



ibm.com/db2/labchats

Data Management

Simplifying Administration and Development with IBM Data Studio

August 27, 2009

ibm.com/db2/labchats

Executive's Message



Sal Vella

**Vice President, Development,
Distributed Data Servers and Data Warehousing**

IBM



IDUG Update



Elizabeth Moore

**President Elect,
Board of Directors**

International DB2 Users Group (IDUG)



***Become an IDUG Member at no cost
and receive many benefits:***

- Online Content with Podcasts/Webcasts
- Code Place and Technical Library
- On-demand technical presentations and archived proceedings from previous IDUG conferences
- IDUG Solutions Journal and eBulletin
- Global and Regional Events
- Discounted books from IBM Press
- Regional Users Group information and resources
- Networking, discussion forums, and the DB2-L list service
- and more...



Visit www.IDUG.org



Access these benefits and more with IDUG Membership

Sign up for free at
www.idug.org



Visit www.IDUG.org



Featured Speaker

Deb Jenson



**Lead Product Manager
Data Studio**

IBM



Agenda

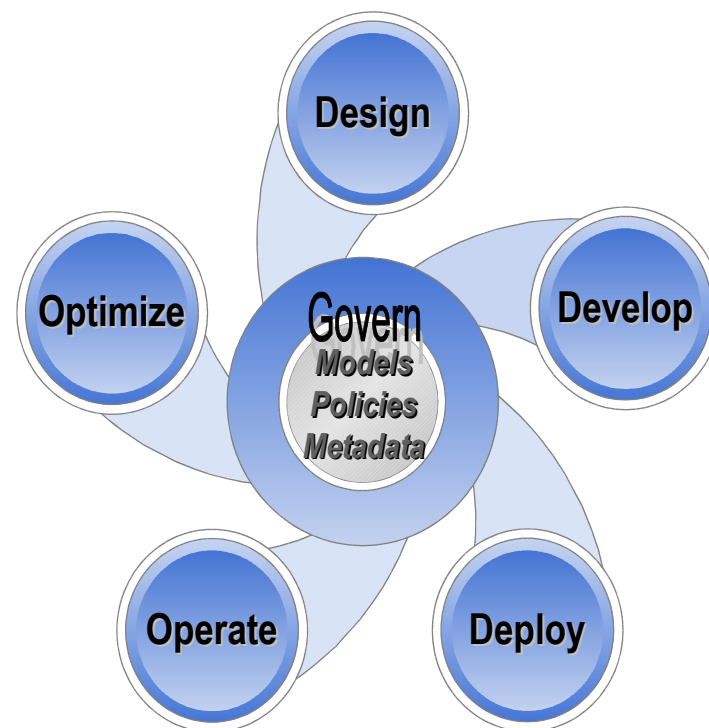
- What is Data Studio?
- The “How To” Series
- Questions



What is Data Studio?

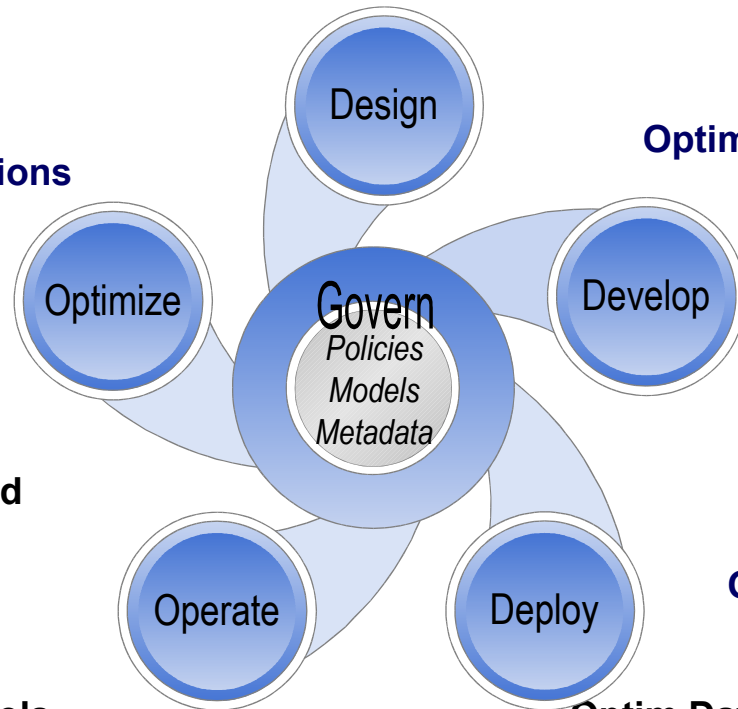
• An integrated Data Management environment

- **Providing end-to-end data lifecycle management**
- **Facilitating cross-organizational collaboration by providing an integrated set of tools for different roles**
 - Data Architect
 - Developer
 - Tester
 - DBA
 - Data Steward



The broadest range of capabilities for managing the value of your data throughout its lifetime

InfoSphere Data Architect



Optim Data Growth Solutions

Optim Development Studio

Optim Query Tuner

Optim Test Data Management

Optimize

Develop

Optim Data Privacy Solutions

Optim Query Workload Tuner

Govern
Policies
Models
Metadata

Optim pureQuery Runtime

DB2 Performance Expert and Extended Insight

Operate

Deploy

Optim Database Administrator

DB2 for z/OS Tools

DB2 Audit Management Expert
Database Encryption Expert

Data Studio – No Charge Offering

** Heterogeneous support*



Where does **free** Data Studio fit in?

- **A base product containing critical functionality from most of the charged Optim offerings**

Data Studio

Design

ER Diagramming
Data Distribution Viewer

Develop

Integrated query editor
SQL & Java Routine + Debugger
XML, XML Schema Editors
Data Web Services
Visual Explain
Export / Import Data
Generate DDL

Administer

Database Object Management
Instance Management
Database Management
Privilege Management
Configuration Management
Back / Restore Database
Reorg and Other Utilities
Run DB2 Commands

Monitor & Tune

Alerts
Problem Determination
Replication Monitor
Recommendations
Limited Historical Information
Visual Explain

Featured Speaker



Manas Dadarkar

Technical Lead
Data Studio

IBM



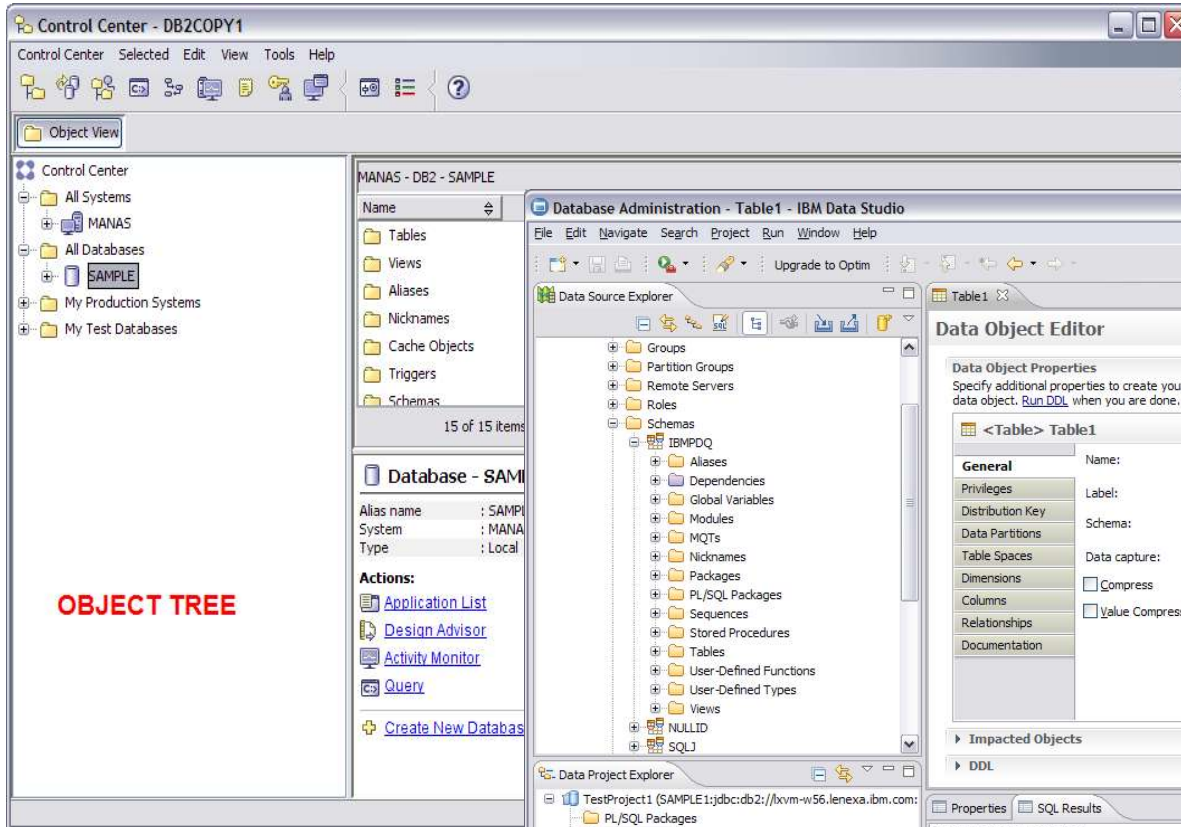
How to download and install free Data Studio ...

- **Download Link:** <http://www-01.ibm.com/software/data/optim/data-studio/>
- **Two flavors**
 - Standalone
 - Smaller footprint (~200 MB)
 - Simple installer
 - Subset of functionality compared to the IDE version
 - Work artifacts like workspaces can be shared with other Optim offerings
 - Integrated Development Environment (IDE)
 - Larger footprint (~800 MB)
 - Powerful Installation Manager allows shell sharing with other Optim offerings

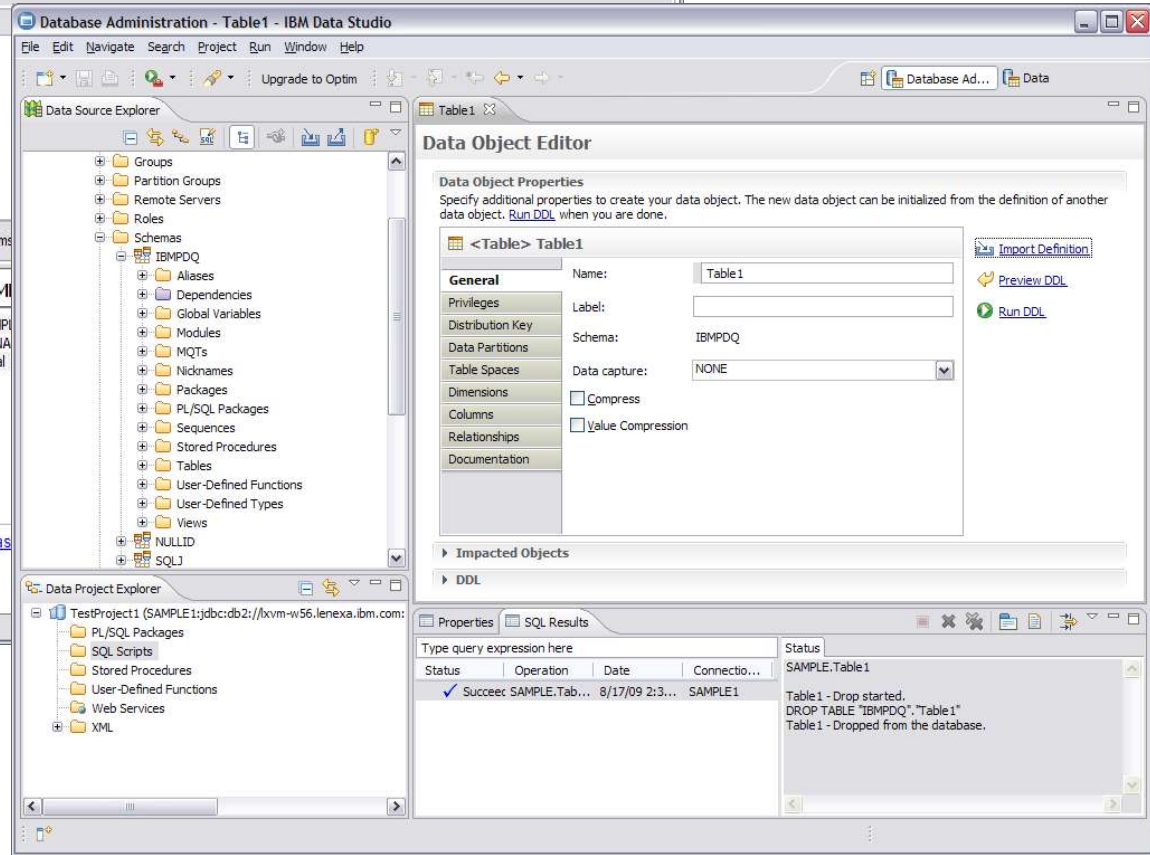


How to navigate

Control Center

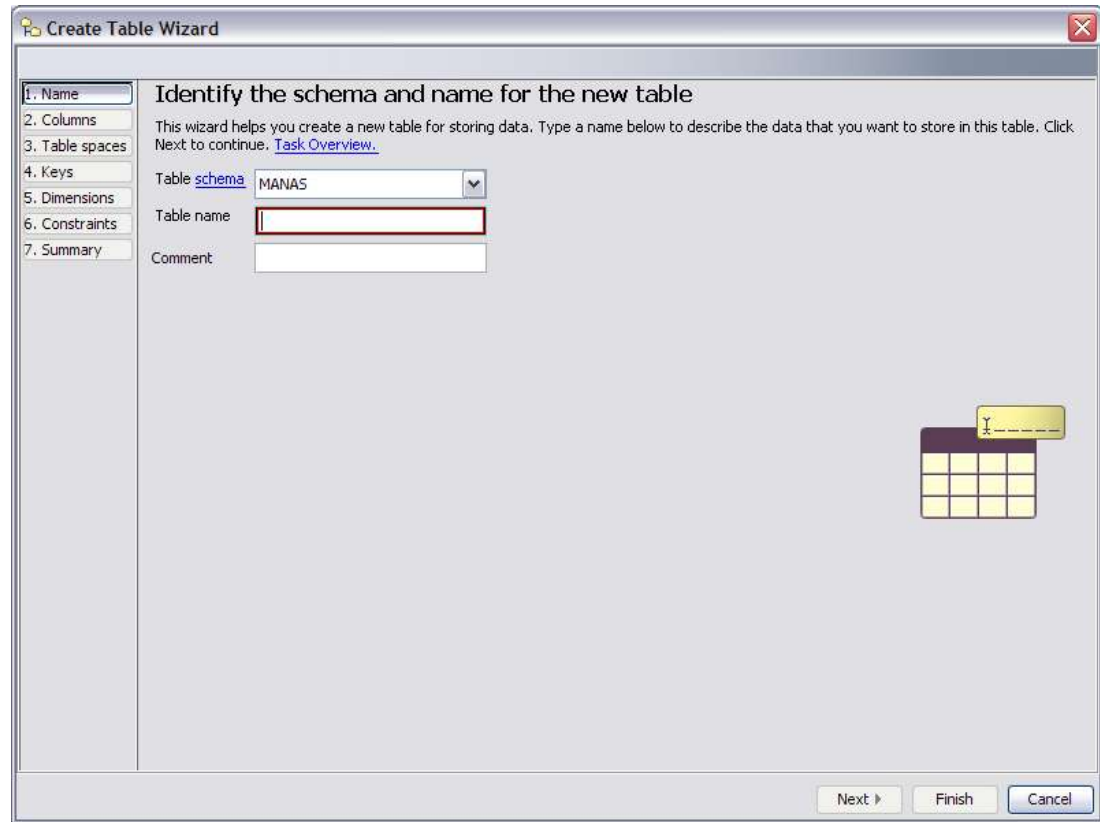
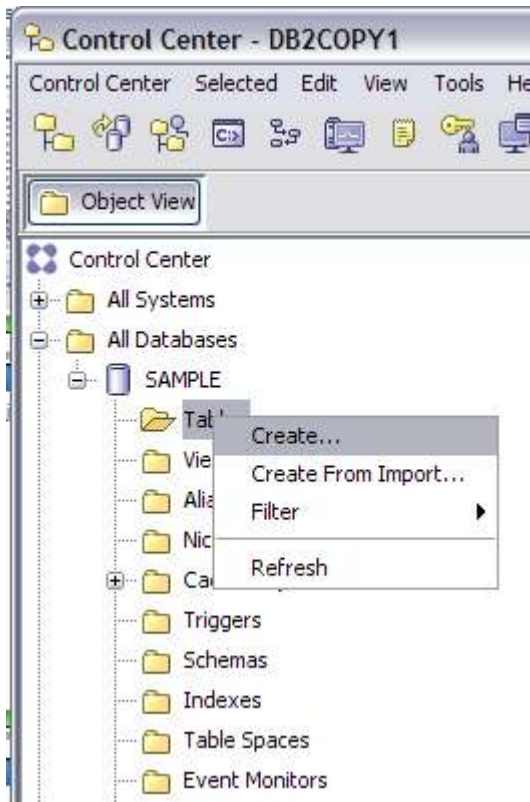


Data Studio



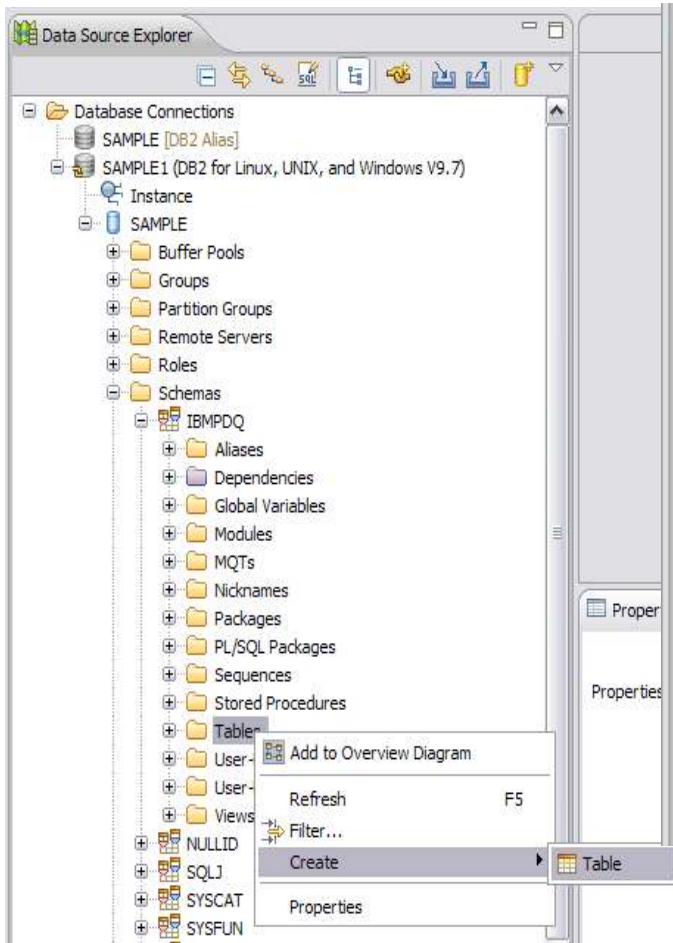
How to manage objects ...

Control Center

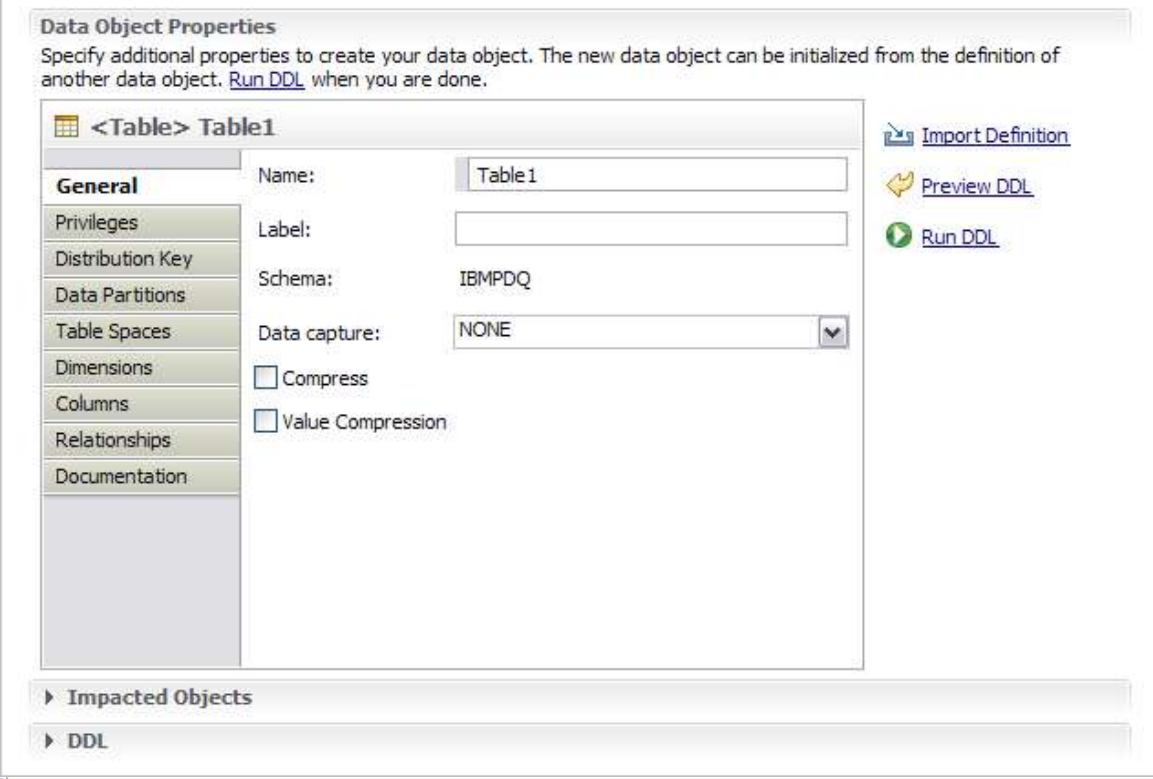


Data Source Explorer/Data Object Editor

Data Studio



Data Object Editor



Tips about Navigation

- **Need to work with large amounts of objects**
 - Use the Object List View (flat view)
 - Easily view related objects
 - Customize object list including sorting and filtering on multiple columns



Data Studio Object Lists

Store and Share Connections

Customize Object List

Database De...

Row ...	Perc...	Log ...	Parti...	Reg...	Inde...	LOB ...
DEJENSON	EXPLAIN_ARG...	17	0		USER...	
DEJENSON	EXPLAIN_OBJE...	3	0		USER...	
DEJENSON	EXPLAIN_OPE...	4	0		USER...	
DEJENSON	EXPLAIN_PRE...		0		USER...	
DEJENSON	EXPLAIN_STRE...	4	0		USER...	
DEJENSON	EXPLAIN_DIAG...	0	0		USER...	
DEJENSON	EXPLAIN_DIAG...	0	0		USER...	
DEJENSON	EMPLOYEE		0		USER...	
IBMPDQ	CONNECTION		0		USER...	
IBMPDQ	CONNECTION...		0		USER...	
IBMPDQ	GROUP		0		USER...	
IBMPDQ	GROUP_CONN...		0		USER...	
SALES	STATE_TAX	2	0		USER...	
SALES	PRODUCT		0		USER...	
SALES	PRODUCT_IM...		0		USER...	
SALES	CUSTOMER	7	0		USER...	
SALES	CREDITCARD		0		USER...	

All Showing 154 of 154 items

Prope Tasks Probl Error Model SQL R Book Conso Impa



How to manage data ...

Control Center

Open Table - EMPLOYEE

MANAS - DB2 - SAMPLE - MANAS.EMPLOYEE

Edits to these results are performed as positioned UPDATES and DELETES. Use the Tools Settings notebook to change the form of editing.

EMPNO	FIRSTNAME	MIDINIT	LASTNAME
000010	CHRISTINE	I	HAAS
000020	MICHAEL	L	THOMPSON
000030	SALLY		
000050	JOHN		
000060	IRVING		
000070	EVA		
000090	EILEEN		
000100	THEODO		
000110	VINCENZ		
000120	SEAN		
000130	DELORES		
000140	HEATHER		
000150	BRUCE		
000160	ELIZABE		
000170	MASATO		
000180	MARILYN		
000190	JAMES		
000200	DAVID		
000210	WILLIAM		
000220	JENNIFE		
000230	JAMES		
000240	SALVATO		
000250	DANIEL	S	SMITH
000260	SYBIL	P	JOHNSON
000270	MARIA	L	PEREZ

Filter

MANAS - DB2 - SAMPLE - MANAS.EMPLOYEE

Locate | **Advanced**

Column	Comparison	Values
EMPNO	LIKE	
FIRSTNAME	LIKE	
MIDINIT	LIKE	
LASTNAME	LIKE	
WORKDEPT	LIKE	

Clear

Maximum number of rows to return:

OK Cancel Delete Help

Commit Roll Back Filter Fetch More Rows

Automatically commit updates

42 row(s) in memory

Close Help



Data Studio Table Editor

The screenshot shows the IBM Data Studio Table Editor interface. The main window displays a table with the following columns: ORDERID [INTEGER], CUSTOMERID [INTEGER], TIMESTAMP [DATE], ORIGINOFSALE [VARCHAR(128)], and NUMB. The table contains 14 rows of data. A context menu is open over the first row (ORDERID 5000), showing options: Revert, Refresh (F5), Edit Value, Set Null, Insert Row, Delete Row (Delete), and Save (Ctrl+S).

ORDERID [INTEGER]	CUSTOMERID [INTEGER]	TIMESTAMP [DATE]	ORIGINOFSALE [VARCHAR(128)]	NUMB
5000	100	2007-08-23		3
5001	104	2007-08-23		5
5002	104	2007-08-23		2
5003	104	2007-08-23		3
5004	103	2007-08-23		2
5005	103	2007-08-23		2
5006	103	2007-08-23		3
5007	103	2007-08-23		6
5008	103	2007-08-23		2
5009	103	2007-08-23		2
5010	103	2007-08-23		5
5011	100	2007-08-23		40
5012	105	2009-05-01		2
5013	106	2009-05-01		2
5014	105	2009-05-01		2

Below the table, the 'Task' pane shows a list of operations:

Status	Operation	Date	Conn
✓	Succesec CALL SYSPR...	5/28/09 7:2...	JKENT
✓	Succesec Run SQL	5/28/09 8:0...	SAMPL
✓	Succesec Return All R...	5/29/09 11:...	JKENT

At the bottom right, a secondary context menu is visible with options: Copy Row(s), Save, Export, Print, and Convert Row(s) To Hexadecimal.

How to run Utilities

Control Center

The screenshot shows the IBM Control Center interface. On the left, a table named 'CATALOG' is selected in the 'MANAS - DB2 - SAMPLE - Tables' view. A context menu is open over the table, with the 'Load...' option highlighted. The 'Load Wizard - CATALOG' dialog box is displayed in the foreground, showing the 'Files' step. The dialog title is 'Load Wizard - CATALOG'. It has a list of steps on the left: 1. Type, 2. Files (selected), 3. Columns, 4. Performance, 5. Recovery, 6. Options, 7. Schedule, 8. Summary. The main area of the dialog contains the text: 'Choose to replace or keep the original table data.' Below this, it says: 'You can choose to append the input data to the end of the table or replace all of the data in the table. [Task Overview.](#)' There are two radio buttons: 'Append data to table' (which is selected) and 'Replace table data'. There is also a checkbox for 'Allow read access durir'. At the bottom of the dialog are 'Next >', 'Finish', and 'Cancel' buttons.



Data Source Explorer/Task Assistants

Data Studio

- Tables
 - ACT
 - CATALOG**
 - CL_SCHED
 - CUSTOMER
 - DEPARTMENT
 - EMP_PHOTO
 - EMP_RESUME
 - EMPLOYEE
 - EMPMDC
 - EMPPROJECT
 - IN_TRAY
 - INVENTORY
 - ORG
 - PRODUCT
 - PRODUCTSUPPLIER
 - PROJECT
 - PROJECT
 - PURCHASEORDER
 - SALES
 - STAFF
 - STAFFG
 - SUPPLIERS
- User-Defined Functions
- User-Defined Types
- Views

- Data
- Drop
- Add to Overview Diagram
- Generate DDL...
- Update Statistics
- Analyze Impact...
- Value Distributions
- Refresh
- High Performance Unload...
- Export Table...
- Import Table...
- Load...
- Reorg Table...
- Reorg Index...
- Set Integrity...
- Alter
- Copy
- Paste...

Load Table CATALOG

Load Table CATALOG

Use the load utility to move large quantities of data into newly created tables or into tables that already contain data. The load utility writes formatted pages directly into the database and is faster than the import utility.

Settings
Specify any additional settings to use. Click Run when you are done.

[Preview Command](#) Run

Files

Mode

Columns

Recovery

Format Options

Options

Specify input and output files

Most load operations will have at least one input or output file. You can find other minor file specifications on the Options tab.

Input file format:
Delimited Text (DEL) ▼

Input file location:
 Full path and file name of input files:
 Browse...



Data Studio – HADR Setup

The screenshot displays the IBM Data Studio interface for setting up High Availability Disaster Recovery (HADR) for a database named JKENT1. The main window is titled "Setup HADR JKENT1".

Setup HADR JKENT1

Before you can use the High Availability Disaster Recovery (HADR) feature, which provides complex logging, failover, and recovery for availability databases in the event of partial or complete site failures, you must initialize HADR. HADR protects against data loss by replicating data from a source database, called the primary database, to a target database, called the standby database. The standby database takes over in the event of a failure.

Settings
Specify any additional settings to use. Click Run when you are done.

Buttons: [Preview Command](#) [Run](#)

Primary Database

- Standby Database
- Copy Objects
- TCP/IP Parameters
- Client Reroute
- Synchronization
- Summary

Confirm the primary database selection

Verify that you want to set up a standby database for the following database. If you do not want to set up a standby database for this database, close the task assistant.

Database information

Primary database:	JKENT1
Host name :	IBM-64DAC294200.svl.ibm.com
DB2 instance:	DB2
Database state:	STANDARD
Logging type:	CIRCULAR
Infinite active log space:	NO

Database configuration

Configure the primary database for archive logging and index management. Also, if a backup of the primary database is required, you must specify the backup location.

The interface also shows a Data Project Explorer on the left with a tree view containing folders like "Change Management", "Instance", and "JKENT1", and sub-items like "Aliases", "Application Objects", "Buffer Pools", "Constraints", "Federated Stored Procedures", "Indexes", "MQTs", and "Nicknames".

Data Studio – Database Configuration

The screenshot displays the 'Configuration parameters' window in IBM Data Studio. The window title is 'Data - JKENT1 - IBM Optim Database Administrator'. The main area shows a table of configuration parameters for the 'JKENT1' instance. The table has columns for Name, Value, New Value, Automatic, and Immediate. The parameters are grouped into three categories: Applications, Deprecated, and Environment.

Name	Value	New Value	Automatic	Immediate
Applications				
AVG_APPLS	1	1	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
DATABASE_CONSISTE	YES	YES		<input checked="" type="checkbox"/>
DLCHKTIME	10000	10000		<input checked="" type="checkbox"/>
LOCKTIMEOUT	-1	-1		<input checked="" type="checkbox"/>
MAXAPPLS	40	40	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
MAXFILOP	32768	32768		<input checked="" type="checkbox"/>
MAXLOCKS	98	98	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
Deprecated				
APP_CTL_HEAP_SZ	256	256		
APPGROUP_MEM_SZ	30000	30000		
GROUPHEAP_RATIO	70	70		
LOGRETAIN	OFF	OFF		
MIN_DEC_DIV_3	NO	NO		
NUMSEGS	1	1		
USEREXIT	OFF	OFF		
Environment				
ALT_COLLATE				
CODEPAGE	1208	1208		
CODESET	UTF-8	UTF-8		
COLLATE_INFO				
COUNTRY	1	1		



How to do SQL Development ...

Control Center

Control Center
Control Center Selected Edit View Tools Help

Object View Command Editor 1 X

Commands Query Results Access Plan

Target SAMPLE Add

```
select * from department
```

A JDBC connection to the target has succeeded.
----- Commands Entered -----
select * from department

Results for a single query are displayed on the Query Results pane.
9 row(s) returned successfully.

Statement termination character ;

SQL Assist

Outline

- SQL statement properties
 - SELECT statement
 - FROM (Source table)
 - SELECT (Result columns)
 - WHERE (Row filter)
 - GROUP BY (Row groups)
 - HAVING (Group filter)
 - ORDER BY (Sort order)

Details

Available tables

- ABBYMAC
 - ADVISE_INDEX
 - ADVISE_INSTAN
 - ADVISE_MQT
 - ADVISE_PARTITI
 - ADVISE_TABLE
 - ADVISE_WORKL
 - CL_SCHED
 - DEPARTMENT
 - EMPLOYEE
 - EMP_ACT
 - EMP_PHOTO
 - EMP_SEQUENCE

Selected source tables

Table	Name
DEPARTMENT	
ABBYMAC.EMPLO...	EMPLOYEE

Join Tables...

SQL code

SQL validated

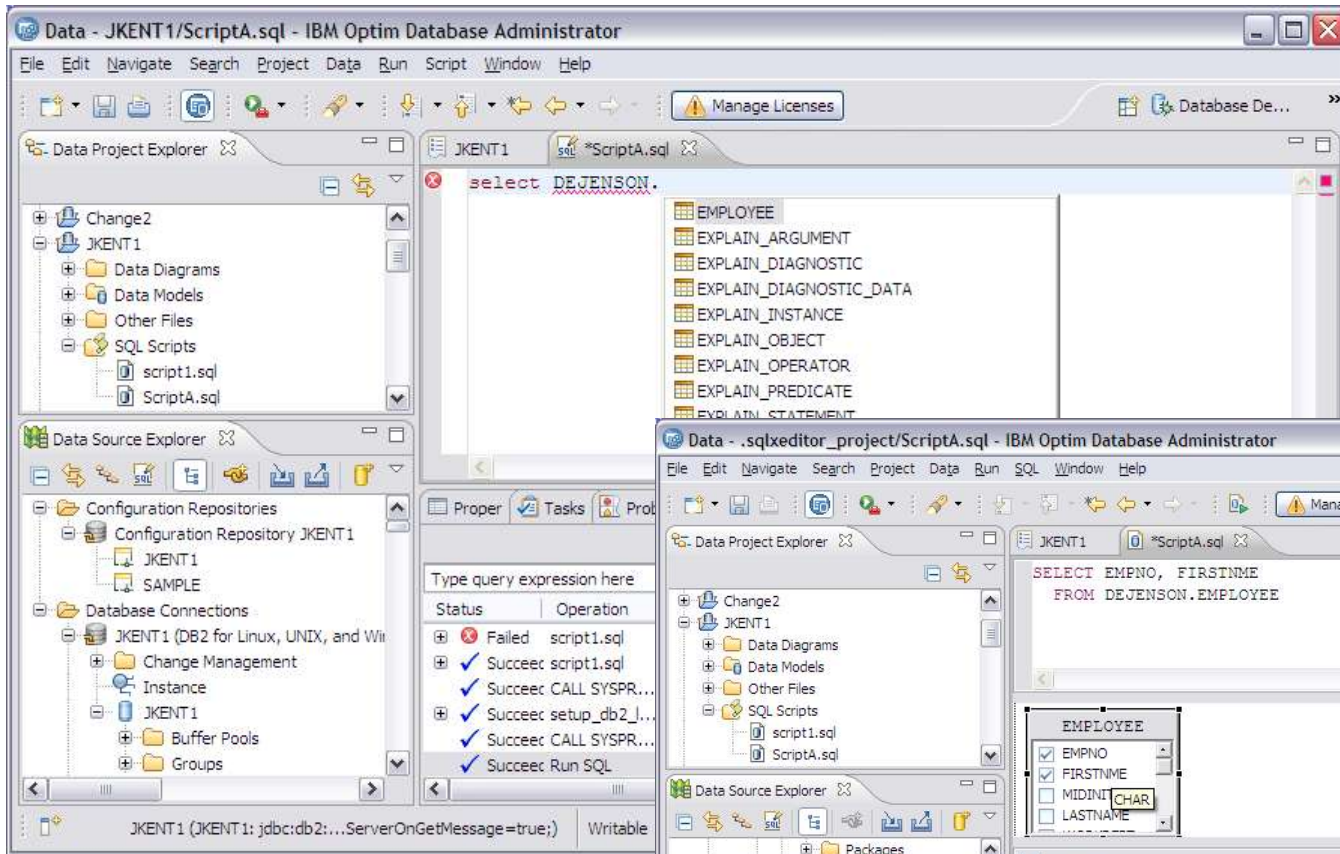
Clear Undo Edit Check Run

```
SELECT *
FROM DEPARTMENT, ABBYMAC.EMPLOYEE AS EMPLOYEE
```

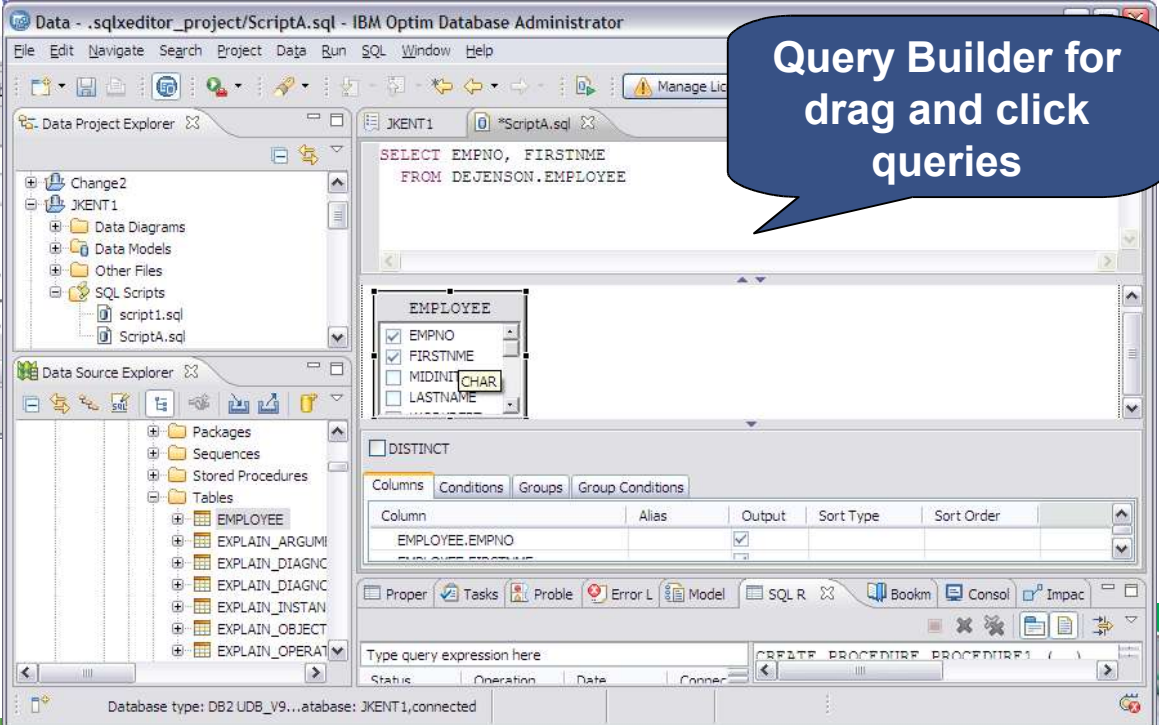
OK Cancel Help



Data Studio SQL Editor



Ctrl-Space for content assist



Query Builder for drag and click queries



Data Studio Stored Procedure Wizard

The screenshot shows the IBM Data Studio interface with the 'New Stored Procedure' wizard open. The 'Construct an SQL Statement' dialog box is the primary focus, displaying a list of available tables and a table for selected tables.

Available Tables:

- DEJENSON
- NULLID
- SALES
- SQLJ
- SYSCAT
- SYSFUN
- SYSIBM
- SYSIBMADM
- SYSIBMINTERNAL
- SYSIBMTS
- SYSPROC
- SYSSTAT
- SYSTOOLS

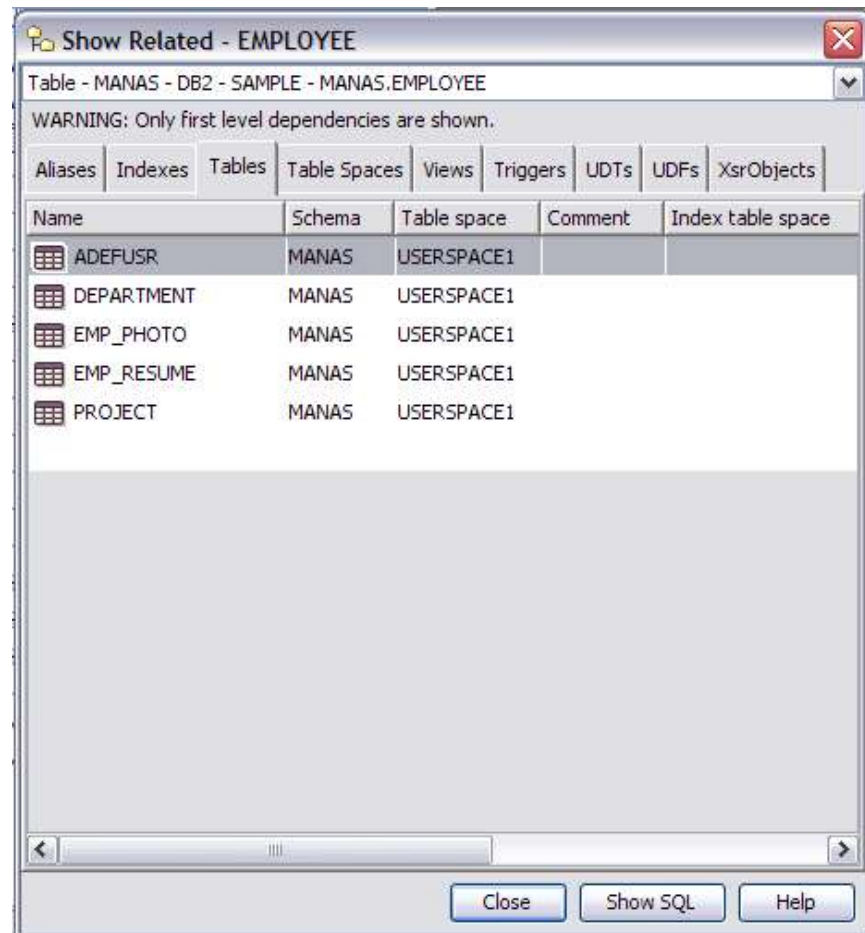
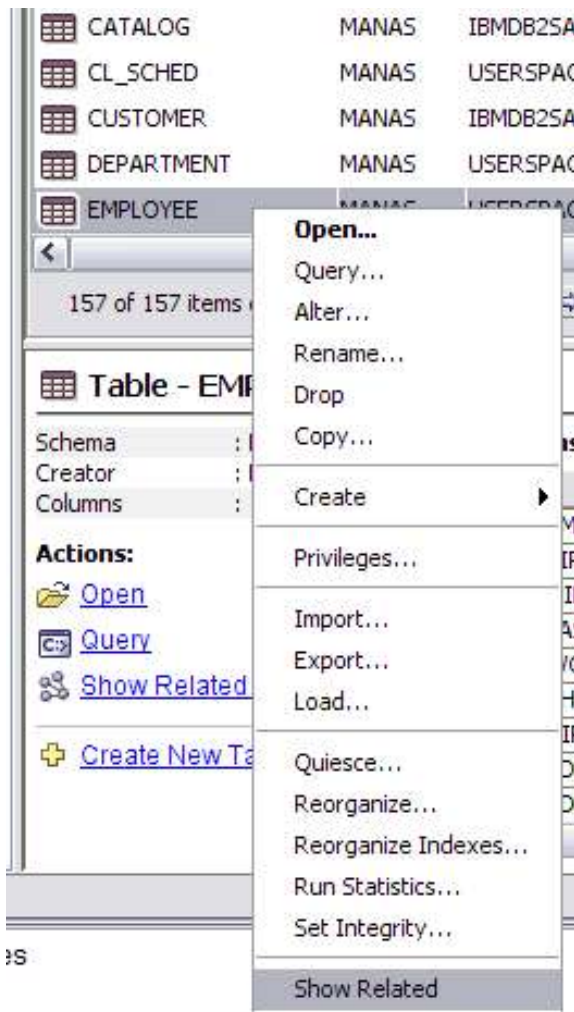
Selected Tables:

Tables	Alias
--------	-------

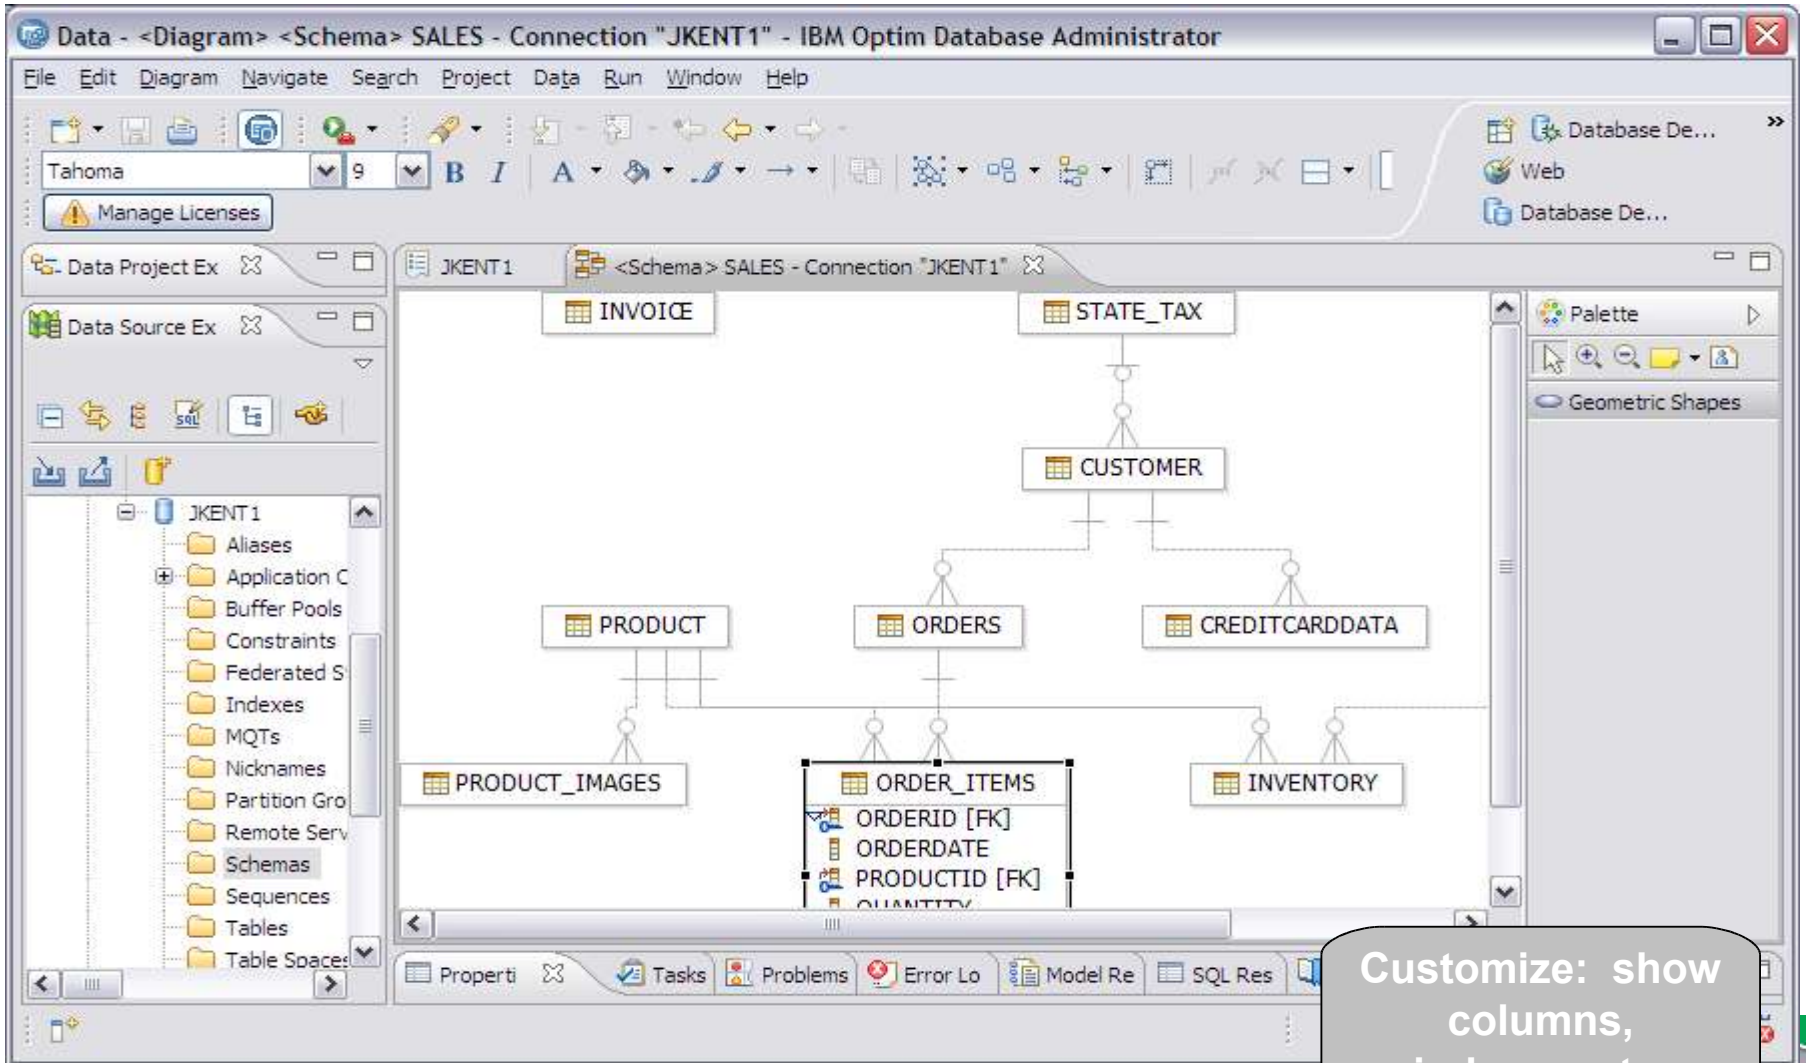
The dialog also includes navigation tabs for 'Tables', 'Columns', 'Joins', 'Conditions', 'Groups', and 'Order'. The 'Tables' tab is currently selected. The 'Result set' is set to 'One'.

How to find relationships/value distributions ...

Control Center



Data Studio Overview Diagram



Customize: show columns, indexes, etc.



Data Studio Column Values

Data - INVENTORY - IBM Optim Database Administrator

File Edit Navigate Search Project Data Run Window Help

Manage Licenses Database De...

Data Project Explorer

Data Source Explorer

Dependencies
 Federated Stored Pro
 Aliases
 MQTs
 Nicknames
 Packages
 Sequences
 Stored Procedures
 Tables
 CREDITCARDAT
 CUSTOMER
 INVENTORY
 INVOICE
 ORDER_ITEMS
 ORDERS
 PRODUCT
 PRODUCT_IMAGE
 STATE_TAX
 SUPPLIER
 SUPPLYORDERS

JKENT1 CUSTOMER INVENTORY

Table statistics (44 rows):

Field name	Data...	Values	Nulls	Mini...	Maxi...	Range	Total	Mean	Stan...	1st ...
PRODUCTID	Cate...	44	0							
SUPPLIERID	Cate...	44	0							
UNITCOST	Conti...	44	0	2	300	298	1,623	36,886	59,343	1...

Charts for selected fields Charts per page: all Show Frequencies in %

SUPPLIERID

Supplier Name	Supplier ID	Percentage
ALVIN-142	BIRKS-811	7%
BIRKS-811	BIRKS-811	14%
GORHA-05	GORHA-05	34%
ONEID-42	ONEID-42	5%
POOLE-22	POOLE-22	2%
REEDB-25	REEDB-25	2%
REVVE-55	REVVE-55	2%
ROGER-55	ROGER-55	5%
STIEF-778	STIEF-778	14%
THMAR-74	THMAR-74	14%
WALLA-4411	WALLA-4411	2%

Proper Tasks Proble Error L Model SQL R Bookm Consol Impac



How to configure Automatic Maintenance ...

Control Center

Configure Automatic Maintenance

1. Introduction
2. Type
3. Timing
4. Notification
5. Activities
6. Summary

Review the automatic maintenance settings

When you click Finish, the wizard updates the automatic maintenance configuration as you have specified. To change any of the settings, go back to the appropriate page in this wizard.

Database: ABBYMAC0 - DB2 - SAMPLE

Backup database (BACKUP)

Automation	Enabled
Notification	Enabled
Criteria	Balance database recoverability with performance
Media	File system
Location	D:\DB2\NODE0000\SQL00001\DB2AUTOBACKUPS;
Mode	Offline

Defragment data (REORG)

Automation	Enabled
Notification	Enabled
Table scope	TABSCHEMA NOT LIKE 'SYS%'

Optimize data access for applications (RUNSTATS)

Automation	Enabled
Notification	Enabled
Table scope	All tables

Maintenance window used for all online activities

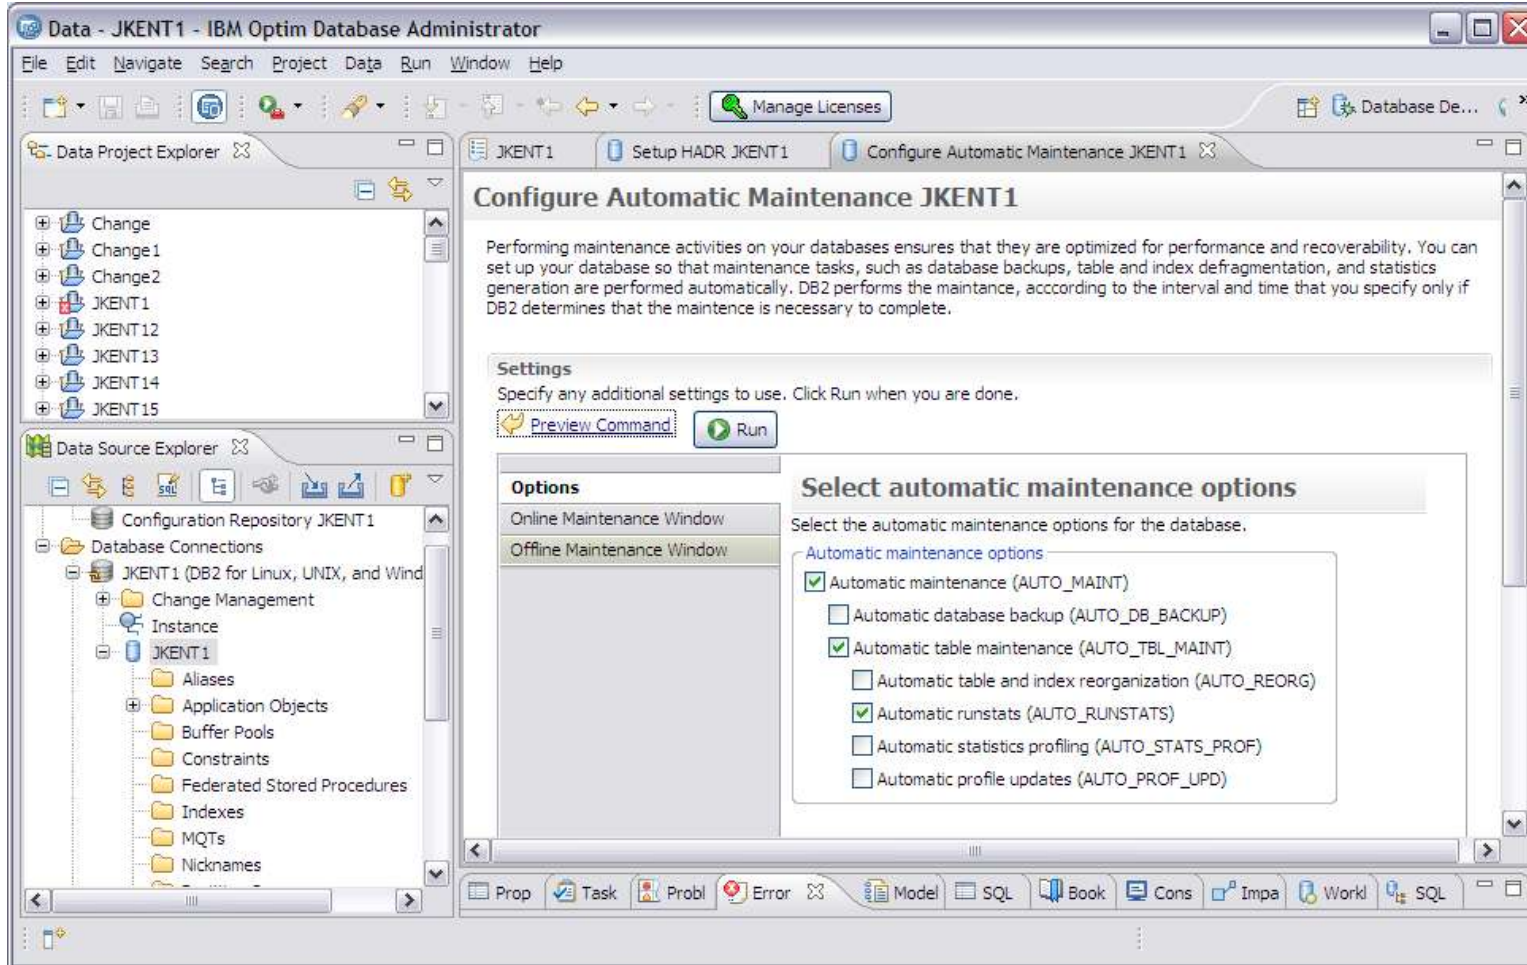
Online automatic maintenance can occur during the following window

Time	00:00 - 23:00 (23 hours)
Days of the week	ALL
Days of the month	ALL
Activities using this window	Optimize data access (RUNSTATS)

Back Finish Cancel

Data Studio – Setup Automatic Maintenance

Data Studio



Data - JKENT1 - IBM Optim Database Administrator

File Edit Navigate Search Project Data Run Window Help

Manage Licenses Database De...

Data Project Explorer

- Change
- Change1
- Change2
- JKENT1
- JKENT12
- JKENT13
- JKENT14
- JKENT15

Data Source Explorer

- Configuration Repository JKENT1
 - Database Connections
 - JKENT1 (DB2 for Linux, UNIX, and Wind
 - Change Management
 - Instance
 - JKENT1
 - Aliases
 - Application Objects
 - Buffer Pools
 - Constraints
 - Federated Stored Procedures
 - Indexes
 - MQTs
 - Nicknames

Configure Automatic Maintenance JKENT1

Performing maintenance activities on your databases ensures that they are optimized for performance and recoverability. You can set up your database so that maintenance tasks, such as database backups, table and index defragmentation, and statistics generation are performed automatically. DB2 performs the maintenance, according to the interval and time that you specify only if DB2 determines that the maintenance is necessary to complete.

Settings

Specify any additional settings to use. Click Run when you are done.

Preview Command Run

Options

Online Maintenance Window

Offline Maintenance Window

Select automatic maintenance options

Select the automatic maintenance options for the database.

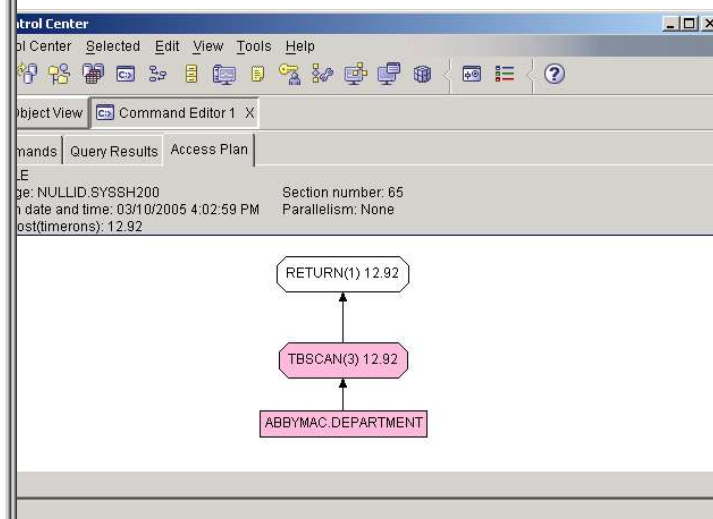
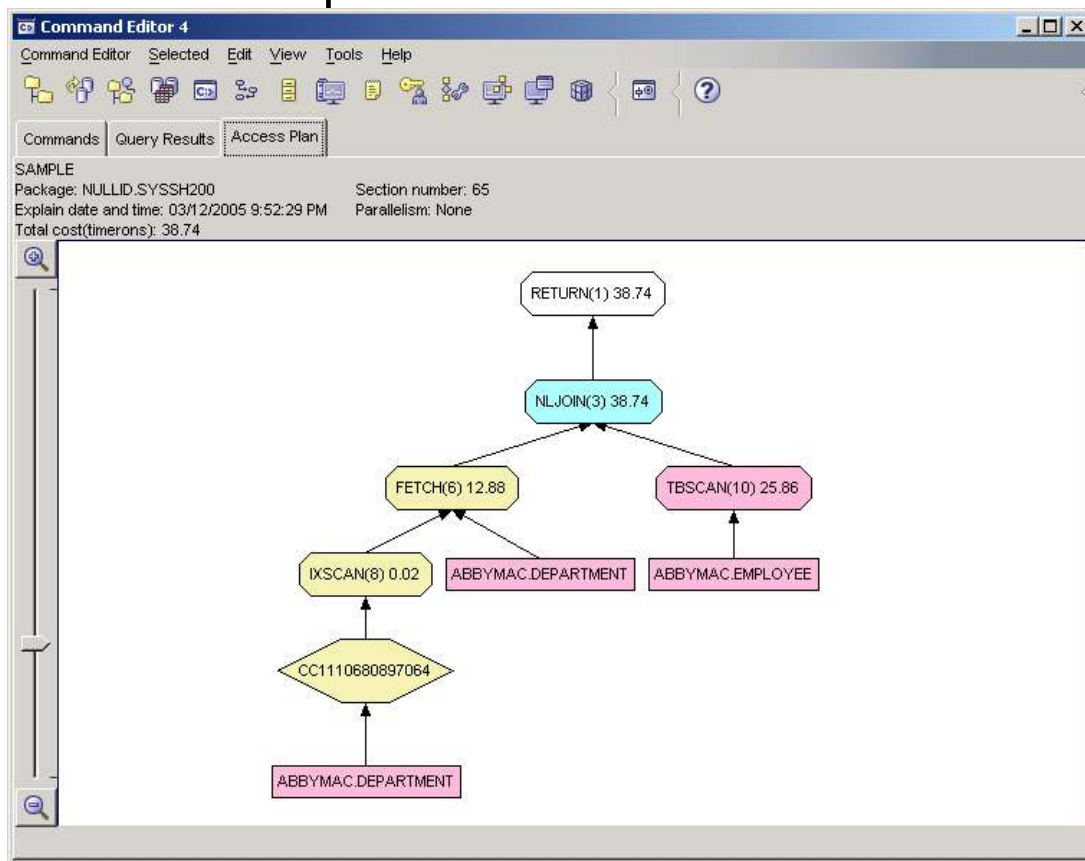
Automatic maintenance options

- Automatic maintenance (AUTO_MAINT)
- Automatic database backup (AUTO_DB_BACKUP)
- Automatic table maintenance (AUTO_TBL_MAINT)
- Automatic table and index reorganization (AUTO_REORG)
- Automatic runstats (AUTO_RUNSTATS)
- Automatic statistics profiling (AUTO_STATS_PROF)
- Automatic profile updates (AUTO_PROF_UPD)

Prop Task Probl Error Model SQL Book Cons Impa Work SQL

How to do Query Tuning ...

- In Control Center, you used Visual Explain analyze queries and look at access plans



- **SELECT EMPLOYEE.LASTNAME, DEPARTMENT.DEPTNAME, DEPARTMENT.DEPTNO**
- **FROM ABBYMAC.DEPARTMENT AS DEPARTMENT, ABBYMAC.EMPLOYEE AS EMPLOYEE**
- **WHERE DEPARTMENT.DEPTNO = EMPLOYEE.WORKDEPT AND DEPARTMENT.DEPTNO = 'D11'**

Data Studio Visual Explain

Access Plan Graph - JKENT1

Statement Node View Help

IBM-64DAC294200 - DB2 - JKENT1
 Package: NULLID.SYSSH200
 Explain date and time: 05/28/2009 8:44:35 PM
 Total cost(timerons): 15.25

Section number: 65
Parallelism: None

```

graph BT
    A[TBSCAN(3) 15.25] --> B[AN(9) 7.64]
    A --> C[TBSCAN(11) 7.64]
    B --> D[HSJOIN(7) 15.24]
    C --> D
    D --> E[Sort(5) 15.25]
    E --> F[TBSCAN(3) 15.25]
        
```

Information

EXP0020W Table has no statistics. The table "SALES"."CREDITCARDATA" has not had runstats run on it. This may result in a sub-optimal access plan and poor performance.

Close

Operator details - TBSCAN(3)

IBM-64DAC294200 - DB2 - JKENT1

Level of details: Overview Full

Cumulative cost	
Total cost	15.25 timerons
CPU cost	346,939.16 instructions
I/O cost	2 I/Os
First row cost	15.25 timerons
Remote communication cost	0 timerons

Cumulative properties	
Tables	SALES.CREDITCARDATA SALES.CUSTOMER
Columns	SALES.CUSTOMER.PHONENUMBER SALES.CUSTOMER.LASTNAME SALES.CUSTOMER.CUSTOMERID
Order columns	Number: 3 Name: SALES.CUSTOMER.LASTNAME
Predicates	Number: 2 Selectivity: 0.04
Cardinality	27.44
Total buffer pool pages used	0
Buffer pool usages	None

Input arguments	
Scan source	Scan over input plan
Columns retrieved	SALES.CUSTOMER.PHONENUMBER SALES.CUSTOMER.LASTNAME SALES.CUSTOMER.CUSTOMERID
Sargable predicates	None
Scan direction	Forward
Residual predicates	None
Block sargable predicates	None

Save As... Print... Close Help



How to monitor the health of the server ...

- After a health alert has occurred on one or more DB2 objects, you can display the details of the alert and bring up the Recommendation Advisor.

The screenshot shows the Health Center Control Center window. The main pane displays a table of alerts for the 'SAMPLE' database. The table has columns for Health Indicator, Value, Unit, and Category. One alert is visible: 'Instance Operational State' with a value of 'Down' and category 'DBMS'. A context menu is open over this alert, showing options: 'Show Details', 'Recommendation Advisor...', and 'Disable Evaluation'. The left pane shows a tree view with 'DB2' and 'SAMPLE' objects. The bottom status bar indicates 'Health monitor data is refreshed. (3/11/05 6:11 PM)'.

Health Indicator	Value	Unit	Category
Instance Operational State	Down		DBMS

Control Center

Show Details
 Recommendation Advisor...
 Disable Evaluation

The dialog box provides detailed information about the 'Instance Operational State' alert. It includes a description of the state and a 'View History' button.

Health indicator name: Instance Operational State
 Description
 An instance is considered healthy if the instance state does not restrict activity or tasks being performed. The state can be one of the following: ACTIVE, QUIESCE PENDING, QUIESCED, or DOWN. A non-Active state might generate an Attention alert. Refer to the DB2 Information Center for information on these states.

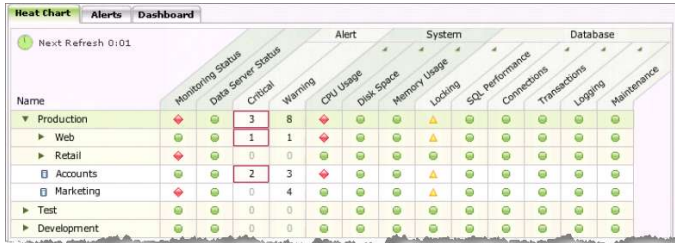
Instance: DB2
 Partition: 0
 Health indicator: db2.db2_op_status
 Timestamp: 03/11/2005 6:11:40 PM
 Current state: Attention

View History

Close Help

Data Studio - Administration Console

Heatchart – Overall Health Status



Where are the most important hotspots that need my attention?

Alert List – Historical Investigation

Severity	Alert Type	Timestamp	Database
Critical	CPU Utilization of LPAR/System	05:31pm	Support(z/OS)
Critical	CPU Utilization of LPAR/System	04:37pm	Account(z/OS)
Critical	CPU Utilization of LPAR/System	08:07am	Account(z/OS)
Warning	Application timeout	07:12pm	Support(z/OS)
Warning	Application timeout	05:44pm	Account(z/OS)
Warning	Application timeout	04:37pm	Marketing(z/OS)
Warning	Application timeout	11:18am	Marketing(z/OS)
Warning	Application timeout	11:07am	Marketing(z/OS)
Warning	Application timeout	2007/05/06	Marketing(z/OS)
Warning	CPU Utilization of LPAR/System	2007/05/06	Account(z/OS)
Warning	CPU Utilization of LPAR/System	2007/05/05	Account(z/OS)

What happened when I was out for lunch? ... Away for weekend?

Dashboard – Adhoc Investigation



Something doesn't seem quite right. I wonder what's happening?

Recommendations – Root Cause Analysis

Alert Detail: Table space TA.TS1 is offline @ Jun 25, 2007 12:48:06 PM

Table space TS1 in the TA database is offline

At the time of the alert, the TS1 tablespace is offline, and, as a result, is inaccessible.

Symptoms
The TS1 tablespace is inaccessible.

Causes
A table space is in this state if there is a problem preventing access to one or more of its containers. This is often caused by media problems that are either permanent (for instance a bad disk) or temporary (for instance an offline disk or unmounted file system). After the problem has been remediated and the containers are all accessible, the table space can be brought back online.

Diagnosing the problem

- Table space container is missing**
A table space container file for TS1 was renamed or moved.
The file system is offline.
The physical disk is unmounted.
The physical disk is corrupted.
- A table space container file for TS1 was deleted**
Table space container is damaged.
The disk sector is damaged.
The container file is tampered.
- Table space container is missing**
If one or more containers of a table space cannot be found by the database management system, the table space will be taken offline and put in an inaccessible state.
- Table space container is damaged**
If one or more containers of a table space are found to be damaged by the data manager, the table space will be taken offline and put in an inaccessible state.

Guide me to the root cause and help me fix it properly; I need to know all the relevant info to make the best decision.



Chat with the Lab

COOL functionality from the priced offerings

ODS → Optim Development Studio

ODA → Optim Database Administrator

OQT → Optim Query Tuner

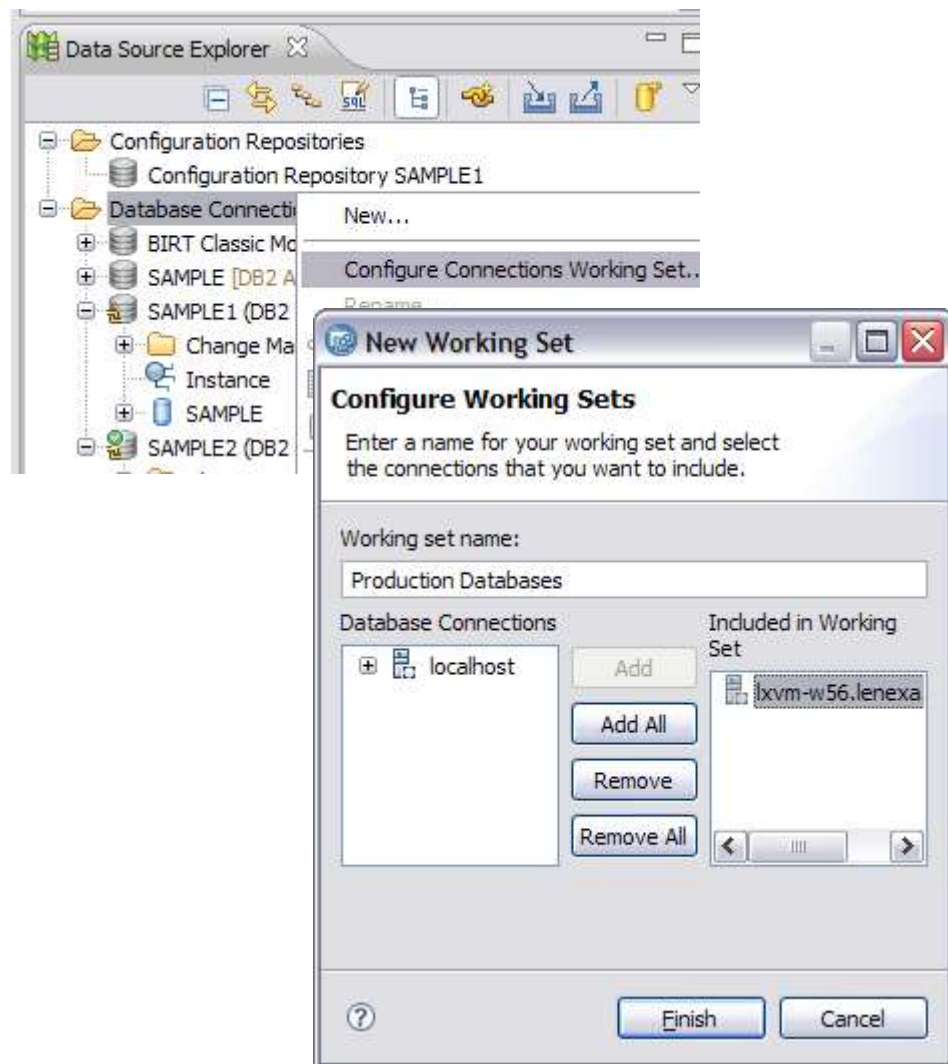
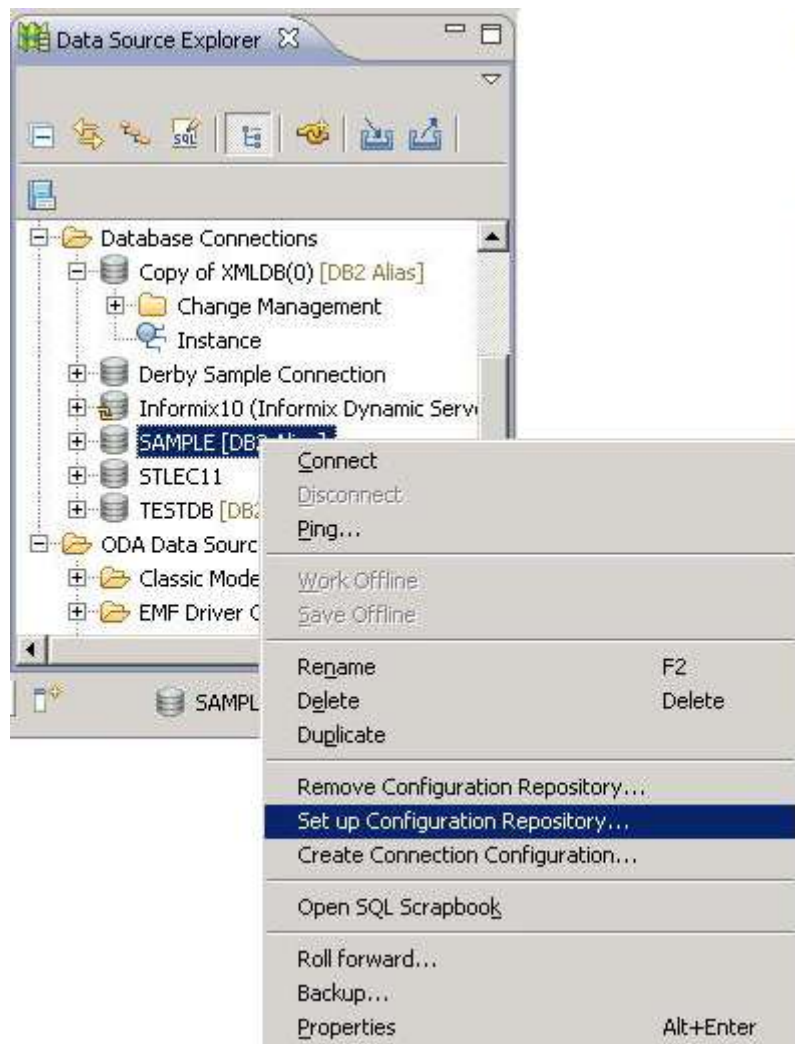
IDA → InfoSphere Data Architect

Navigation (In ODA, ODS and OQT)

- Need to share connectivity information with various workstations
 - Create Configuration Repository to store connection information
 - Export to file and share with team members
- Need to group my database server connections for easier access
 - Create a Working Set containing the connections you are working with.



Configuration Repositories/Working Sets



SQL Development (In ODS)

- **SQL Outline View**

The screenshot displays the IBM Data Studio Developer interface. The main editor shows a Java file named `ReportApplication.java` with the following code snippet:

```
statement.close();

if(customerList.size() > 0) {

    sql = "SELECT ORD_NBR AS ORDER_NUMBER, " +
        "ORD_NBR_OF_ITEMS AS NUMBER_OF_ITEMS, " +
        "ORD_NBR_OF_PRODS AS NUMBER_OF_PRODUCTS " +
        "FROM GOSALESCST.CUST_ORD " +
        "WHERE CUST_CODE = ?";

    pStatement = connection.prepareStatement(sql);

    for(int i = 0; i < customerList.size(); i++) {
        pStatement.setInt(1, customerList.get(i).intValue());
        pStatement.execute();
        resultSet = pStatement.getResultSet();
        while(resultSet.next()) {
            System.out.println(resultSet.getInt("ORDER_NUMBER") +
                resultSet.getInt("NUMBER_OF_ITEMS") +
                resultSet.getInt("NUMBER_OF_PRODUCTS"));
        }
    }

    resultSet.close();
}
```

Annotations on the screenshot include:

- Correlate with issuing code**: A blue speech bubble pointing to the SQL query string in the code editor.
- Visualize execution metrics**: A blue speech bubble pointing to the SQL Outline View table at the bottom.
- Correlate with data source attributes**: A blue speech bubble pointing to the Data Source Explorer on the left, which shows a table named `CUST_ORD` with columns like `ORD_NBR`, `ORD_DATE`, and `ORD_NBR_OF_ITEMS`.
- Execute, tune, share, trace, explore SQL**: A blue speech bubble pointing to the context menu of the SQL Outline View table.

The SQL Outline View table at the bottom shows the following data:

SQL Statement	Number of Times Run	Total Time	Max Time	Average Time	Min Time
SELECT ORD_NBR AS ORDER_NUMBER, ...	234	16489.00	75.00	70.47	68.00
SELECT CUST_CODE, STDE...	4	287.00	73.00	71.75	71.00
SELECT CUST_CODE, AV...	4	294.00	73.00	73.00	73.00
SELECT CUST_CODE, COR...	3	234.00	73.00	74.00	74.00
SELECT count(CUST_CODE)...	14	1236.00	73.00	68.00	68.00
SELECT CUST_CODE, SUM...	3	222.00	73.00	73.00	73.00



Query Tuning (In OQT)

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Data - Project1_Query Group 1_Query Before Tuning - IBM Optim Query Tuner Client".

Query Format and Annotation: Shows the parsed and formatted query:

```
SELECT *
FROM TPCDS.STORE_SALES
, TPCDS.DATE_DIM
WHERE TPCDS.DATE_DIM.D_YEAR = 2002
AND TPCDS.DATE_DIM.D_MOY BETWEEN 1 AND 1 + 3
AND TPCDS.STORE_SALES.SS_SOLD_DATE_SK = TPCDS.DATE_DIM.D_DATE
```

Access Plan Graph: A hierarchical diagram showing the execution plan. At the top is a RETURN node (730.40). Below it is an HSJOIN node (4747.05). This join node branches into two TBSCAN nodes: one for STORE_SALES (2.88014e+006) and one for DATE_DIM (120.399). Both TBSCAN nodes are connected to their respective table nodes (STORE_SALES and DATE_DIM).

Advisor Recommendation Overview: A table listing recommendations from various advisors:

Advisor	Priority	Description
Query Advisor	LOW	Consider replacing the asterisk (*) or the long column list of table STORE_SALES by a list of columns.
Access Path Advisor	LOW	The TPCDS.STORE_SALES table is accessed by a relational scan (Operator ID = 1).
Access Path Advisor	LOW	The TPCDS.DATE_DIM table is accessed by a relational scan (Operator ID = 2).
Index Advisor	LOW	Index recommendations found.
Statistics Advisor	MAINTENANCE	Gather and recalculate all of relevant statistics for this query.

Query Recommendation Detail: Shows the SQL text of the query being analyzed, identical to the one in the Query Format and Annotation pane.

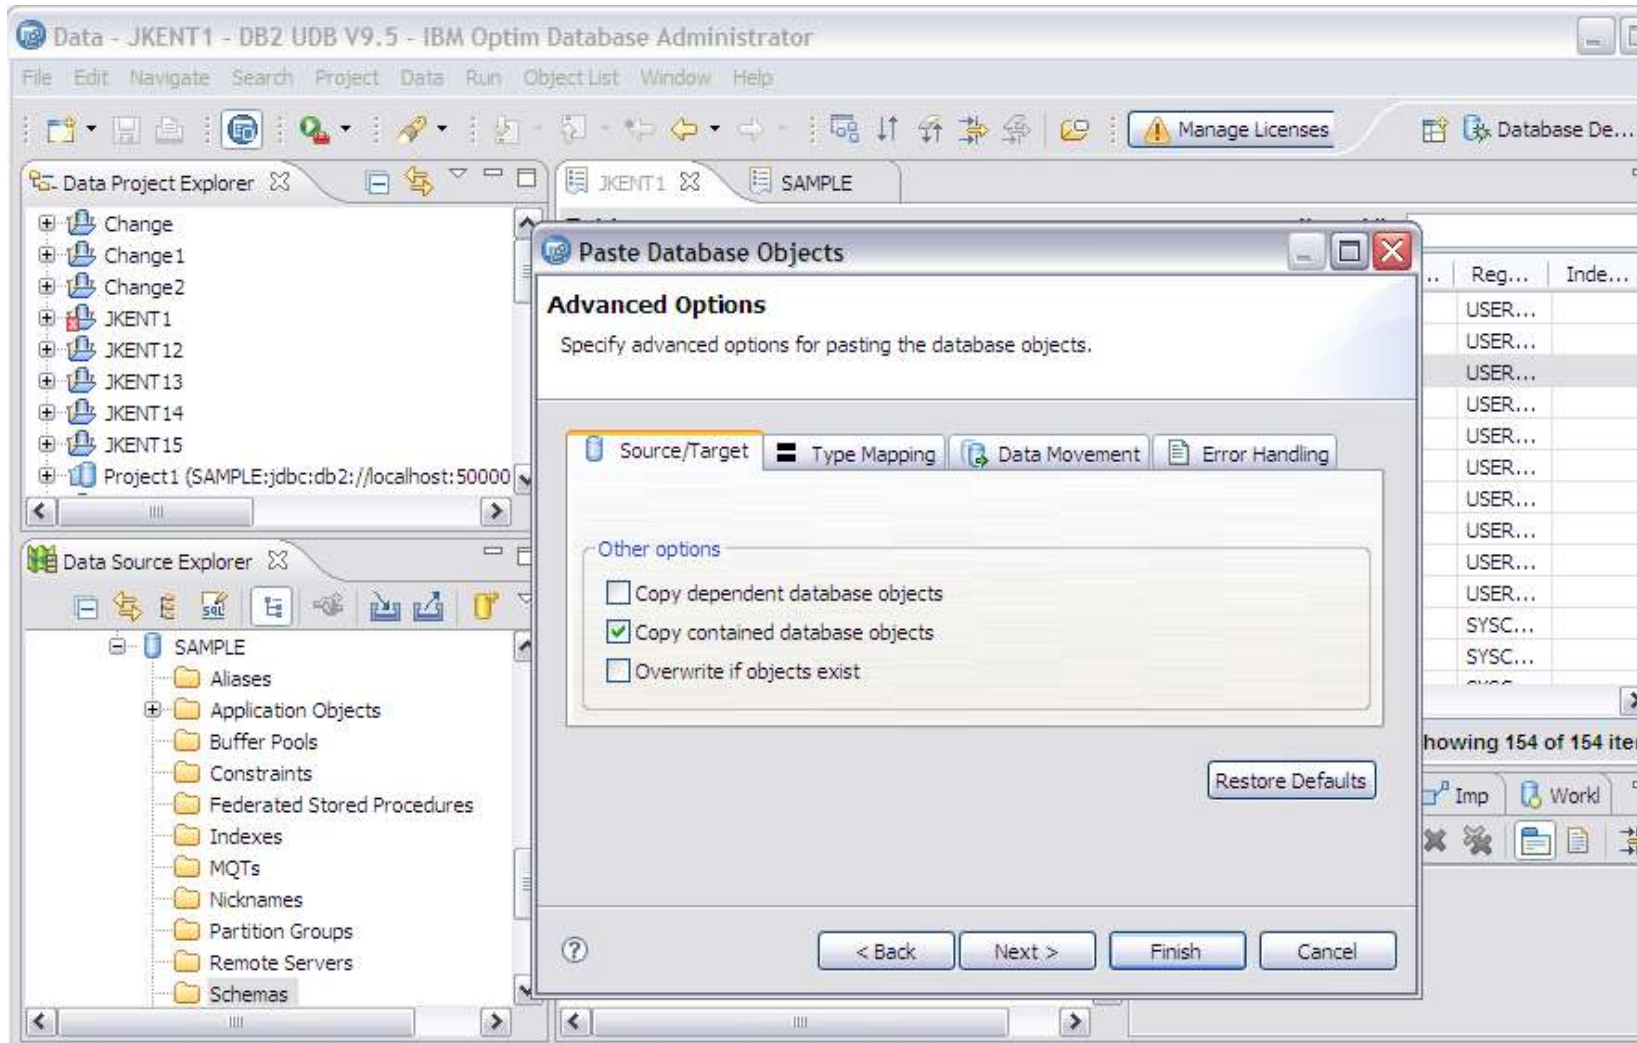


Design (In IDA)

- Use InfoSphere Data Architect to create or work with existing logical or physical data models. Deploy the model by having IDA automatically create the DDL for the physical data model.

The screenshot displays the IBM InfoSphere Data Architect (IDA) interface. On the left, the 'Data Project Explorer' shows a hierarchical tree of objects under 'Database Model.dbm', including 'Data Models', 'SQL Statements', and various tables like 'ACT', 'CUSTOMER', and 'EMPLOYEE'. A context menu is open over the 'EMPLOYEE' table, with 'Generate DDL...' selected. On the right, the 'Properties' pane shows the configuration for the '<Table> EMPLOYEE' table, including fields for Name, Label, Schema, and Data capture. Below this, the 'Physical Data Model Editor' is visible, showing 'Database Information' (Vendor: DB2 UDB, Version: V9.7) and 'Data Model Information' (Name, Location, Size, Last modified, Editable).

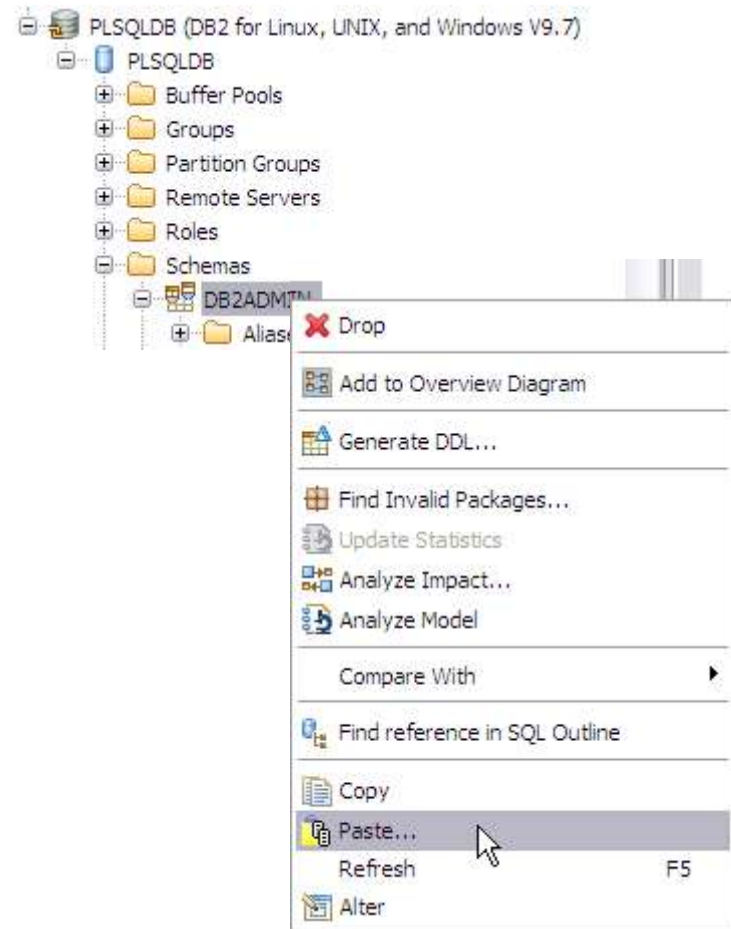
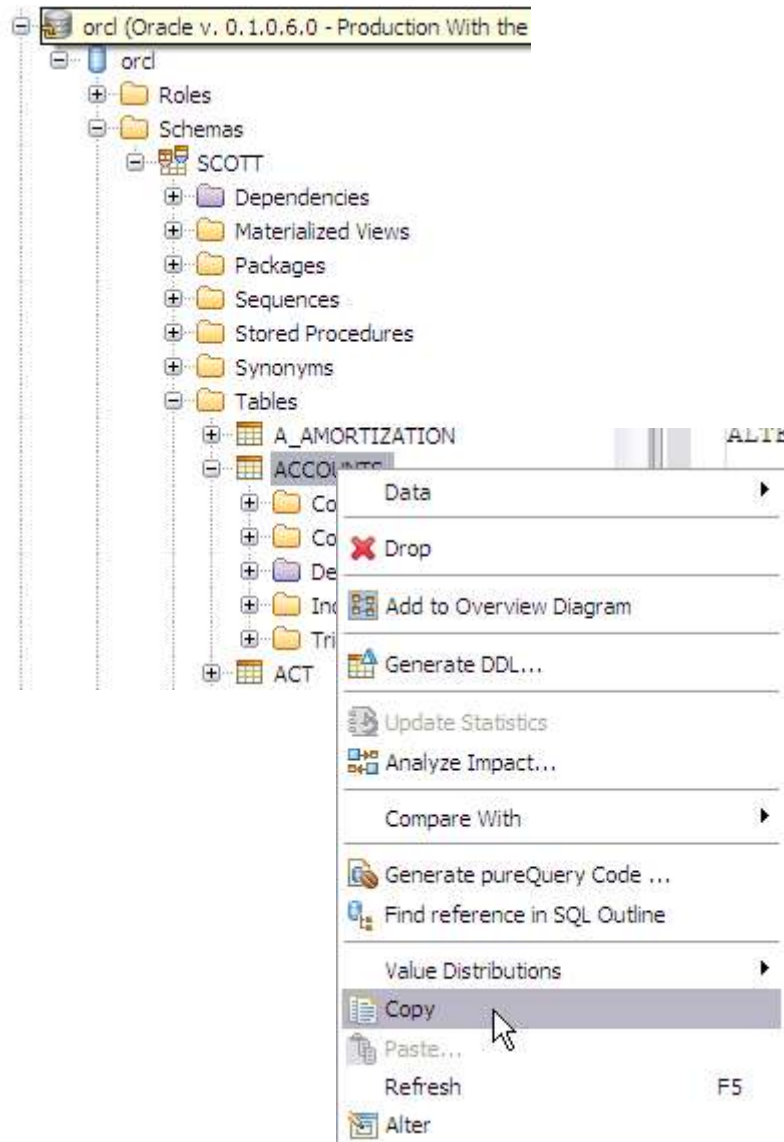
Data Management – Copy/Paste (In ODS)



What's new for DB2 9.7?



ODS Data Object Copy from Oracle to DB2



Support for DB2 Version 9.7 for Linux, UNIX, and Windows PL/SQL

- You can create PL/SQL modules and new data types for PL/SQL modules:
- Row data type
 - A structure composed of multiple fields each with their own name and data type that you can use to store the column values of a row in a result set or table.
- Array data type
 - A structure that contains an ordered collection of data elements, in which each element can be referenced by its ordinal position in the collection.



Questions



Thank You!

ibm.com/db2/labchats



Thank you for attending!

