep updated and tends to age rapidly
- And now.....
 - ▶ DB2 Grouper Metadata

Usages for RI Information

- The information about relationships can be used for various kinds of administration activities, such as:
 - ▶ Planning effective backup and recovery strategies
 - ▶ Archiving and Relocation of data
 - ▶ Replicating data
 - ▶ Test Data generation
 - ▶ Planning capacity and analyzing performance
 - ▶ Application object change management, migration and alteration



Additional Usages

- Recovery Expert:
 - ▶ Grouper responds to the need to recover data to a specific point in time, with related data recovered to that same point in time.
 - ▶ Grouper component aids in determining the set of related objects that need to be recovered simultaneously.
- Test Database Generator
 - ▶ Understand what constitutes a intact and complete set of test data from a related tables standpoint and manage the extracting of this subset of data from the operational data store.
- General Data Administration/Understanding
 - ▶ There is a need to understand relationships between various tables.
 - ▶ Much of information is in the catalog, but there is a need to retrieve it and view it in a coherent manner.



Problem Statement

- There are many relationships between the various DB2 objects, such as tables, in a business application. Some of these relationships can be discovered easily, while others cannot be easily found.
- As the complexity of data relationships increases, it becomes increasingly difficult to discover and keep track of this information.
- These relationships can be “dynamic” or “transient”, causing difficulties in managing business applications efficiently.
- Finding these relationships can be critical to treating portions of application data in a coherent manner.

What is Grouper?

- DB2 Grouper – 5799-GXQ. A component of the below indicated tools, with it's own PID but cannot be ordered as a separate, standalone product.
- Grouper is a client-server component that can be used to group tables for a single DB2 for z/OS subsystem.
- Grouper can discover, record, allow administration for, and provide group relationship information to other IBM DB2 tools as a basis for database administration activities.



What is Grouper?

- These activities can be performed manually by the user or by tools that work with Grouper – 5799-GXQ:
 - ▶ DB2 Data Archive Expert for z/OS – 5655-I95
 - ▶ DB2 Test Database Generator for z/OS – 5697-I02
 - ▶ DB2 Test Database Generator for MP – 5724-E4
 - ▶ DB2 Recovery Expert for z/OS – 5697-H74

- And more to come later.....



DB2 Grouper Components

- Grouper Administration
 - ▶ Schedule and execute the Grouper Functions via the use of the Client UI
 - ▶ Create specifications which control Group Discovery
 - ▶ Administration of the group contents
 - ▶ Define non DB2-enforced relationships
- Unit of Work Discovery
 - ▶ Searches archived DB2 log records on the server for relationships. A group of tables are group of tables are suggested as being related if they were all updated in the same unit of work.
- Group Discovery
 - ▶ A stored procedure that runs in the server database and discovers the relationships between tables in the database
- Group View
 - ▶ An SQL view that provides information to examine the various groups



DB2 Grouper Prerequisites

- DB2 Grouper Client UI
 - ▶ Runs on Windows XP or similar
 - ▶ Pentium III class machine
 - ▶ Memory of 256MB or Greater
 - ▶ Requires DB2 Connect 7.2 FP 10 or greater
- DB2 Grouper Metadata
- DB2 Grouper Stored Procedures
 - ▶ DB2 V7 with UQ73301
 - ▶ OS/390 2.10 or Greater with:
 - USS
 - WLM in GOAL Mode
 - RRS
 - Java JRE 1.3.1 +



Grouper - Terminology

The screenshot shows the Grouper interface. On the left is a tree view of the database structure. On the right is a table titled 'NDCDB202 - testprojectempact - NEWVERSION_1 - Group_1' listing tables with their creator names, tablespaces, and boundaries.

Table name	Creator Name	Tablespace	Boundary
ACT	DEMO18	DSN8S51P	N
DEPT	DEMO18	DSN8S51D	N
EMP	DEMO18	DSN8S51E	N
EMPPROJECT	DEMO18	DSN8S51P	N
PROJ	DEMO18	DSN8S51P	N
PROJECT	DEMO18	DSN8S51P	N

The sets appear in date of creation sequence

- Group - A collection of related tables (Group_1 in the above)
- Version - A collection of groups (NEWVERSION_1 shown above)
- Set - A collection of one or more versions (testprojectempact)
- Starting point tables - Tables that the user specifies in the group discovery options to limit the scope of the group discovery (DEMO18.PROJ) in the above

DB2 Grouper CUI Choices

DB2 Grouper Launchpad

POPO - VERSION1

(Optional inputs for all sets)

Define non-enforced RI

Run a Unit of Work Discovery

Check the status of a UOW Discovery

Preparing for a Group Discovery

Create a set to hold groups

Name a version of the set

Select starting points

Select boundary objects

Select relationships to ignore

Add relationships

Discovering Groups

Run a Group Discovery

Check the status of a Group Discovery

Close launchpad and view the results

Submit a job to be run on the targeted server that will search the DB2 logs for Unit of Work relationships. These UOW relationships are used globally as input for all subsequent Group Discoveries. If you want to have UOW relationships considered, then run this job before the Group Discovery job, as this job provides input to Group Discovery.

To perform this task from the Grouper tree view, select Run Unit of Work Discovery in the main menu.

Global inputs

Context inputs

Group Discovery

Unit of Work relationships

DSNB

VERSION1 POPO

Do not show the launchpad again when DB2 Grouper opens.

You have a choice of CUI “look and feel’s”

Launchpad

Started up by default at product initiation

Provides a “step by step” prompting through various Grouper activities. Intended for first time users...a hand holding UI.

Can be turned off This is an optional screen and can be disabled by clicking on the radio button on the bottom of the Panel.

Example one - Set Creation with RI

Click on the “Create a set to hold groups” button from the launchpad.

DB2 Grouper Launchpad ✕

POPO - VERSION1

(Optional inputs for all sets)

- Define non-enforced RI
- Run a Unit of Work Discovery
- Check the status of a UOW Discovery

Preparing for a Group Discovery

- Create a set to hold groups
- Name a version of the set
- Select starting points
- Select boundary objects
- Select relationships to ignore
- Add relationships

Discovering Groups

- Run a Group Discovery
- Check the status of a Group Discovery
- Close launchpad and view the results

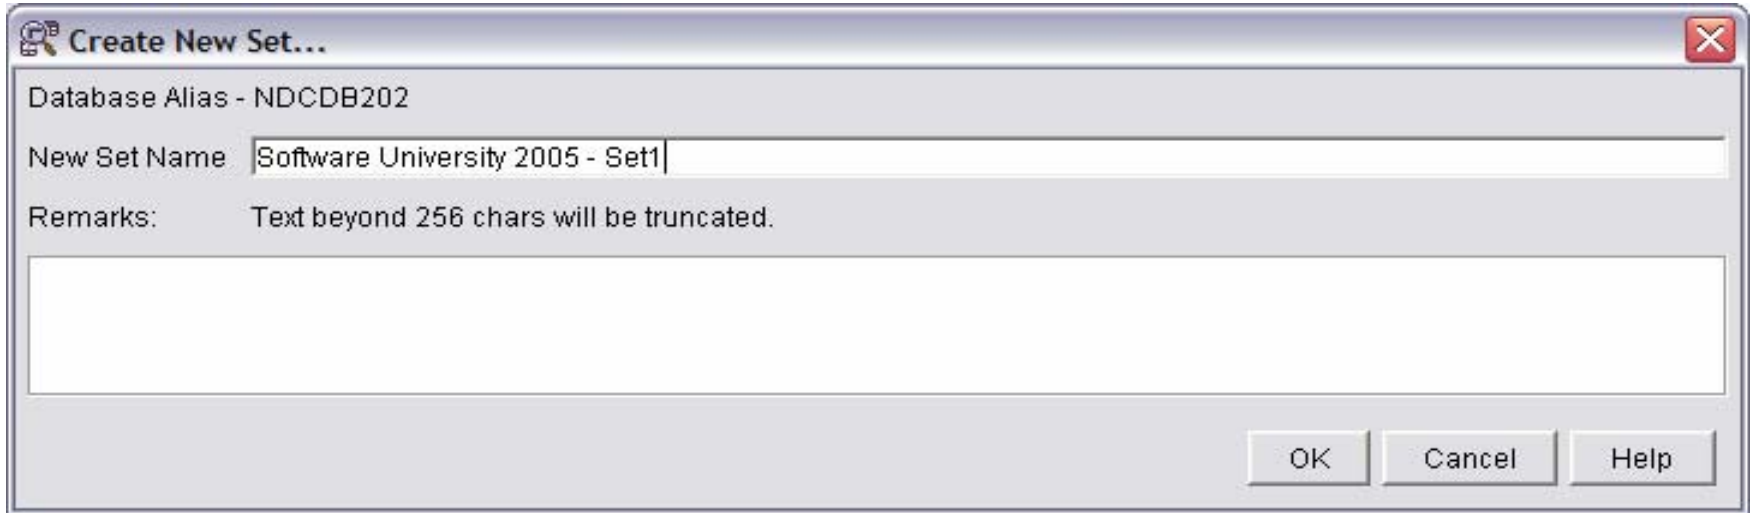
Submit a job to be run on the targeted server that will search the DB2 logs for Unit of Work relationships. These UOW relationships are used globally as input for all subsequent Group Discoveries. If you want to have UOW relationships considered, then run this job before the Group Discovery job, as this job provides input to Group Discovery.

To perform this task from the Grouper tree view, select Run Unit of Work Discovery in the main menu.

The diagram illustrates the data flow in the DB2 Grouper Launchpad. It is divided into three main sections: Global inputs, Context inputs, and Group Discovery. In the Global inputs section, 'Unit of Work relationships' (represented by a magnifying glass over a grid) provides input to the Context inputs section. The Context inputs section shows a 3D cube and a cylinder representing data sources. These inputs feed into the Group Discovery section, which includes a 'DSNB' (Data Set Name Base) table and a 'VERSION1 POPO' (Person-Organization) table. Arrows indicate the flow of information from the inputs through the processing stages to the final discovery results.

Do not show the launchpad again when DB2 Grouper opens.

DB2 Grouper – Set Creation



Create New Set...

Database Alias - NDCDB202

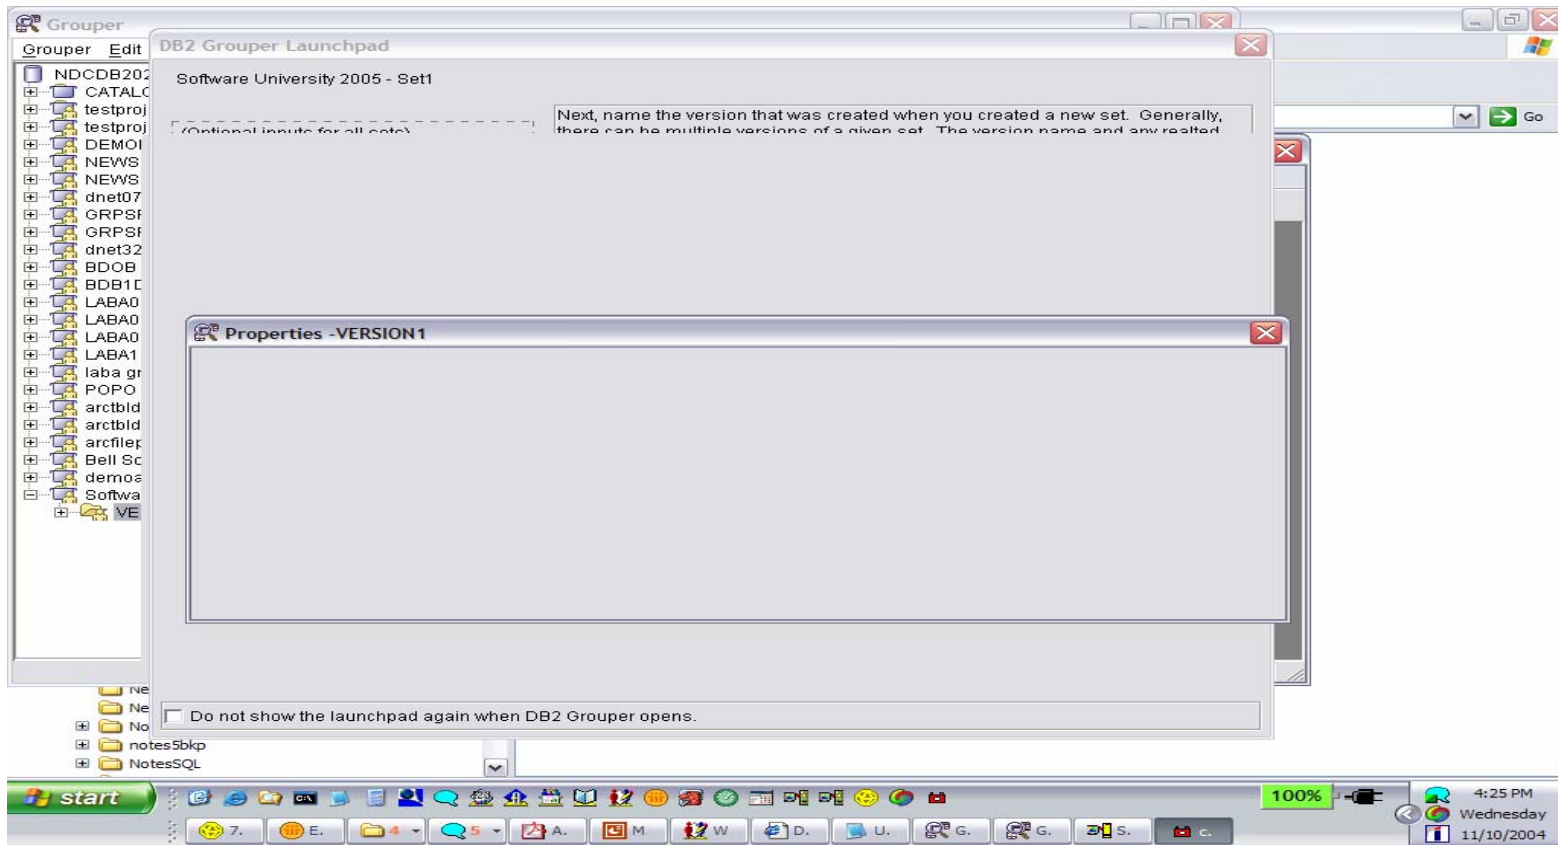
New Set Name: Software University 2005 - Set1

Remarks: Text beyond 256 chars will be truncated.

OK Cancel Help

Provide a meaningful Set name, and an optional description in the remarks field. Note: all of the enterable fields are case sensitive so any object name will be stored in the metadata in mixed case format

DB2 Grouper – Create Version



In our example, we'll default the version name to VERSION1

DB2 Grouper – Group creation

Configure Group Discovery Options

Database Alias - NDCDB202 , Set - Software University 2005 - Set1 , Versi...

Starting Point Tables | Boundary Objects | Ignored Relationships | Additional Relationships

Specify key tables within a subset to limit the scope of the Group Discovery.

Table Name	Schema Name	Type
------------	-------------	------

Add Table

Remove

Relationship Types Used For Grouping

All Relationship Types Referential Integrity Relationships Only

OK Cancel Help

Select this radio button for only look for RI relationships

Next, we'll specify the group by selecting a starting point table. Also notice that we've selected the RI option only

DB2 Grouper – Add table to Group



Add Tables

Creator Name (use '*' as wildcard character)

DEMO06

Table Name (use '*' as wildcard character)

PROJ

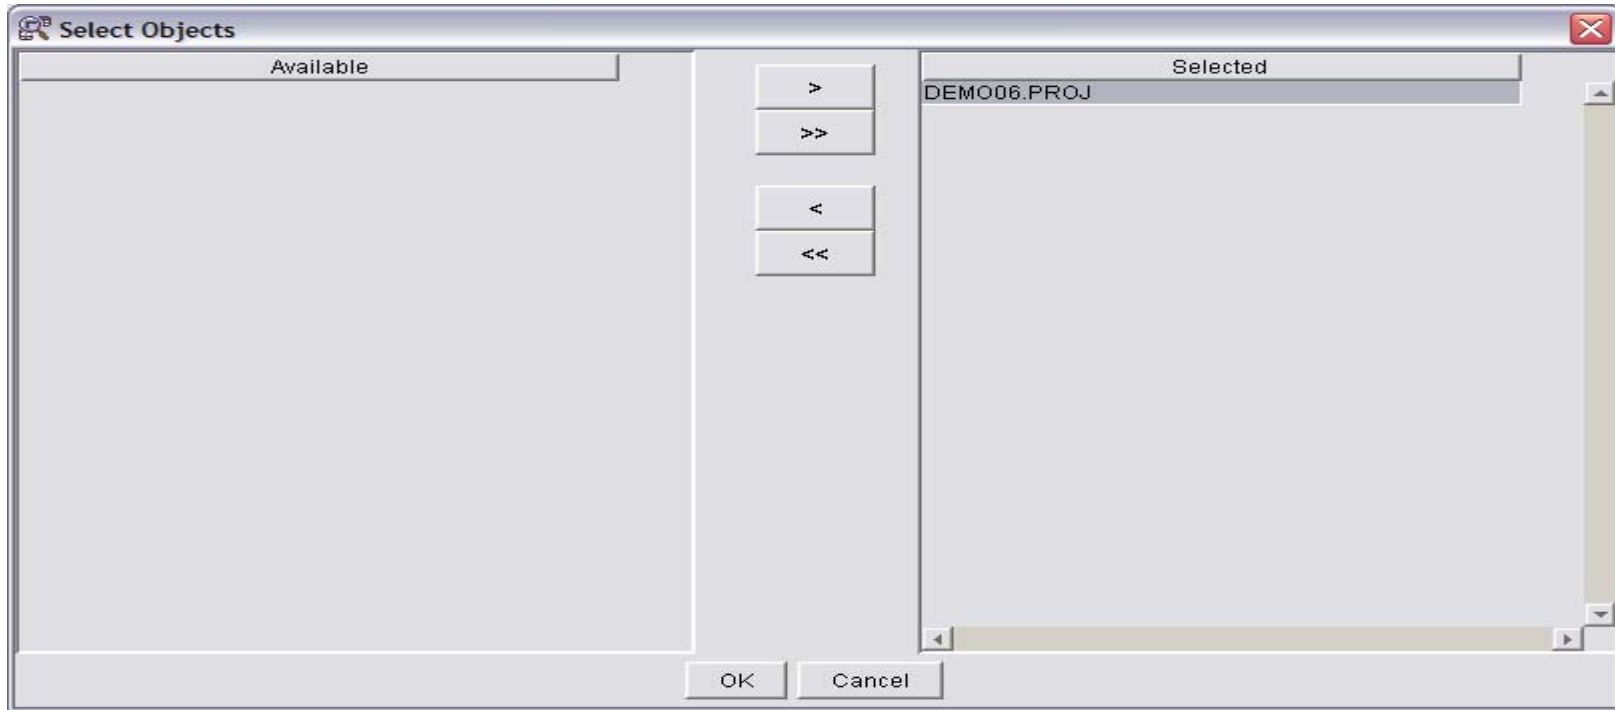
Select from search list

Directly add named table

OK Cancel Help

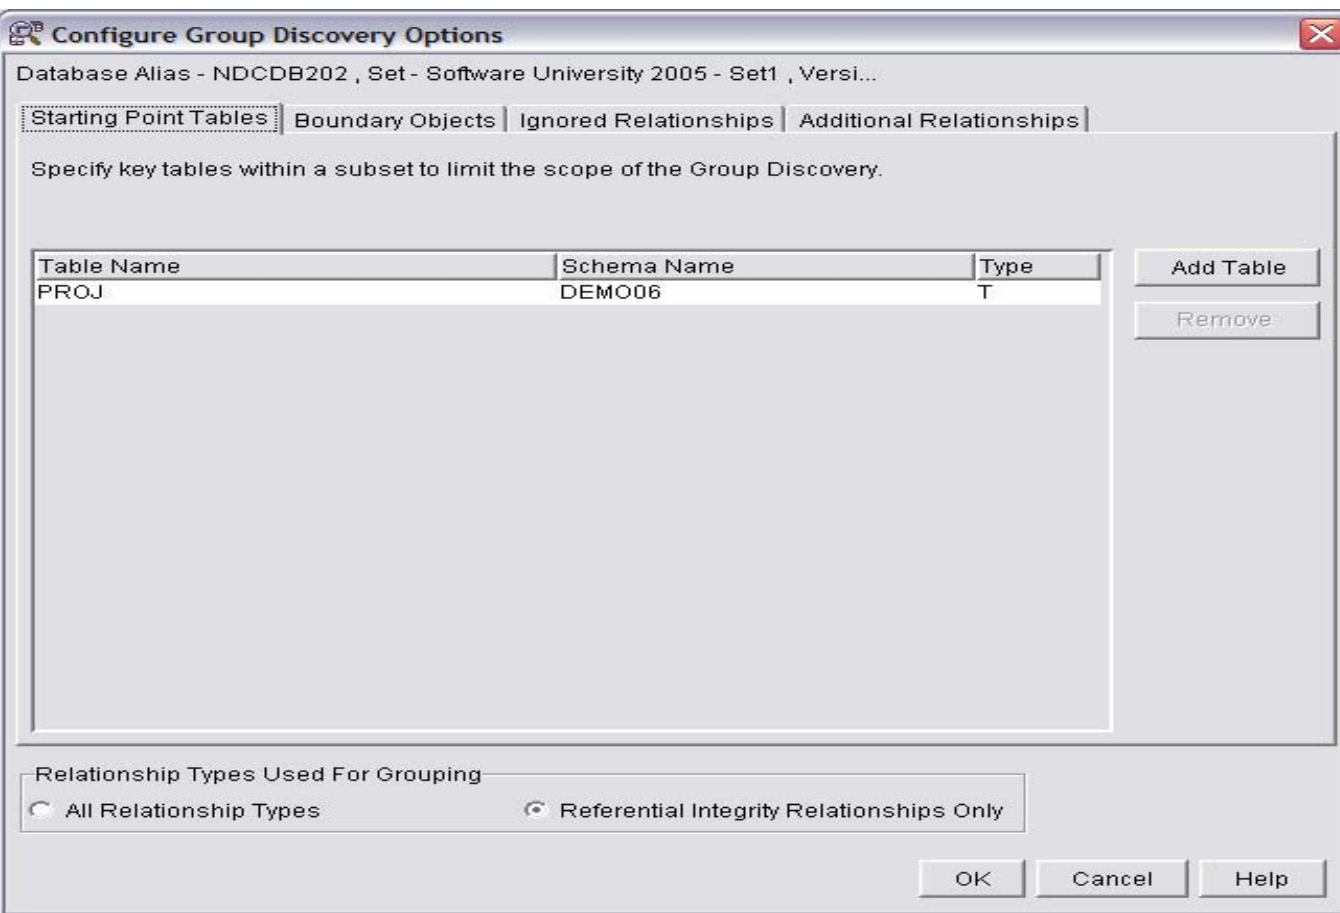
For this drop-down you can specify a complete or partial owner and table name. For large catalogs, better performance will result with a fully qualified table name and selecting the “Directly add named table” radio button will cause an equal clause to be used instead of LIKE.

DB2 Grouper – Add table to Group



The results will be presented in a select window, select one of the available column objects and press the “>” arrow to move it to the selected column

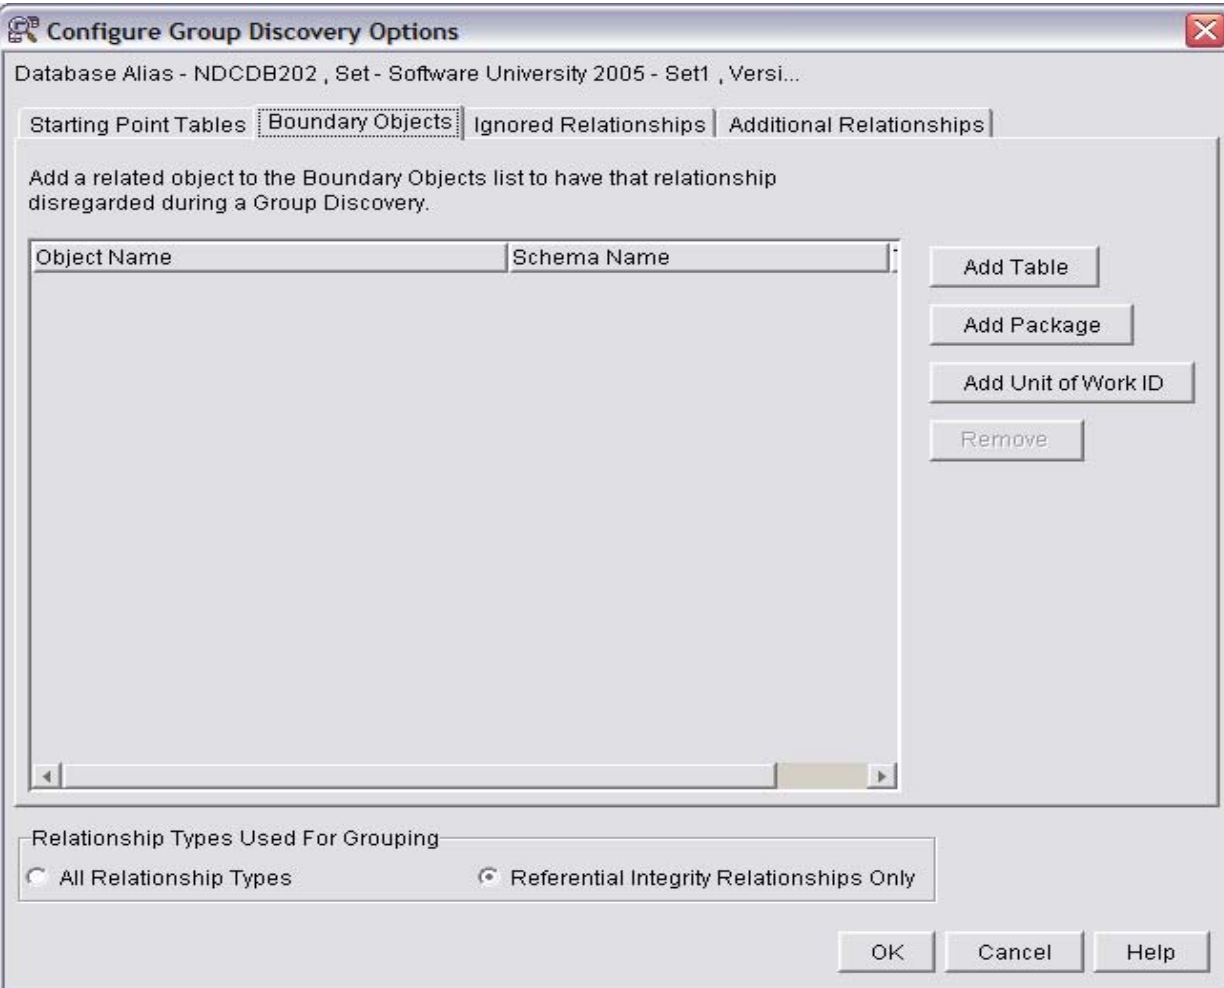
DB2 Grouper – Starting point table panel



Additional tables can be added ...there can be more than 1 starting point table, for this definition there is only one starting point table.

A special point.....for DAE there is only one Starting point table.

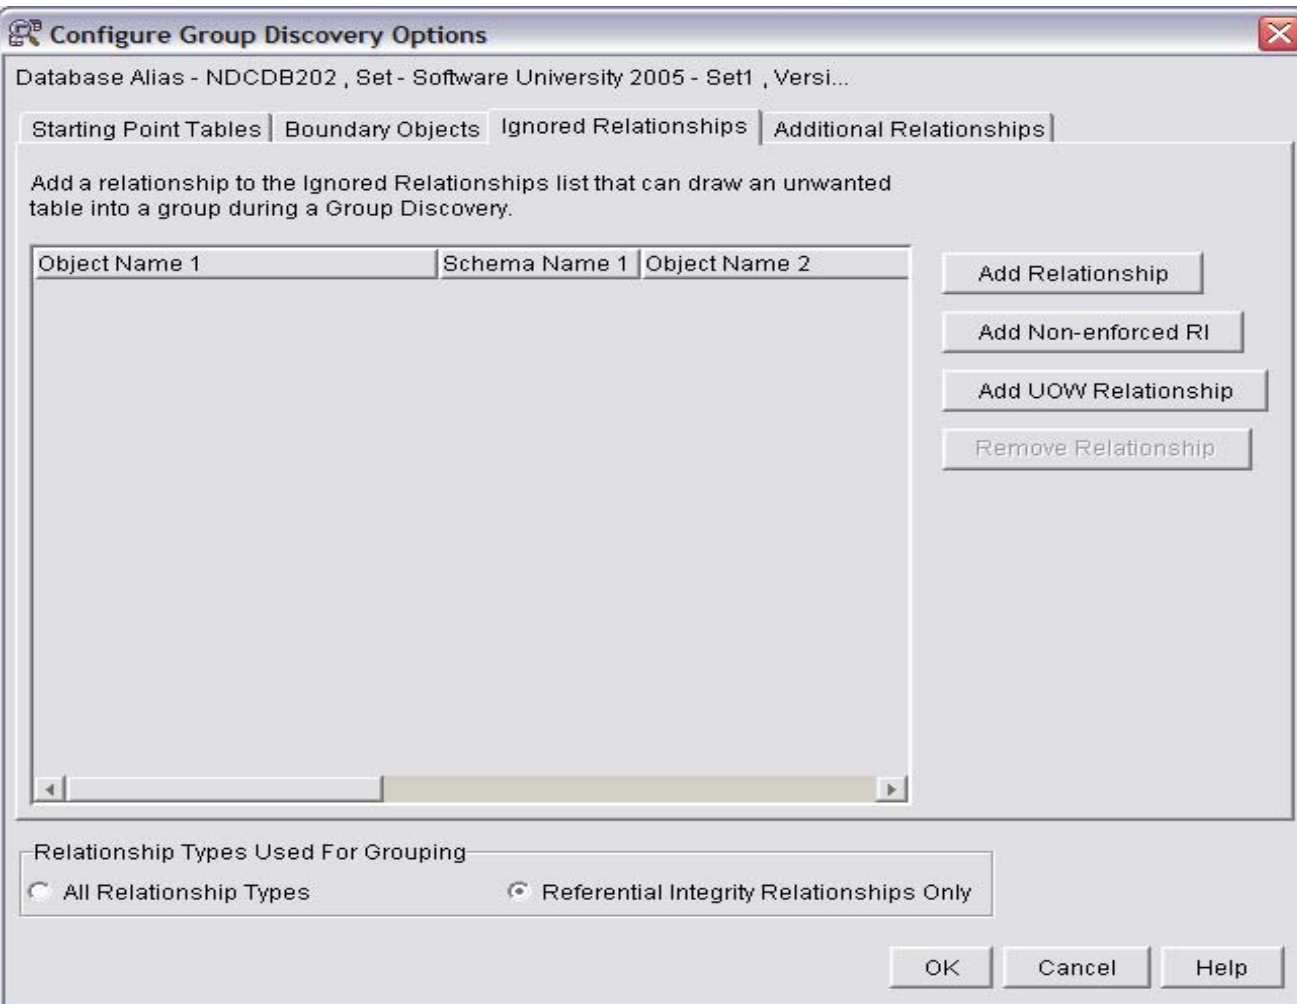
DB2 Grouper – Boundary Objects Panel



A boundary object directs the group discovery to include the object in the group, but not to consider it as a source for drawing additional tables into that group.

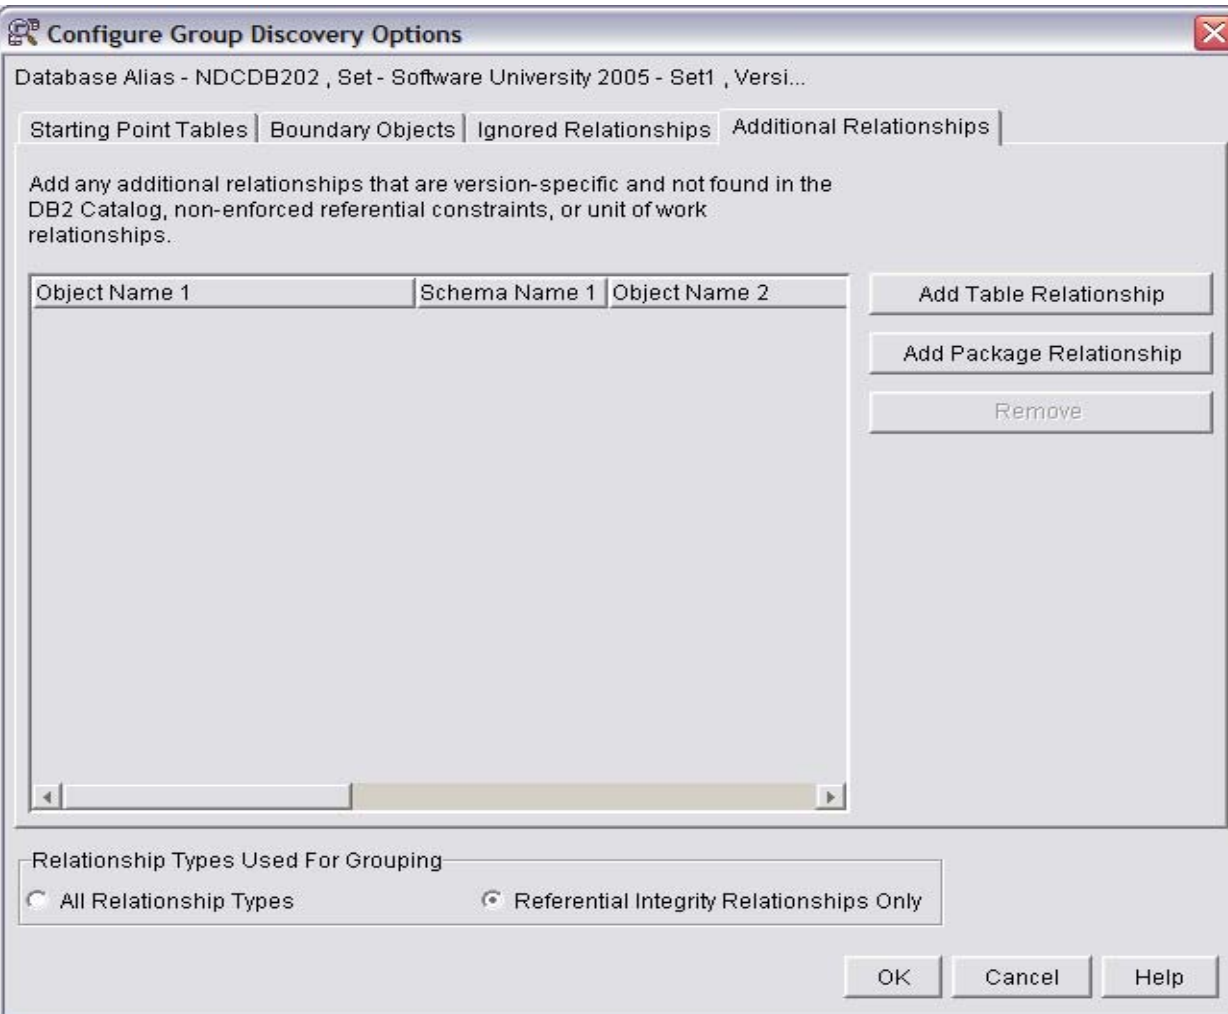
If you specify starting points *a*, *b*, and *c*; and if *b* is related to *c*, they'll be in the same group and the version will result in 2 groups (one with *a* and its related tables, another with *b*, *c*, and their related tables).

DB2 Grouper – Ignored Relationships Panel



If a relationship will draw an unwanted table into a group, you can exclude it for that discovery by adding it to the Ignored Relationships list. In this case, during a group discovery, the unwanted table would not be drawn into a group (unless it is drawn in as a result of other relationships)

DB2 Grouper – Additional Relationships Panel



If during a group discovery, you want to have a specific relationship such as a table or package relationship considered that you do not believe will be found in the DB2 catalog, non-enforced referential constraints, or unit of work relationships, you can add that relationship to the Additional Version-Specific Relationships list

DB2 Grouper – Launchpad “Run a Group Discovery”

DB2 Grouper Launchpad

Software University 2005 - Set1 - VERSION1 - 11-10-2004

(Optional inputs for all sets)

- Define non-enforced RI
- Run a Unit of Work Discovery
- Check the status of a UOW Discovery

Submit a job to be run on the targeted server that will search the DB2 logs for Unit of Work relationships. These UOW relationships are used globally as input for all subsequent Group Discoveries. If you want to have UOW relationships considered, then run this job before the Group Discovery job, as this job provides input to Group Discovery.

To perform this task from the Grouper tree view, select Run Unit of Work Discovery in the main menu.

Preparing for a Group Discovery

- Create a set to hold groups
- Name a version of the set
- Select starting points
- Select boundary objects
- Select relationships to ignore
- Add relationships

Discovering Groups

- Run a Group Discovery
- Check the status of a Group Discovery
- Close launchpad and view the results

Do not show the launchpad again when DB2 Grouper opens.

As defined earlier, This component (Group Discovery) is a stored procedure that runs in the server database and discovers the relationships between tables in the database. It then groups the tables based on these relationships.

DB2 Grouper – Group Discovery

Run Group Discovery

Database Alias - NDCDB202 , Set - Software University 2005 - ...
Version - VERSION1 - 11-10-2004

Version-specific settings

These inputs only affect Group Discovery for this version

Starting points: 1
Boundary objects: 0
Ignored Relationships: 0
Additional relationships: 0

Global relationship inputs

These inputs affect Group Discovery on all versions

Non-enforced RI
Last edited: 2004-09-09 06:07
Total number of relationships: 3

Unit of Work relationships
Last run: 2004-08-24 09:46
Total number of relationships: 19

OK Cancel Help

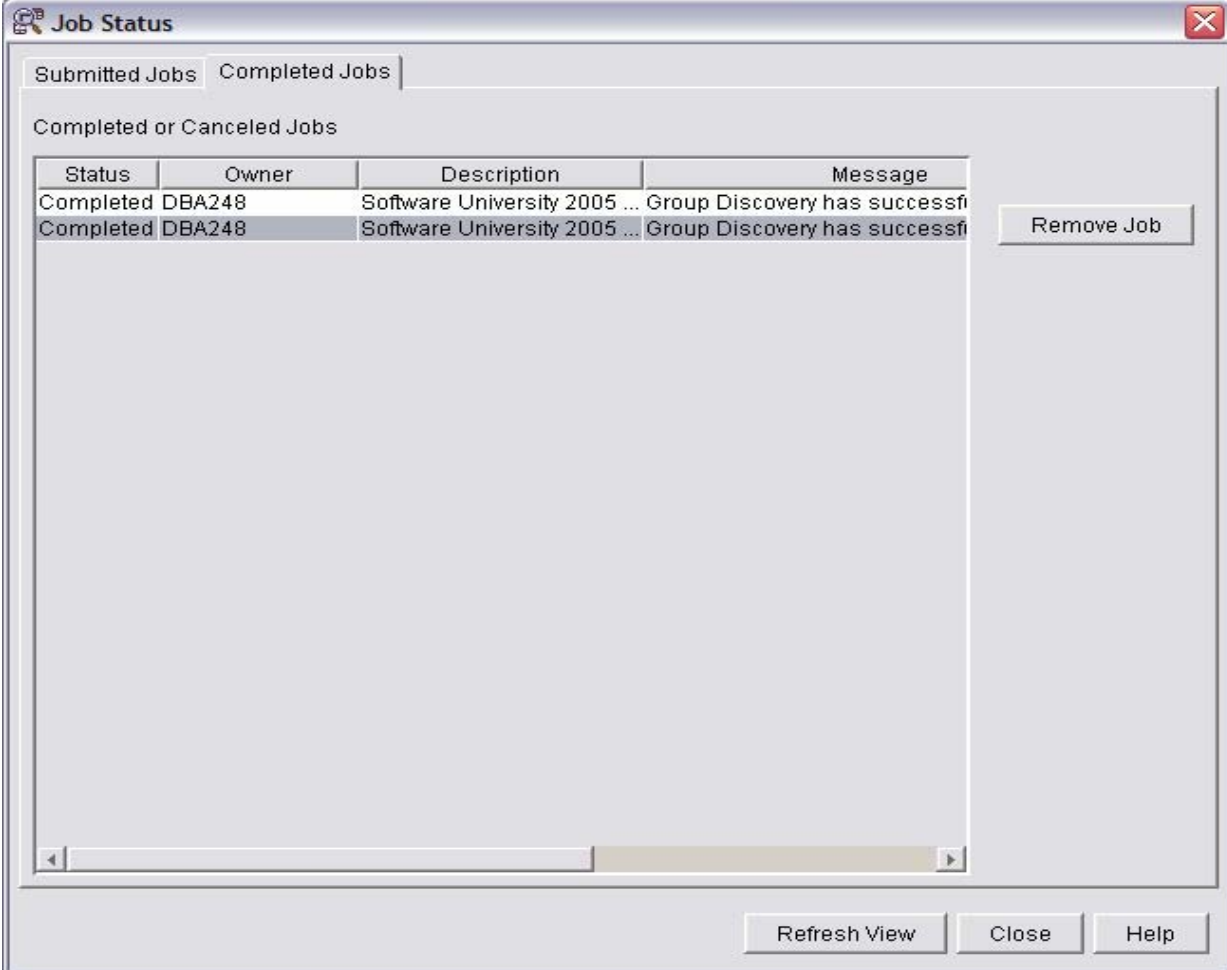
Version specific settings are only for this Version/Set and were previously defined

Global relationships are created from prior UOW and manually entered RI and can potentially affect all Group Discovery runs

This popup is informational and recaps the settings which will influence this Group Discovery run

DB2 Grouper Discovery Job Status

Monitor job execution
Either using ISPF and
Jobnumber from CLP or
From the Job Status
window



The screenshot shows a window titled "Job Status" with two tabs: "Submitted Jobs" and "Completed Jobs". The "Completed Jobs" tab is active, displaying a table of jobs. The table has four columns: "Status", "Owner", "Description", and "Message". Two rows are visible, both with a status of "Completed" and owner "DBA248". The description for both is "Software University 2005 ..." and the message is "Group Discovery has successf". A "Remove Job" button is located to the right of the table. At the bottom of the window, there are three buttons: "Refresh View", "Close", and "Help".

Status	Owner	Description	Message
Completed	DBA248	Software University 2005 ...	Group Discovery has successf
Completed	DBA248	Software University 2005 ...	Group Discovery has successf

DB2 Grouper – Display Discovery Results

The screenshot shows the IBM DB2 Grouper application window. The left pane displays a tree view of groups, with 'VERSION1 - 11-10-2004' selected. A context menu is open over this selection, listing various actions such as 'Configure Group Discovery Options...', 'Run Group Discovery...', and 'View Group Discovery Results'. A blue callout box points to the 'View Group Discovery Results' option, stating: 'From the taskbar dropdown, select View Group Discovery Results'. Another blue callout box points to the 'VERSION1 - 11-10-2004' group in the tree, stating: 'Note: You need to Highlight the Group Version that you wish to review'. The right pane shows a table with the following data:

CDB202 - Software University 2005 - Set1 - VERSION1 - 11-10-2004			
Group Name	Remarks	Created By	Create Time
up_1		DBA248	2004-11-12 04:05

At the bottom of the window, there are 'Next' and 'Filter' buttons.

DB2 Grouper – Group/Version/Set Tree View

Grouper

Grouper Edit Selected Help

NDCDB202

- CATALOGWIDE
- testproj
- testprojectcompact
- DEMOFORFIELDTELECO
- NEWSET1
- NEWSET2 Unit of Recovery Discovery
- dnet072a1
- GRPSPACE
- GRPSPACE2
- dnet329a
- BDOB ARCHIVE1
- BDB1DEMO
- LABA02
- LABA03
- LABA05
- LABA16
- laba grouper
- POPO
- arctbldelemprjact
- arctbldelempprjact
- arcfileprojectdel
- Bell South
- demoarctblproject
- Software University 2005 - Set1
 - VERSION1 - 11-10-2004
 - Group_1
- ARCHDEMO15CITIGR
- ARCHDEMO1CITIGRP
- archdemo1citigrp
- ARCHDEMO1CG
- ARCHDEMO2CG
- ARCHDEMO3CG
- ARCHDEMO3

NDCDB202 - Software University 2005 - Set1 - VERSION1 - 11-10-2004

Group Name	Remarks	Created By	Create Time
Group_1		DBA248	2004-11-12 04:05:

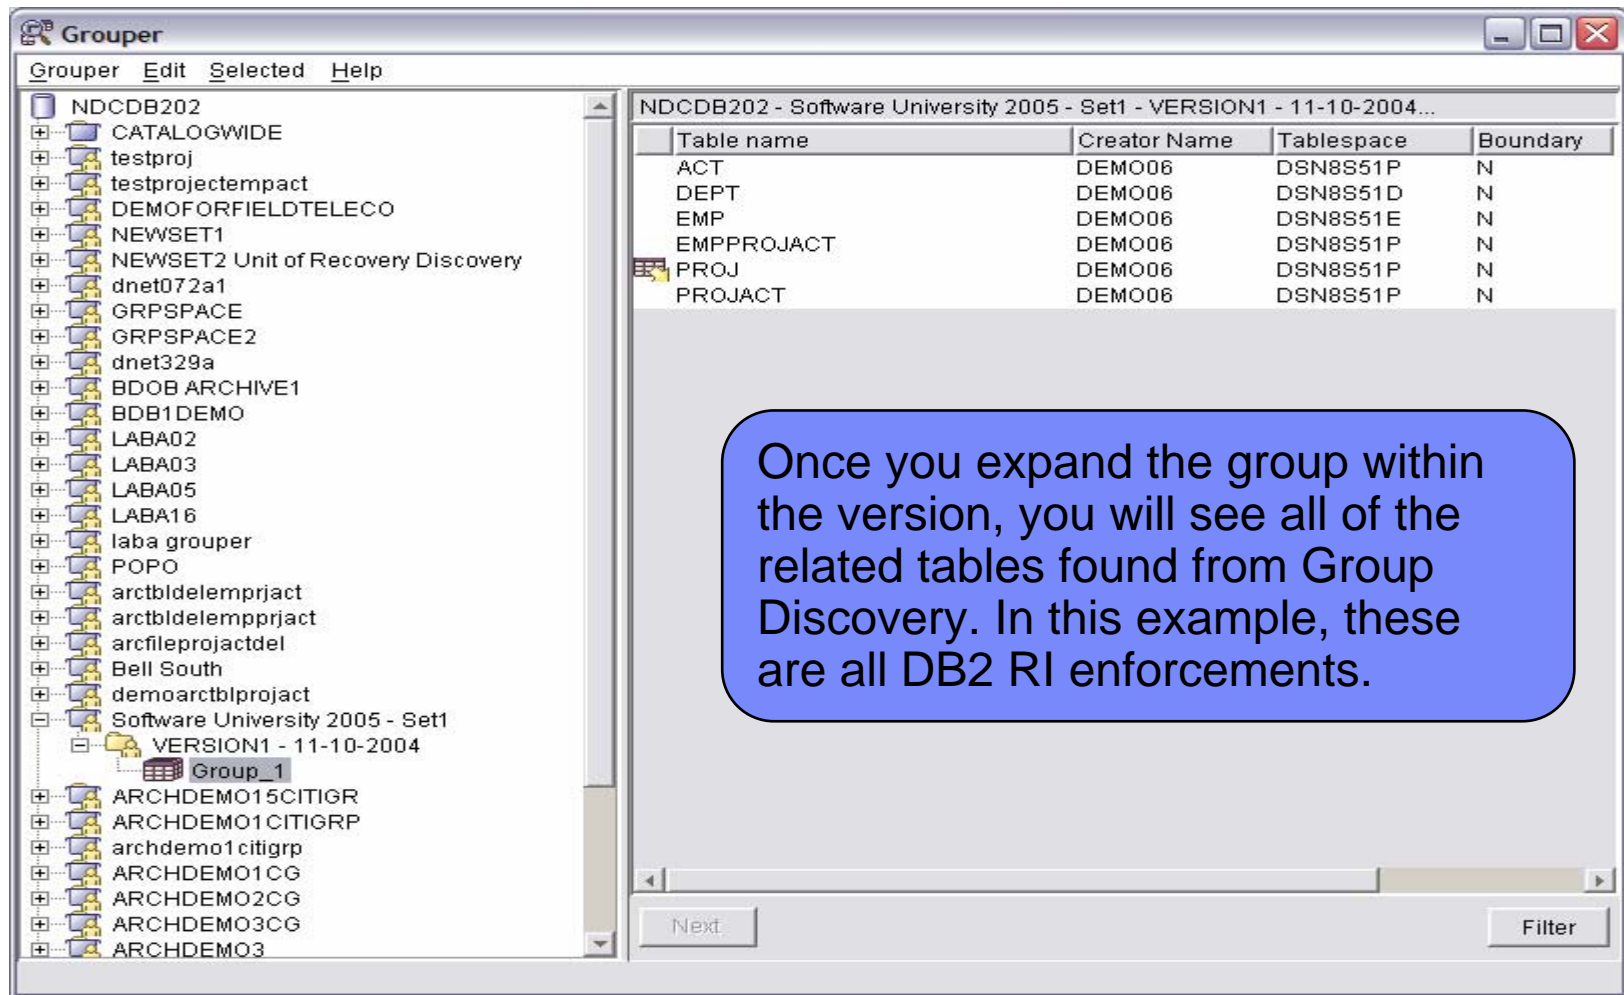
Software University 2005 - Set1

- VERSION1 - 11-10-2004
 - Group_1

Note: The view will contain a group icon upon completion of a Group Discovery Run. Note: there could be more than one group in a version

Next Filter

DB2 Grouper – Discovered group relationships



Once you expand the group within the version, you will see all of the related tables found from Group Discovery. In this example, these are all DB2 RI enforcements.

Table name	Creator Name	Tablespace	Boundary
ACT	DEMO06	DSN8S51P	N
DEPT	DEMO06	DSN8S51D	N
EMP	DEMO06	DSN8S51E	N
EMPPROJACT	DEMO06	DSN8S51P	N
PROJ	DEMO06	DSN8S51P	N
PROJACT	DEMO06	DSN8S51P	N

DB2 Data Archive Expert

```

Search for related Tables?

Find related tables? ==> Y (Yes/No)

Starting point table: PROJ
Creator . . . . . : DEM006
Database name . . . : DEM006
DB2 system . . . . : DSNB

Command ==> _____
    
```

Specify "Y" to call Grouper for discovery

```

AHXV11 ----- Select Related Tables ----- Row 1 to 5 of 5

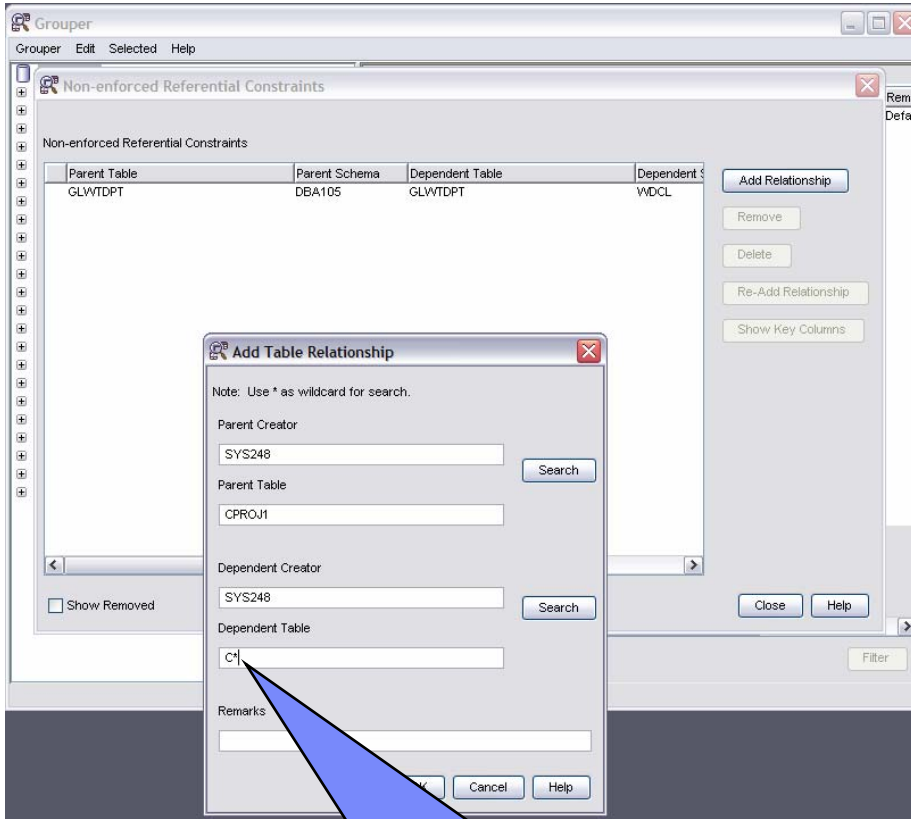
Archive Specification : newone
Starting point table : PROJ
Creator . . . . . : DEM006
DB2 system . . . . : DSNB

Line commands are:
S - Select table for archive unit  S* - Select all tables
D - Deselect table  D* - Deselect all tables

Cmd * Table name      Creator  Database Table space
-----
___  ACT              DEM006  DEM006  DSN8S51P
___  DEPT             DEM006  DEM006  DSN8S51D
___  EMP              DEM006  DEM006  DSN8S51E
___  EMPPROJECT       DEM006  DEM006  DSN8S51P
___  PROJECT          DEM006  DEM006  DSN8S51P
***** Bottom of data *****
    
```

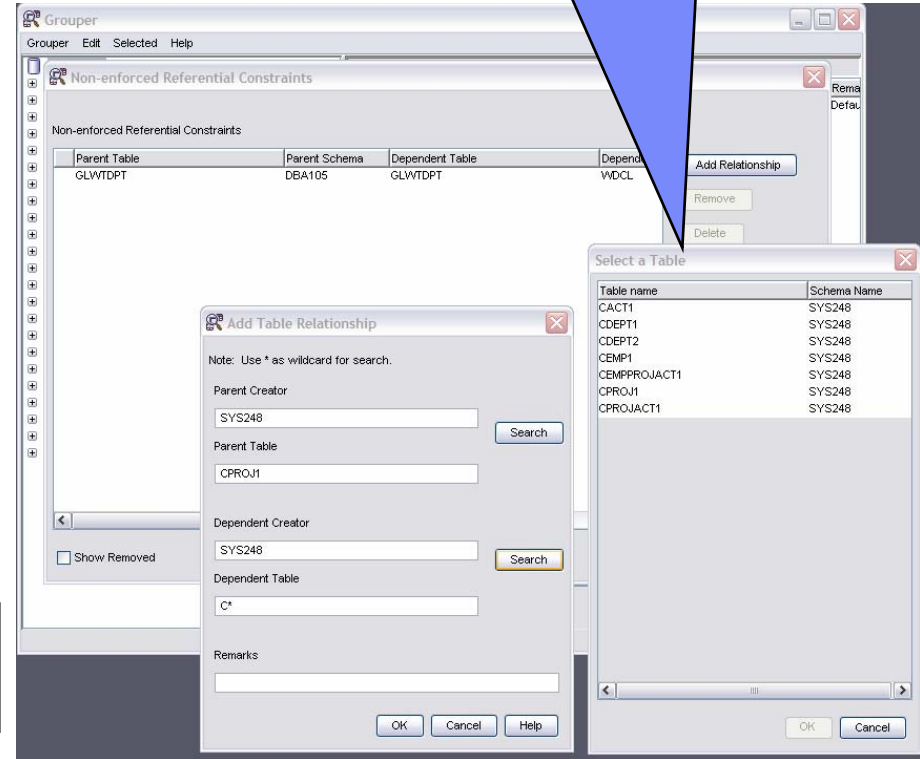
Tables related to PROJ as discovered

Second Example – Un-enforced Relationships

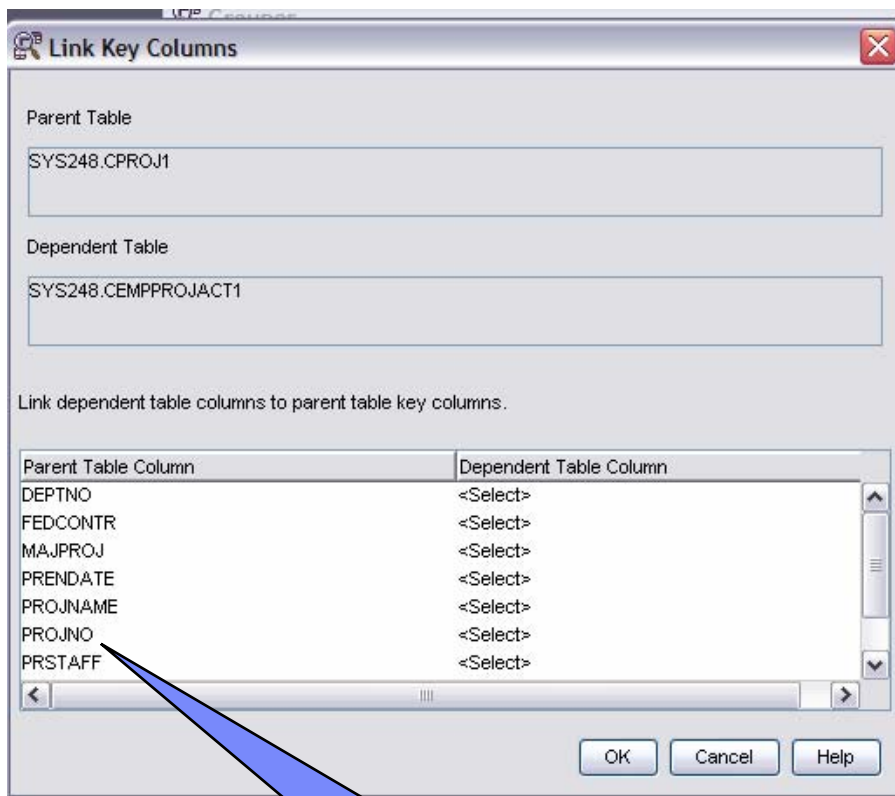


Specific qualified parent table but wildcard for dependents

Result is a dropdown selection list for the dependent table

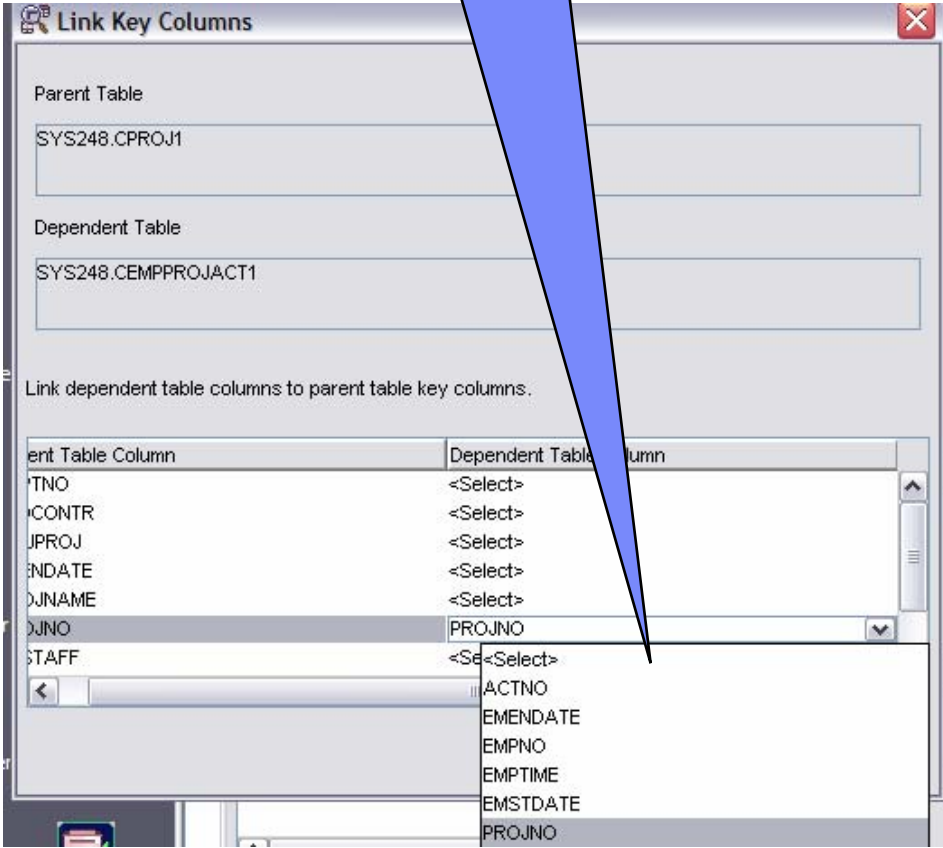


Un-enforced relationship – Link Key Columns



Columns of the starting point (parent) table

Linked dependent table columns from list



Un-enforced Relationship – A second table

Add Table Relationship

Note: Use * as wildcard for search.

Parent Creator
SYS248

Parent Table
CPROJ1

Dependent Creator
SYS248

Dependent Table
CPROJACT1

Remarks

OK Cancel Help

Link Key Columns

Parent Table
SYS248.CPROJ1

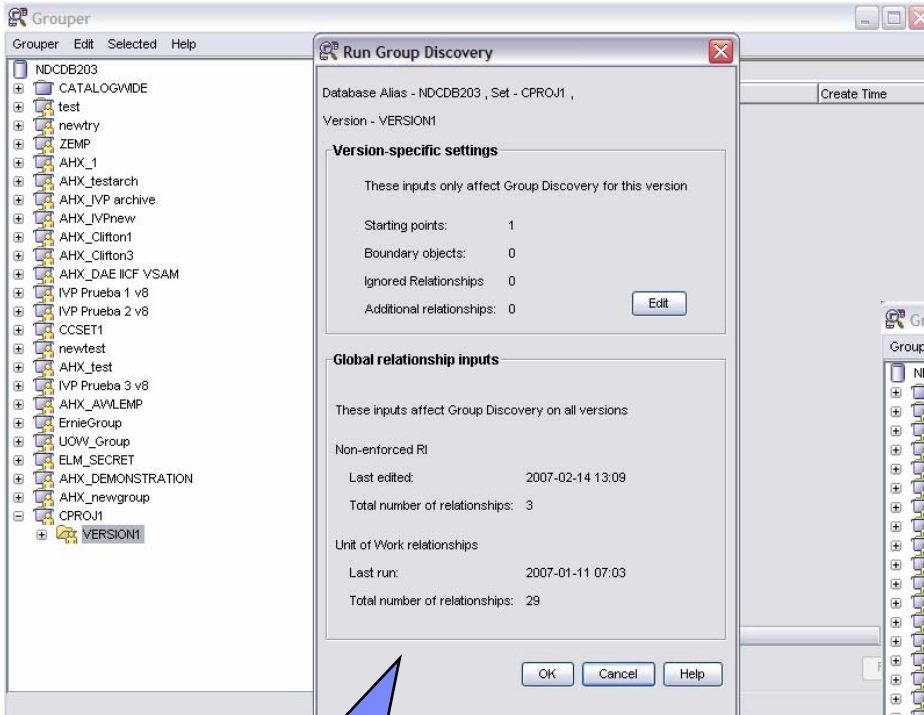
Dependent Table
SYS248.CPROJACT1

Link dependent table columns to parent table key columns.

Parent Table Column	Dependent Table Column
MAJPROJ	<Select>
PRENDATE	<Select>
PROJNAME	<Select>
PROJNO	PROJNO
PRSTAFF	<Select>
PRSTDATE	<Select>
RESPEMP	<Select>

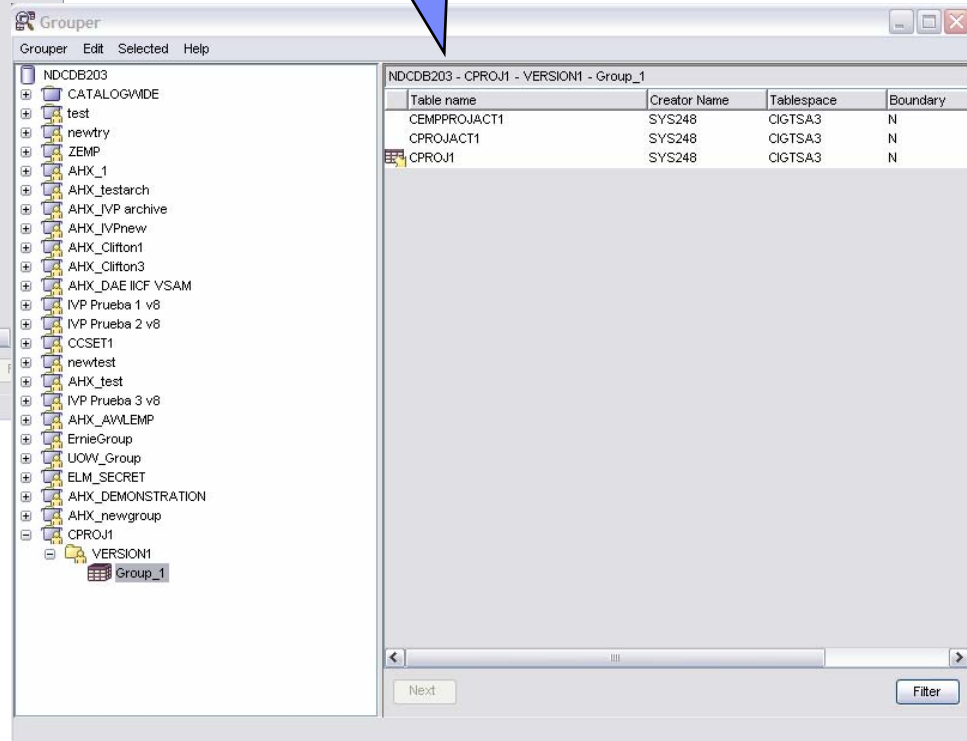
OK Cancel Help

Un-enforced Relationships – Result



Options in effect for this group discovery

Result group including starting point and dependent children



DB2 Recovery Expert for z/OS

- Provide Expert assist for performing many types of DB2 recoveries
- DB2 Recovery Expert analyses the requested recovery
 - ▶ Providing a selection of possible recovery plans
 - Selecting for you, the needed recovery assets and utilities
 - Assists in selecting recovery points
 - Builds required JCL



Select objects to recovery

- Selecting Groups of Objects to Recover
 - ▶ Recovery Expert provides 4 ways to specify a group of objects to be recovered
 - Explicit selection (one by one)
 - By pattern
 - **Via RI relationships (Grouper and catalog RI)**
 - **Grouper groups can be selected in the object tree**
 - If Automation Tool is available, via Automation Tool object profiles
 - Automation Tool profiles can be selected in the object tree



Recover Points

- Specifying a recovery point
 - ▶ Recover to current
 - ▶ Recover to a timestamp
 - Can be selected from
 - Object version
 - Recovery history events
 - Quiet points
 - ▶ Recover to an RBA/LRSN
 - Can be selected from
 - Recovery history events
 - Quiet points



DB2 Recovery Expert for z/OS

- Recovery Expert provides
 - ▶ Application recovery
 - ▶ Drop recovery
 - ▶ Single or multiple objects
 - ▶ Dependency analysis
 - ▶ Quiet point analysis
 - ▶ SYSTEM RESTORE

- Provides DB2 Log Analysis services
 - ▶ undo or redo recoveries
 - ▶ Quite times



DB2 Recovery Expert for z/OS

- **Recovery of a set of Grouper-related objects**
 - ▶ A key feature is the ability to easily recover groups of objects which are related in some way
 - ▶ DB2 Grouper is a key component of this support
 - ▶ Recovery Expert uses the object relationship information provided by Grouper in two different ways:
 - allow user to browse existing Grouper sets, versions, and groups in order to make a selection
 - automatically augmenting the set of objects to be recovered with other objects in containing Grouper Sets



RE z/OS starting screen

DB2 Recovery Expert for z/OS

File Edit Tools Help

Launchpad

Recovery Select to recover data or dropped objects in an operational database system.

Log Analysis Select to analyze database log files to determine points of consistency (quiet times) for an object or objects.

Specifications Select to work with previously saved recovery and log analysis specifications.

Messages

Status	Location	Specification	Type	Timestamp	Description
success	I71A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 8:14:37 AM	Recovery plan generation for DEMODB07.DEMOTS07 has completed successfully.
success	I81A on RS01	(unsaved)	Log analysis	Jun 18, 2006 9:37:33 AM	The job ARYJOB / JOB00743 has ended on RS01 with MAXCC=0.
success	I81A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 9:46:38 AM	Recovery plan generation for DEMODB03.DEMOTS03 has completed successfully.
warning	I81A on RS01	(unsaved)	Recovery plan execution	Jun 18, 2006 9:49:00 AM	The job ARYJOB / JOB00744 has ended on RS01 with MAXCC=4.
success	I81A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 10:32:57 AM	Recovery plan generation for TDBG01.DEPT_EMP.Group_1 has completed successfully.
success	I81A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 4:08:51 PM	The job ARYJOB / JOB00756 has ended on RS01 with MAXCC=0.
success	I81A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 4:30:37 PM	Recovery plan generation for ALA23SMP has completed successfully.
success	I81A on RS01	(unsaved)	Recovery plan generation	Jun 18, 2006 4:34:08 PM	Recovery plan generation for ALA23SMP has completed successfully.

Ready

Select Location

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Select the location that contains the objects you want to recover.
This advisor helps you recover data and dropped objects.
To start the recovery process, select a location. A location is a DB2 UDB for z/OS and OS/390 subsystem (or data sharing group).

Locations

- z/OS Subsystems
 - I71A
 - I71K
 - I81A**
 - I81E
- z/OS Data Sharing Groups
- z/OS Systems
- z/OS Sysplexes

Properties

Name	Value
Active?	true
DB2-established stored procedures address...	I81.ASPAS
Database services address space	I81.ADBM1
Distributed data facility address space	I81.ADIST
Host name	rs01.rocketsoftware.com
IP address	10.1.1.107
Mode	NFM
Port number	3800
Resync port number	3801
Subsystem ID	I81A
System services address space	I81.AMSTR
Version	DSN08015

Status

Name	Value
DB2 restart RBA	0000C5304000
DB2 restart time	18:33:01 JUN 15, 2006
Last checkpoint	10:09:08 JUN 16, 2006
Log copy 1	I81A.LOGCOPY1.DS03 IS 6% FULL
Log copy 2	I81A.LOGCOPY2.DS03 IS 6% FULL
Log high offloaded RBA	0000C586FFFF

◀ Back Next ▶ Save... Close Help

Expand Grouper Sets from Object Tree

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

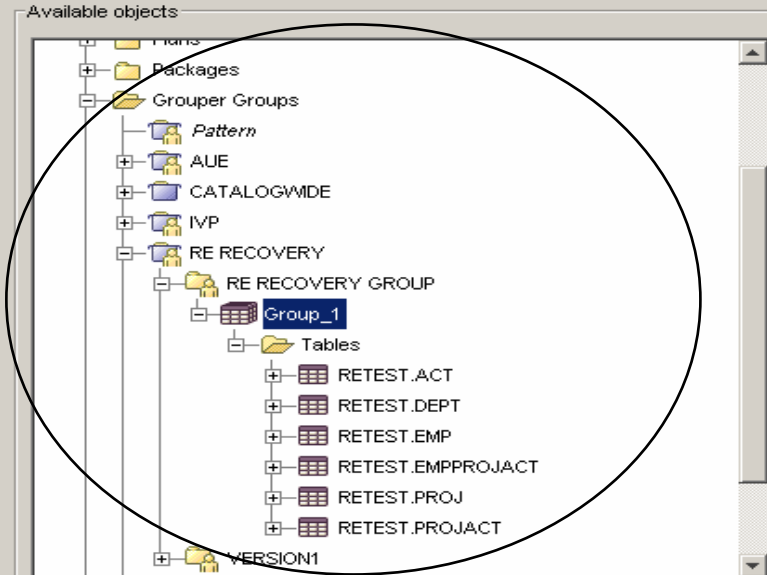
Select the objects that you want to recover.

Select the objects you wish to recover in the list of available objects, and click the right arrow to add them to the list of objects to recover. The list at the bottom of the window will help you by showing detailed properties for the objects you click on.

In order to define a pattern matching multiple objects, select the node labeled "Pattern" under the type of object you wish to select, then click the right arrow. A window will display where you can type the specific pattern values.

The objects you select, as well as objects related to them, will be recovered. For example, if you select a table space, the tables it contains will also be recovered. On the Recovery Plan page, you should carefully review the complete list of objects under Recovered Objects in the Recovery plans tree.

Available objects



Selected objects

Properties

Name	Value
Created timestamp	2007-02-15-09.02.40.347623
Creator	PDDAVI
Group ID	682
Remarks	
Version ID	186

◀ Back Next ▶ Save... Close Help

Select Objects from Object Tree

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Select the objects that you want to recover.

Select the objects you wish to recover in the list of available objects, and click the right arrow to add them to the list of objects to recover. The list at the bottom of the window will help you by showing detailed properties for the objects you click on.

In order to define a pattern matching multiple objects, select the node labeled "Pattern" under the type of object you wish to select, then click the right arrow. A window will display where you can type the specific pattern values.

The objects you select, as well as objects related to them, will be recovered. For example, if you select a table space, the tables it contains will also be recovered. On the Recovery Plan page, you should carefully review the complete list of objects under Recovered Objects in the Recovery plans tree.

Available objects

- Tables
- Packages
- Grouper Groups
 - Pattern
 - AUE
 - CATALOGWIDE
 - IVP
 - RE RECOVERY
 - RE RECOVERY GROUP
 - Group_1
 - Tables
 - RETEST.ACT
 - RETEST.DEPT
 - RETEST.EMP
 - RETEST.EMPPROJECT
 - RETEST.PROJ
 - RETEST.PROJECT

- VERSION1

Selected objects

- Table RETEST.ACT
- Table RETEST.PROJECT
- Table RETEST.DEPT
- Table RETEST.EMP
- Table RETEST.EMPPROJECT
- Table RETEST.PROJ

Properties

Name	Value
------	-------

◀ Back Next ▶ Save... Close Help

Select Point in time to recover to – current

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Select the point in time to which you want to recover.

Select one of the available options for the point in time to which to recover.

Point in time

Current

Select this option when you want to recover data that is logically consistent, but has a physical error or other corruption in how it is stored.

Timestamp 2006-06-19 10.57.12 000000 ... (explicit)

Select this option when you want to recover previous versions of the selected objects, or when data has become logically inconsistent and you can identify a prior timestamp when the data was consistent.
Use the browse button for help in selecting the timestamp.

Log RBA 0000C5E5B24A ... (explicit)

Select this option when data has become logically inconsistent and you can identify a prior log RBA when the data was consistent.
Use the browse button for help in selecting the log RBA.

◀ Back Next ▶ Save... Close Help

Or Timestamp and click on (explicit)

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Select the point in time to which you want to recover.

Select one of the available options for the point in time to which to recover.

Point in time

Current

Select this option when you want to recover data that is logically consistent, but has a physical error or other corruption in how it is stored.

Timestamp: 2006-06-18 17.53.12 000000 ... (explicit)

Select this option when you want to recover previous versions of the selected objects, or when data has become logically inconsistent and you can identify a prior timestamp when the data was consistent.
Use the browse button for help in selecting the timestamp.

Log RBA: 0000C58E6288 ... (explicit)

Select this option when data has become logically inconsistent and you can identify a prior log RBA when the data was consistent.
Use the browse button for help in selecting the log RBA.

◀ Back Next ▶ Save... Close Help

Click Refresh, select Quiet Time Timestamp, OK

Select from

Object definition levels Recovery history events Quiet times

Scanned ranges

Start Timestamp	End Timestamp
2006-06-18-08.36.12.000000	2006-06-18-09.36.12.000000
2006-06-18-17.11.32.000000	2006-06-18-18.11.32.000000

Run New Report...

Quiet times found

Start Timestamp	End Timestamp	Start LRSN	End LRSN
2006-06-18-17.11.32.000000	2006-06-18-17.33.28.000000	0000C5A30A4C	0000C5A95AC8
2006-06-18-17.33.36.000000	2006-06-18-17.35.23.000000	0000C5AD2750	0000C5B35539
2006-06-18-17.35.28.000000	2006-06-18-17.40.09.000000	0000C5B6E150	0000C5BD1539
2006-06-18-17.40.14.000000	2006-06-18-17.43.17.000000	0000C5C0D8D0	0000C5C378AB
2006-06-18-17.43.17.000000	2006-06-18-17.46.18.000000	0000C5C39B79	0000C5C40448
2006-06-18-17.46.19.000000	2006-06-18-17.49.20.000000	0000C5C42831	0000C5C490C4
2006-06-18-17.49.20.000000	2006-06-18-18.11.32.000000	0000C5C4A388	0000C5C610FA

Filter... Refresh

OK Cancel Help

or 'Recovery history events', refresh, select

Select Point in Time

Select from

Object definition levels Recovery history events Quiet times

Events

Start RBA	Timestamp	Database	Space	Space Type	Type	Secondary Type	Share Level	DSNUM	File Sequence N
0000C5CCABEC	2006-06-18-18.48.43.100760	DEMODB04	DEMOTS04	T	F		R	0	0
0000C5D0E000	2006-06-18-18.48.49.235404	DEMODB04	DEMOTS04	T	I		R	0	0
0000C5D17646	2006-06-18-18.51.50.536655	DEMODB04	DEMOTS04	T	Q	W		0	0

Show events that are not points of consistency

Filter... Refresh

OK Cancel Help

Generate and Review chose of recovery plans

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Generate and execute a plan to recover the selected objects.

Click Generate to generate one or more recovery plans for the selected objects.
 After generating, you can select a plan to review its details.
 If it has been some time since you generated the recovery plans, you can click Validate to check whether the external resources required by the plans (such as image copy data sets) are still available.
 Finally, select the plan you wish to use for recovery, and click Run to run it. To instead view or edit the JCL generated to run the plan, click View JCL.

Recovery plans

Generate...

Plans

- Plan 1: Using DSN1 COPY and RECOVER LOGONLY (cost = 24)
- Plan 2: Using DSN1 COPY of IC and redo SQL (cost = 24)
- Plan 3: Using RECOVER (cost = 24)
 - Job 1
 - Check Status
 - START DATABASE
 - Check Status
 - RECOVER
 - Table space DSNDB04.ACTA1X3Z
 - Table space DSNDB04.DEPT
 - Table space DSNDB04.EMPA1E2V
 - Table space DSNDB04.EMPP1ZH1
 - Table space DSNDB04.PROJ1RO0
 - Table space DSNDB04.PROJ1V9W
 - REBUILD INDEX
 - CHECK DATA
 - COPY
 - Check Status
 - START DATABASE
 - Check Status
 - Plan 4: Using RECOVER to IC and redo SQL (cost = 24)
 - Recovered Objects

Properties

Name	Value
Plan cost	24
Plan name	Using DSN1 COPY and RECOVER LOGONLY

Validate **Run** **View JCL...**

◀ Back Next ▶ Save... Close Help

Run or Export to z/OS file

Recovery plan: Plan 1: Using DSN1COPY and RECOVER LOGONLY (cost = 24)

Statements

```

//RCVRO010 EXEC PGM=DSN1COPY,COND=(4,LT),
// PARM='CHECK,PAGESIZE(4K),FULLCOPY'
//STEPLIB DD DISP=SHR,DSN=I81A.SDSNEXIT
// DD DISP=SHR,DSN=RSRTE.DSN.V810.SDSNLOAD
//SYSUT1 DD DISP=SHR,DSN=PDDAVI.I81A.PROJ1V9W.D070215.T212641
//SYSUT2 DD DISP=SHR,DSN=I81A.DSNDBC.DSNDBO4.PROJ1V9W.I0001.A001
//SYSPRINT DD SYSOUT=*
/*
//RCVRO010 EXEC PGM=DSNUTILE,PARM=(I81A),COND=(4,LT)
//STEPLIB DD DISP=SHR,DSN=I81A.SDSNEXIT
// DD DISP=SHR,DSN=RSRTE.DSN.V810.SDSNLOAD
//SYSPRINT DD SYSOUT=*
//SYSIN DD *

RECOVER
  TABLESPACE DSNDBO4.ACTALX3Z
  TABLESPACE DSNDBO4.DEPT
  TABLESPACE DSNDBO4.EMPALR2V
  TABLESPACE DSNDBO4.EMPPLZH1
  TABLESPACE DSNDBO4.PROJ1R00
  TABLESPACE DSNDBO4.PROJ1V9W
  LOGONLY

/*
/*
//RBLD0011 EXEC PGM=DSNUTILE,PARM=(I81A),COND=(4,LT)
//STEPLIB DD DISP=SHR,DSN=I81A.SDSNEXIT
// DD DISP=SHR,DSN=RSRTE.DSN.V810.SDSNLOAD
//SYSPRINT DD SYSOUT=*
//UTPRINT DD SYSOUT=*
//SYSOUT DD UNIT=SYSALLDA,SPACE=(CYL,(1,1))
//SYSIN DD *

```

copy data sets) are

OVER LOGONLY

View JCL...

Run Export... Close Help

Close Help

Example 2 – Select a single table, relationship unknown

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Select the objects that you want to recover.

Select the objects you wish to recover in the list of available objects, and click the right arrow to add them to the list of objects to recover. The list at the bottom of the window will help you by showing detailed properties for the objects you click on.

In order to define a pattern matching multiple objects, select the node labeled "Pattern" under the type of object you wish to select, then click the right arrow. A window will display where you can type the specific pattern values.

The objects you select, as well as objects related to them, will be recovered. For example, if you select a table space, the tables it contains will also be recovered. On the Recovery Plan page, you should carefully review the complete list of objects under Recovered Objects in the Recovery plans tree.

Available objects

- I81A
 - Storage Groups
 - Databases
 - Table Spaces
 - Tables
 - Pattern
 - RETEST
 - ACT
 - DEPT
 - EMP
 - EMPPROJECT
 - PROJ
 - PROJACT
 - Plans
 - Packages
 - Groupers Groups
 - Pattern
 - Δ IIF

Selected objects

- Table RETEST.DEPT

Properties

Name	Value
Altered timestamp	2007-02-15-08.59.09.869515
Created by	PDDAVI
Created timestamp	2007-02-15-08.58.45.336151
DBID	4
Database	DSNDB04
ORID	390

Navigation buttons: Back, Next, Save..., Close, Help

Example 2 – Allow Grouper to check relationships

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Generate and execute a plan to recover the selected objects.

Click Generate to generate one or more recovery plans for the selected objects.
After generating, you can select a plan to review its details.
If it has been some time since you generated the recovery plans, you can click Validate to check whether the external resources required by the plans (such as image copy data sets) are still available.
Finally, select the plan you wish to use for recovery, and click Run to run it. To instead view or edit the JCL generated to run the plan, click View JCL.

Recovery plans
Generate...
Plans

Properties
Name Value

Recovery Plan Generation Options

Restricted Objects | RECOVER Parameters | COPY Parameters | **Grouper Parameters**

Information from a DB2 Grouper set is used to determine additional tables related to those that are directly selected.

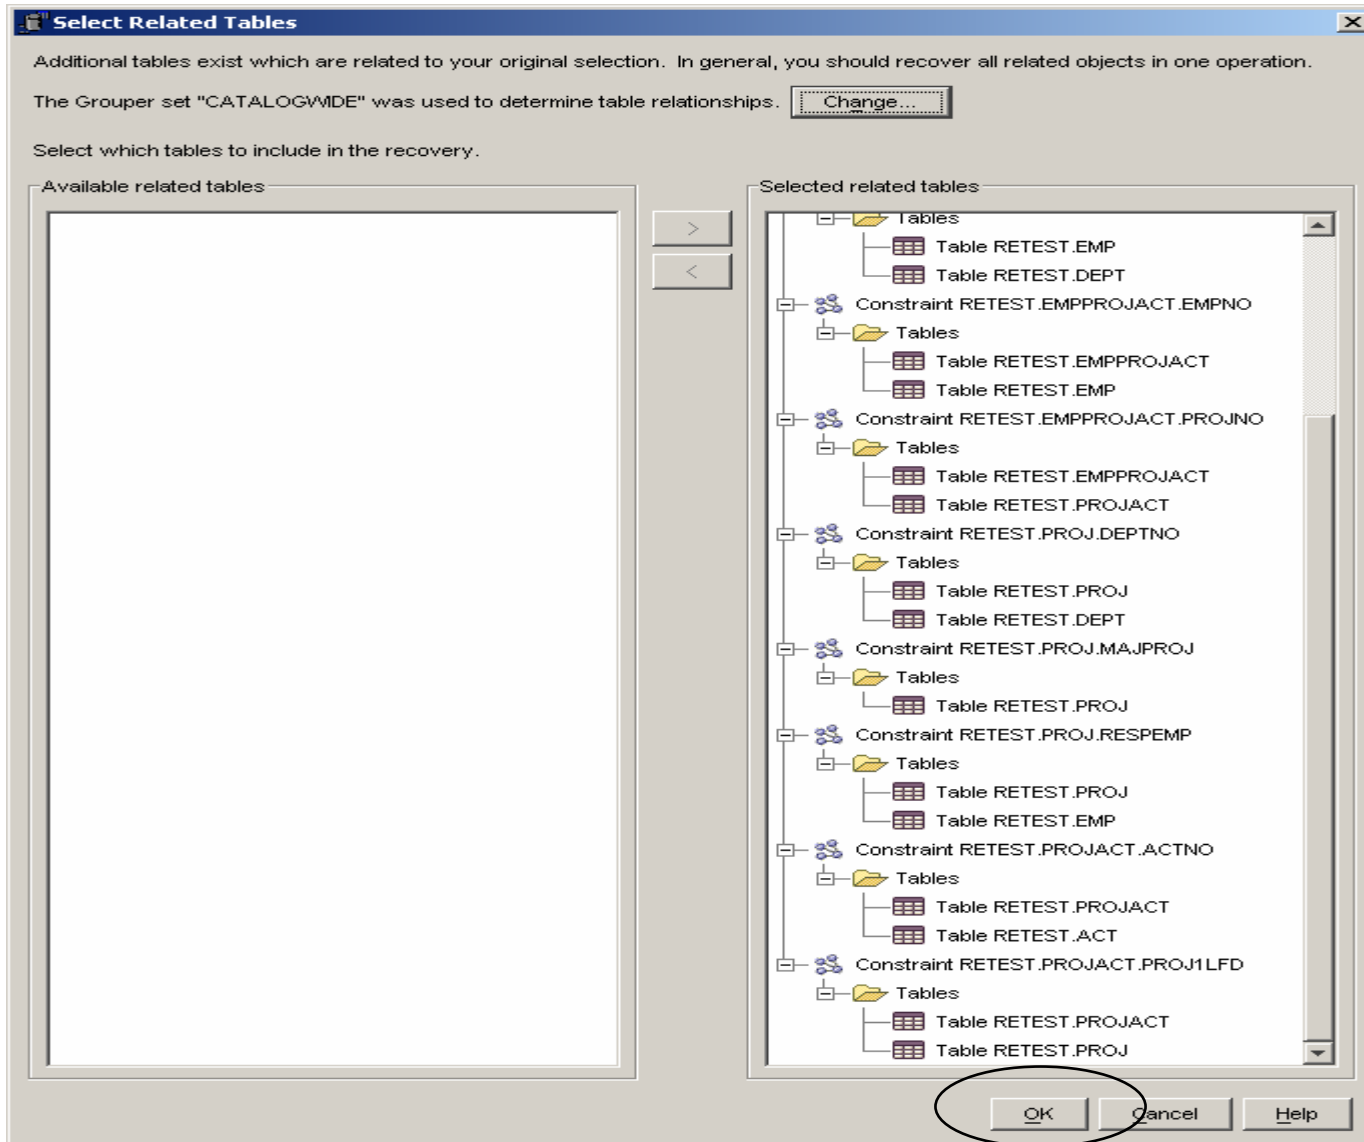
Grouper set name CATALOGWIDE

OK Cancel Help

Validate Run View JCL...

◀ Back Next ▶ Save... Close Help

RE added related tables from Grouper for recovery



recovery plans that automatically contain related objects

Recovery Advisor

1. Location
2. Objects
3. Point in Time
4. Recovery Plan

Generate and execute a plan to recover the selected objects.

Click Generate to generate one or more recovery plans for the selected objects. After generating, you can select a plan to review its details. If it has been some time since you generated the recovery plans, you can click Validate to check whether the external resources required by the plans (such as image copy data sets) are still available. Finally, select the plan you wish to use for recovery, and click Run to run it. To instead view or edit the JCL generated to run the plan, click View JCL.

Recovery plans

Generate...

Plans

- Plan 1: Using DSN1COPY and RECOVER LOGONLY (cost = 24)
- Plan 2: Using DSN1COPY of IC and redo SQL (cost = 24)
- Plan 3: Using RECOVER (cost = 24)
 - Job 1
 - Check Status
 - START DATABASE
 - Check Status
 - RECOVER
 - Table space DSNDDB04.ACTA1X3Z
 - Table space DSNDDB04.DEPT
 - Table space DSNDDB04.EMPA1E2V
 - Table space DSNDDB04.EMPP1ZH1
 - Table space DSNDDB04.PROJ1ROD
 - Table space DSNDDB04.PROJ1V9W
 - REBUILD INDEX
 - CHECK DATA
 - COPY
 - Check Status
 - START DATABASE
 - Check Status
- Plan 4: Using RECOVER to IC and redo SQL (cost = 24)
- Recovered Objects

Properties

Name	Value
Plan cost	24
Plan name	Using DSN1COPY and RECOVER LOGONLY

Validate Run View JCL...

◀ Back Next ▶ Save... Close Help

DB2 Grouper – Publications

- Redbooks
 - ▶ IBM DB2 Data Archive Expert for z/OS, SG24-7080-00
 - ▶ DB2 for z/OS Stored Procedures: Through the CALL and Beyond, SG24-7083
 - ▶ Cross-Platform Stored Procedures: Building and Debugging, SG24-5485-01
- Publications
 - ▶ DB2 Data Archive Expert Program Directory, GI10-8511
 - ▶ DB2 Data Archive Expert User's Guide and Reference, SC18-7344
 - ▶ DB2 Grouper User's Guide, SC18-7409
 - ▶ DB2 Grouper Program Directory, GI10-8569-00
 - ▶ z/OS V1R4.0 MVS Planning Workload Management, SA22-7602
 - ▶ IBM DB2 Recovery Expert for z/OS User's Guide, SC18-9822-00

