



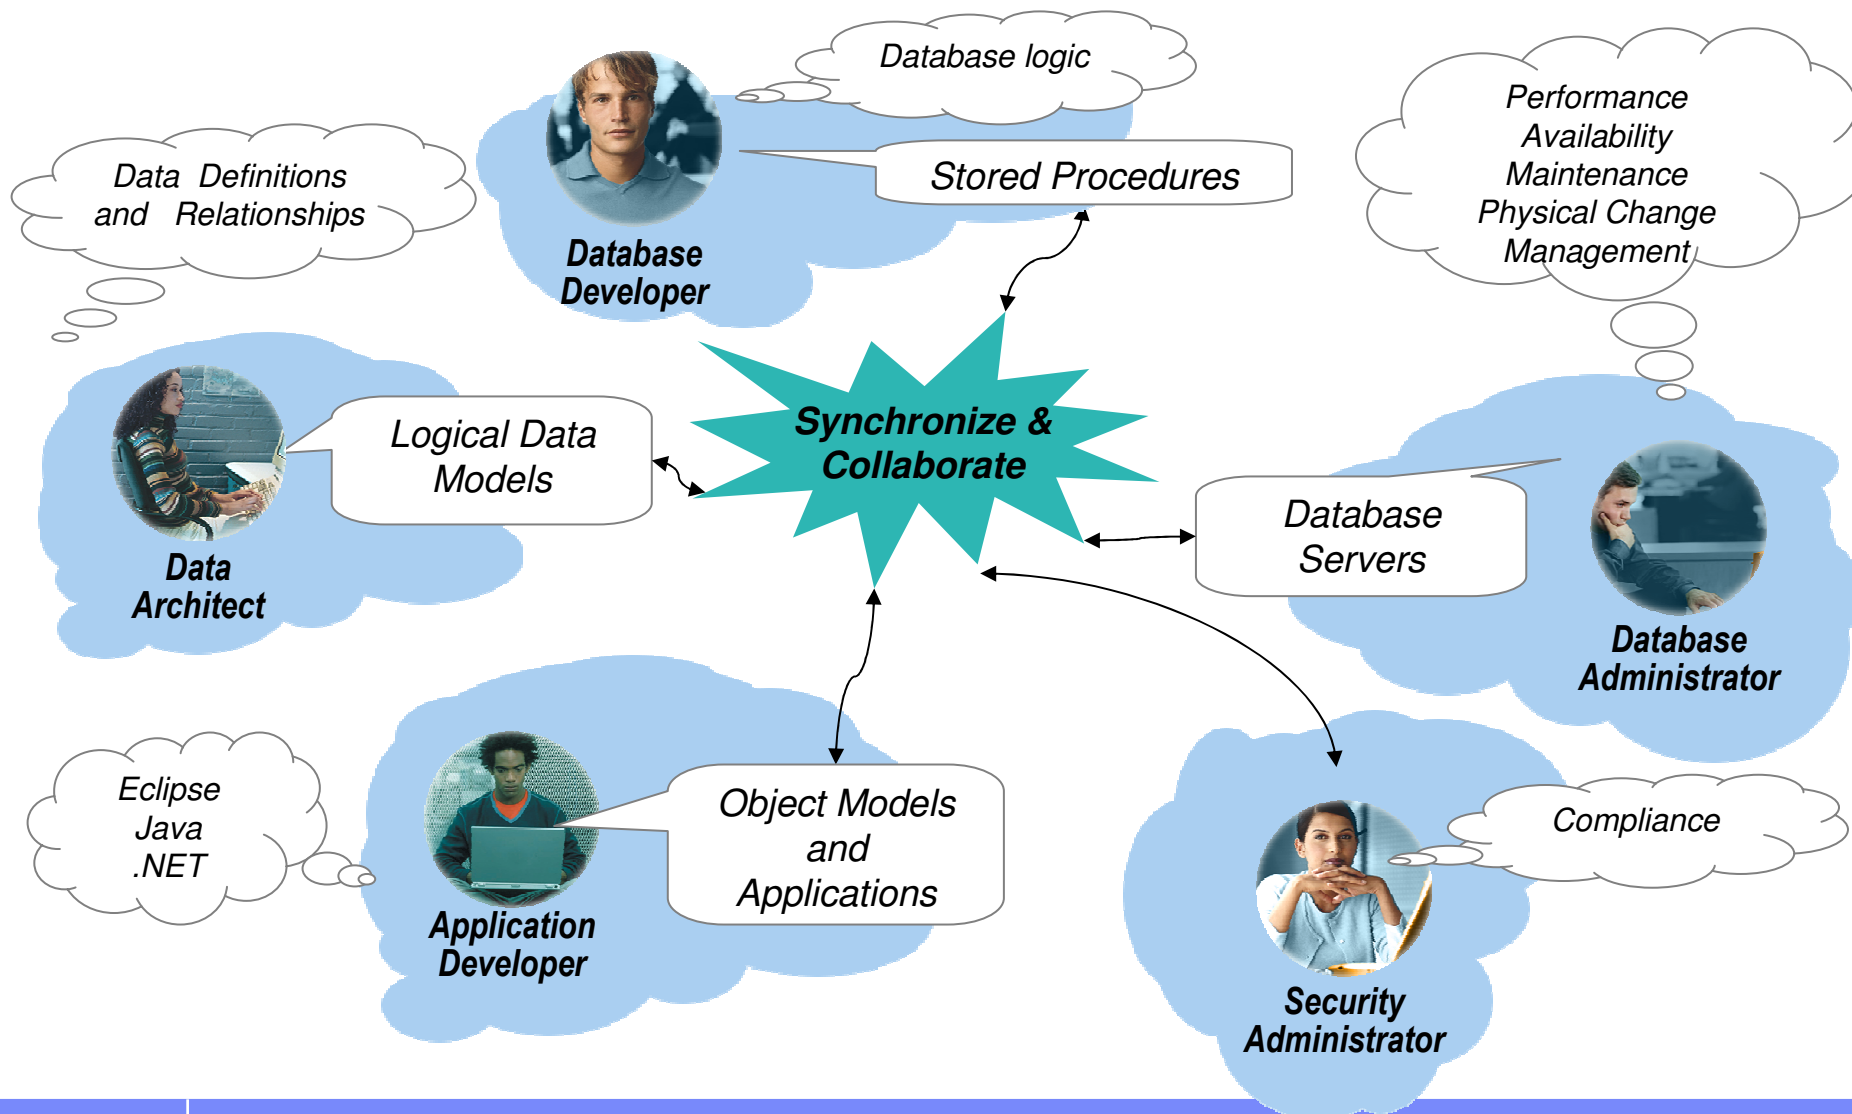
Integrated Data Management – Data Studio

Alfred Horng

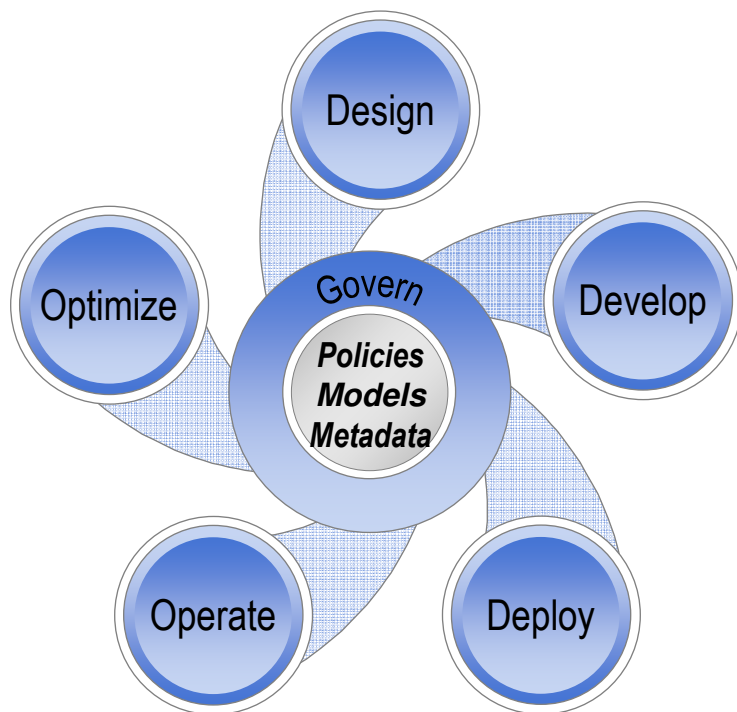
IBM SWG

kfhorng@tw.ibm.com

Integrated Data Management



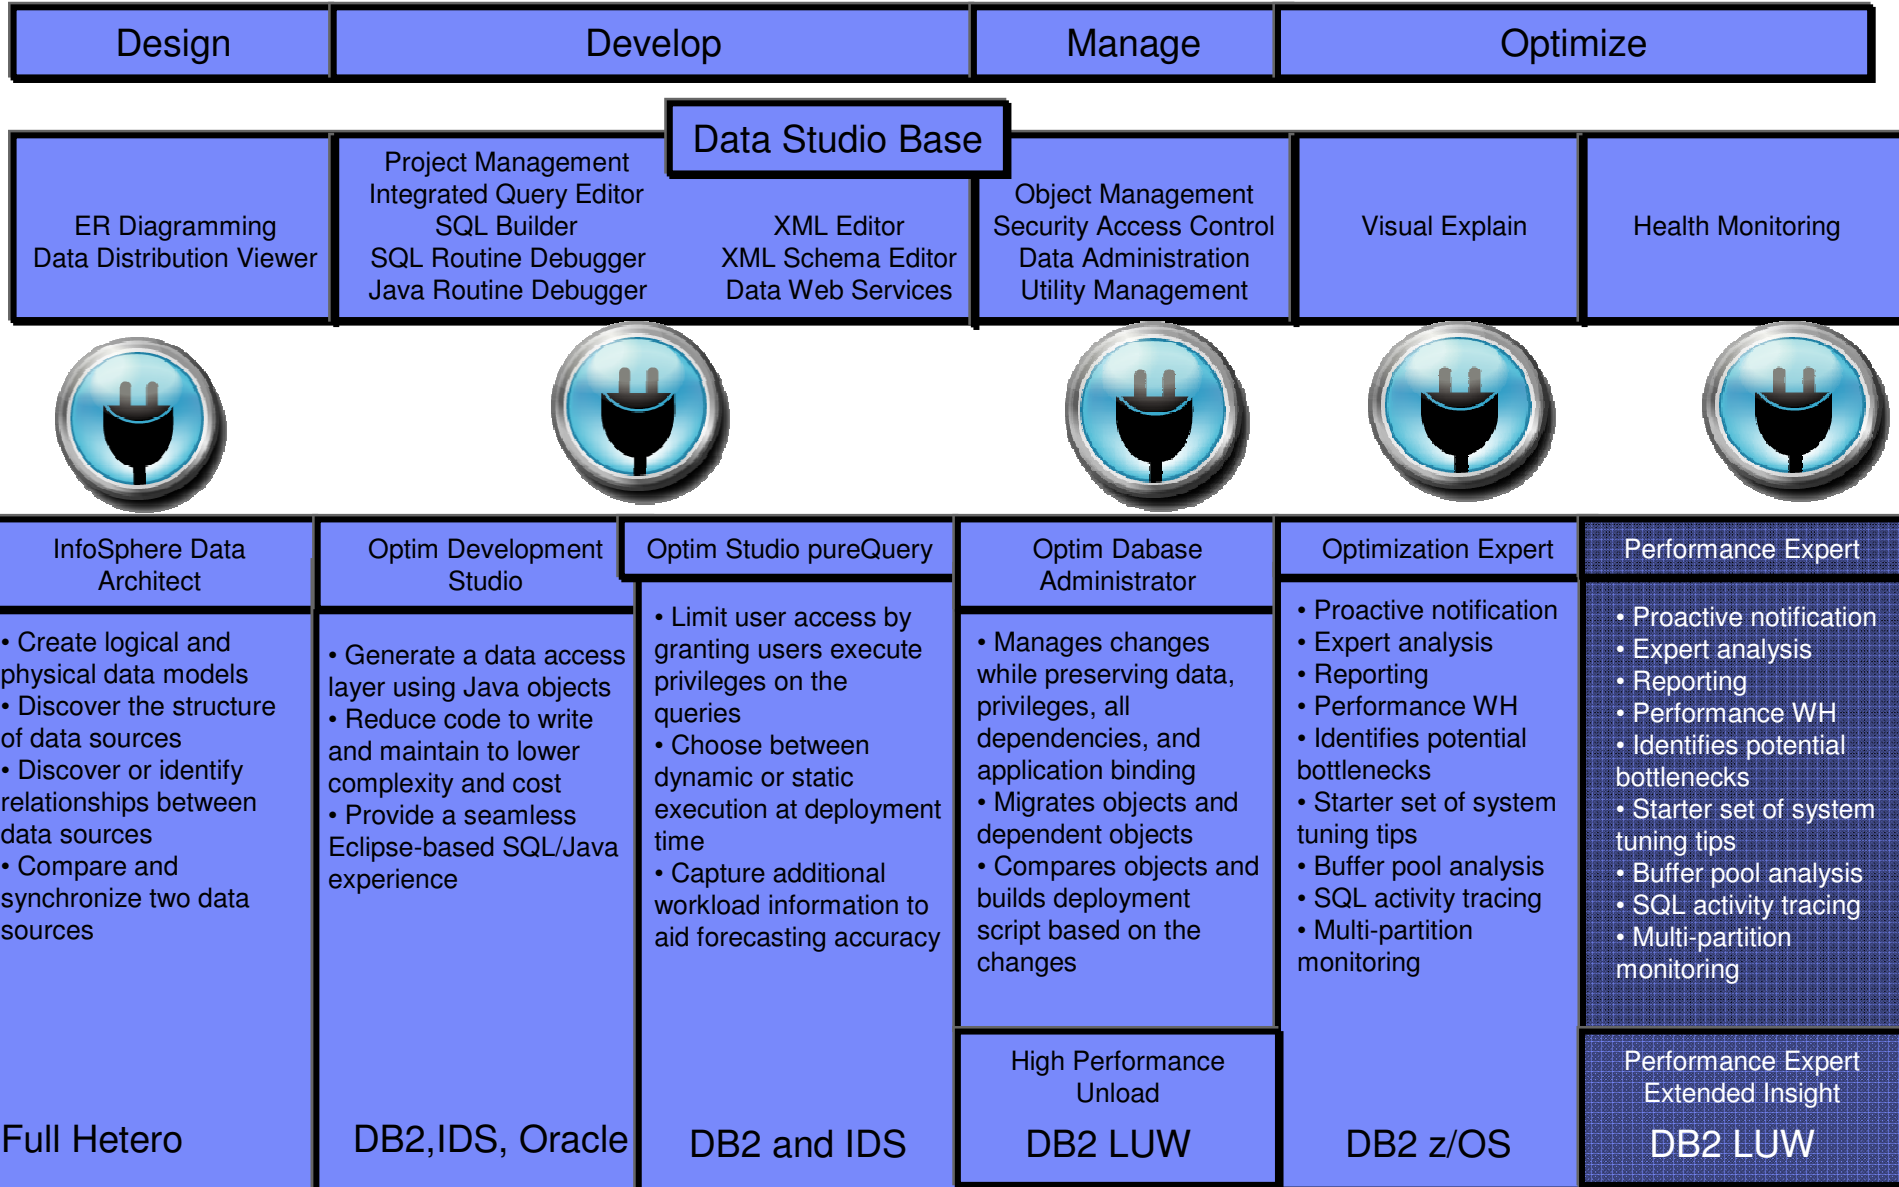
Integrated Data Management – What's Different?



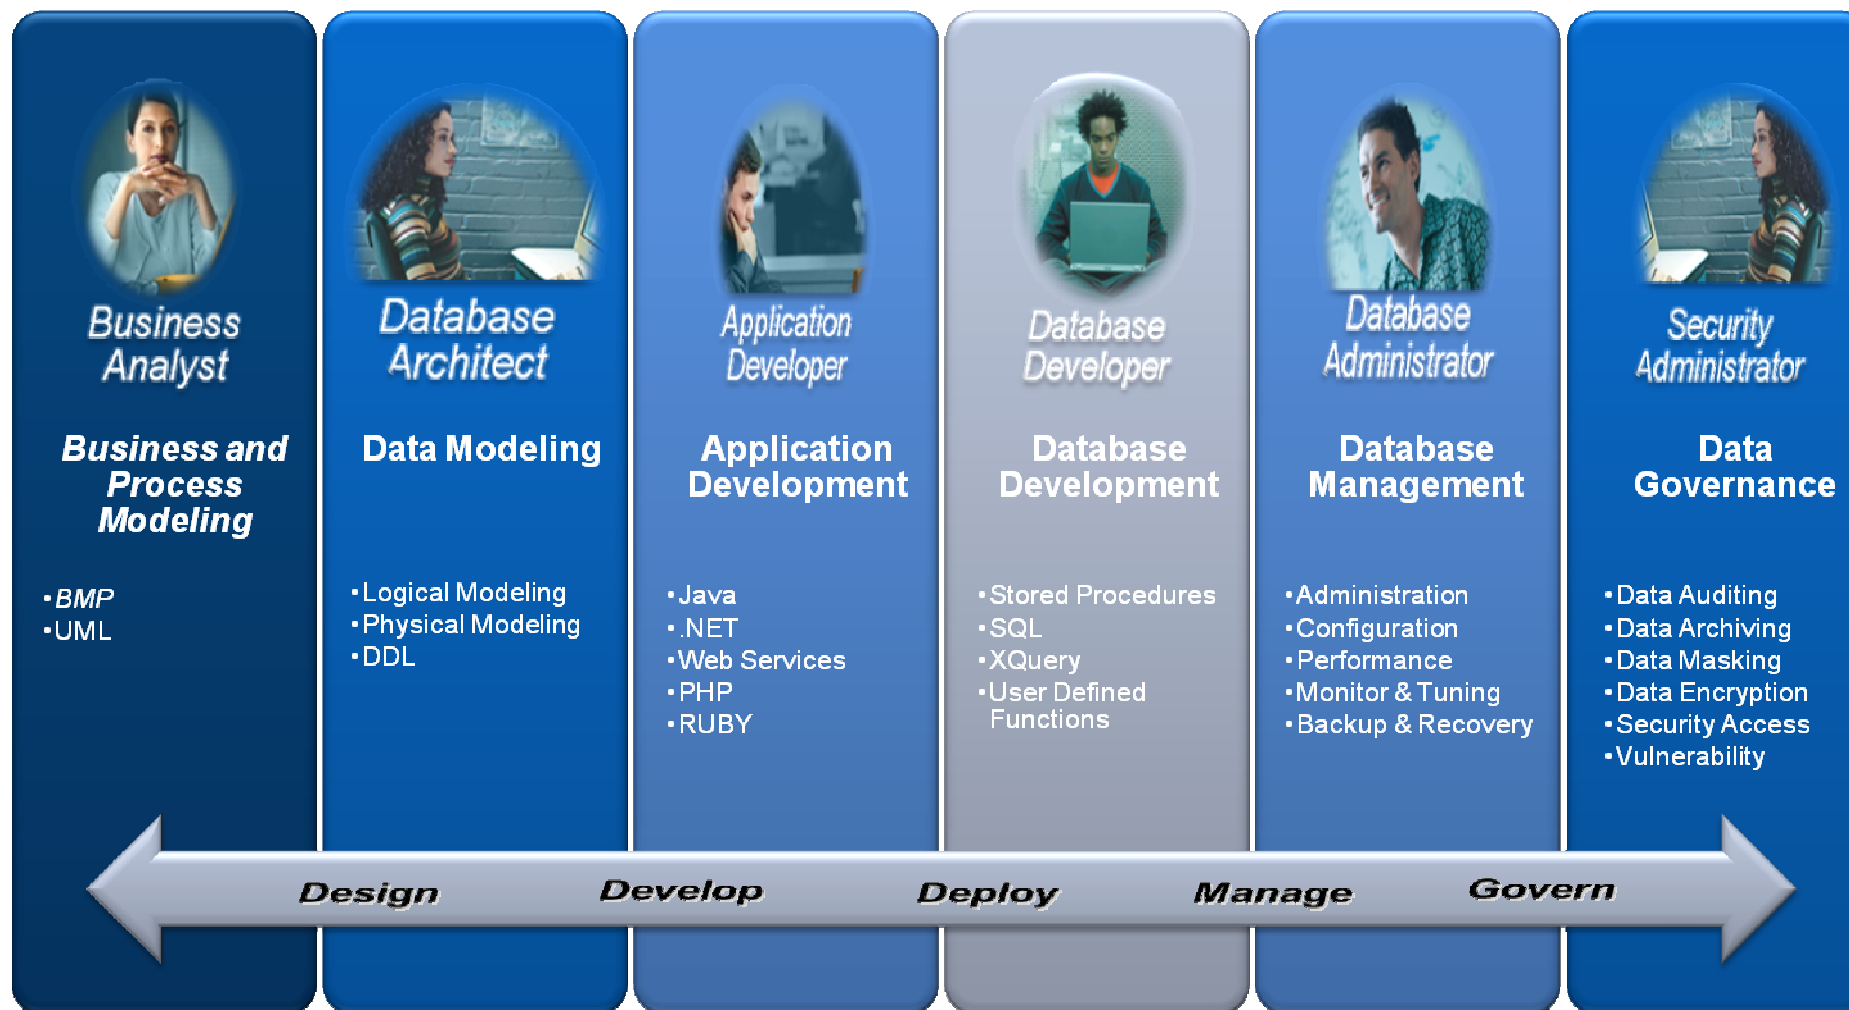
- Support business growth
 - Accommodate new initiatives without expanding infrastructure
 - Simplify application upgrades, consolidations & retirements
- Produce enterprise-ready applications faster
 - Improve data access, speed iterative testing
 - Empower collaboration between architects, developers & DBAs
- Consistently achieve service level targets
 - Automate and simplify operations
 - Provide contextual intelligence across the solution stack
- Facilitate alignment, consistency & governance
 - Define business policies and standards up front; share, extend, and apply throughout the lifecycle



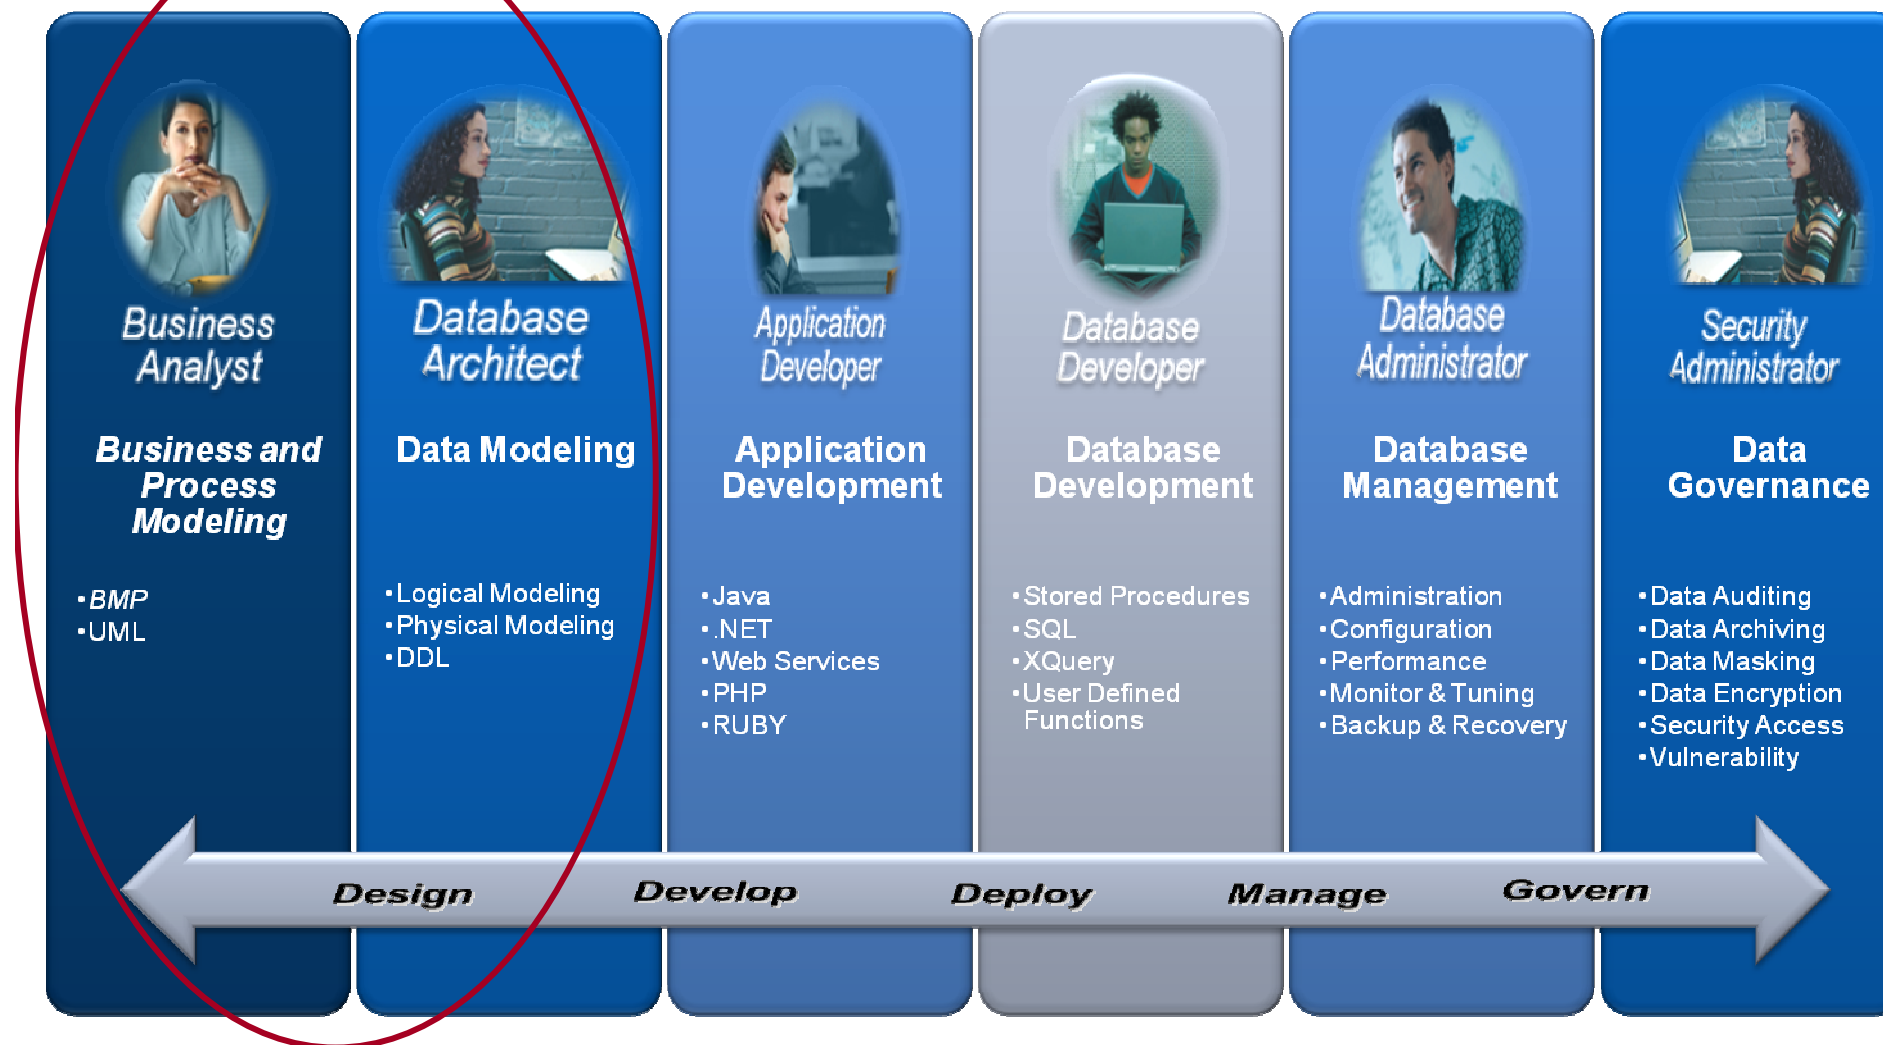
Data Studio Packaging



Core Users of the IBM Data Studio



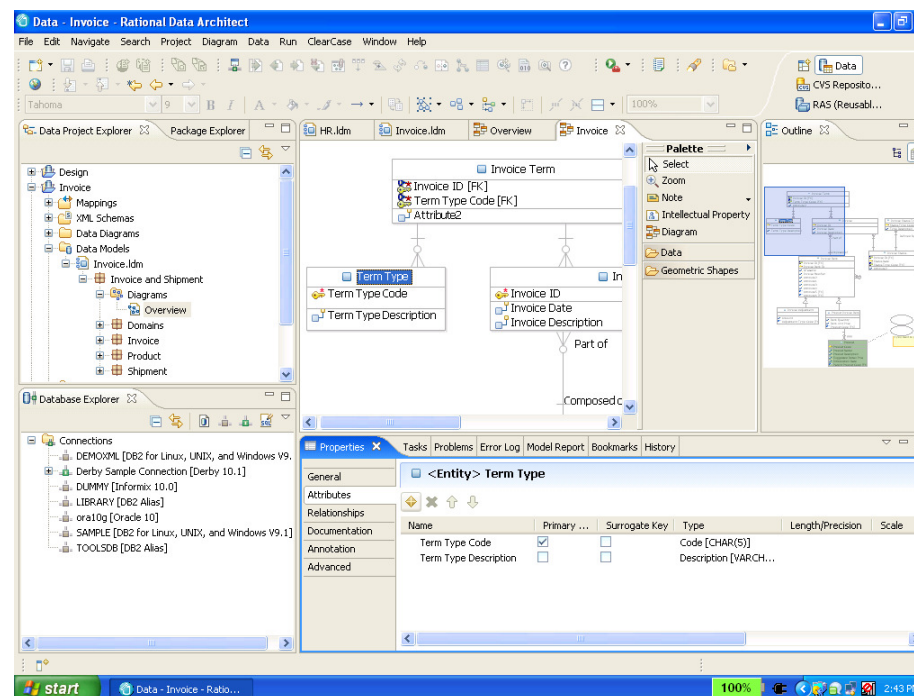
Users of the IBM Data Studio



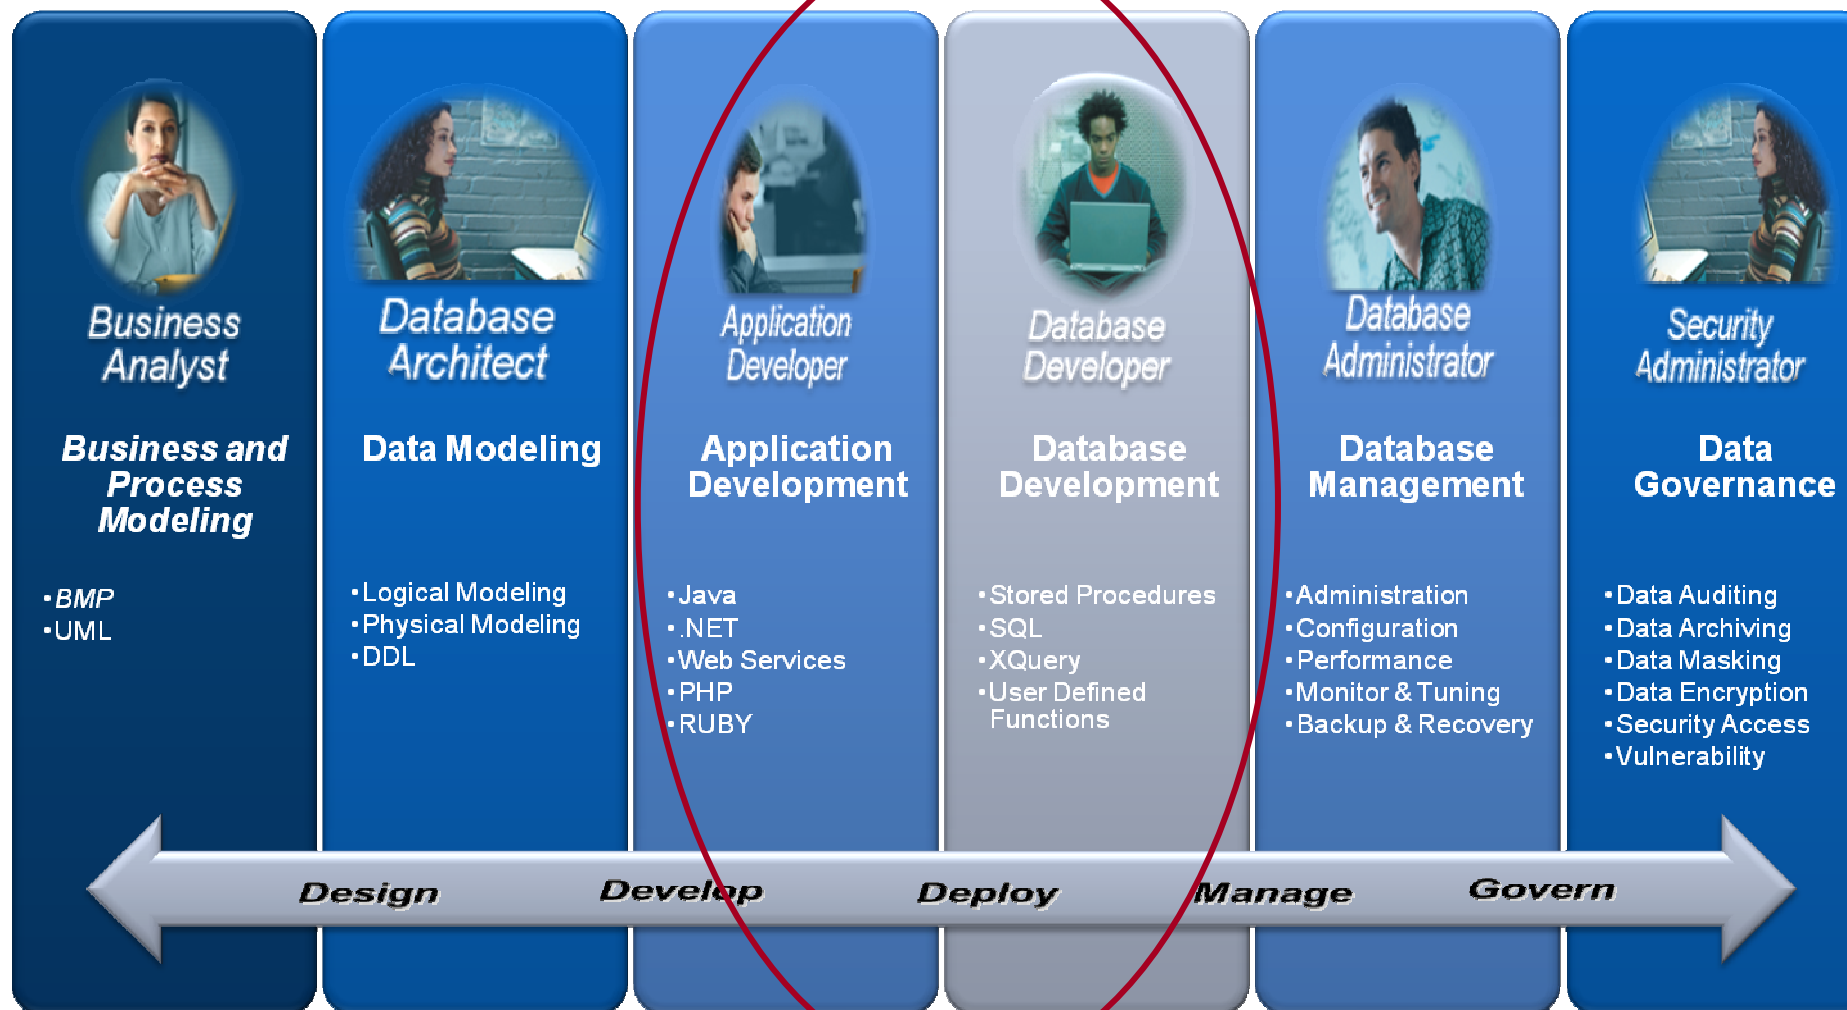
InfoSphere Data Architect

InfoSphere Data Architect is a collaborative, data design solution to discover, model, relate, and standardize diverse data assets.

- Increase Data Quality and Integrity
 - Analyze and enforce compliance to enterprise standards and data privatization
 - Support business and IT collaboration via a common business glossary
 - Use with IBM Industry Models for industry-specific best practices
 - Facilitate model-driven development via seamless integration with Rational Software Delivery Platform
 - Automate transformations between the application model and the data model



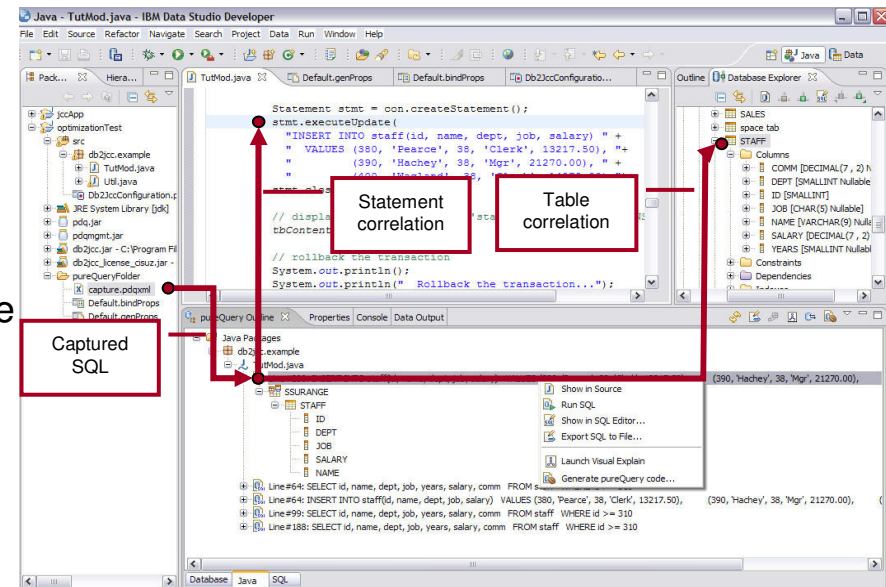
Users of the IBM Data Studio



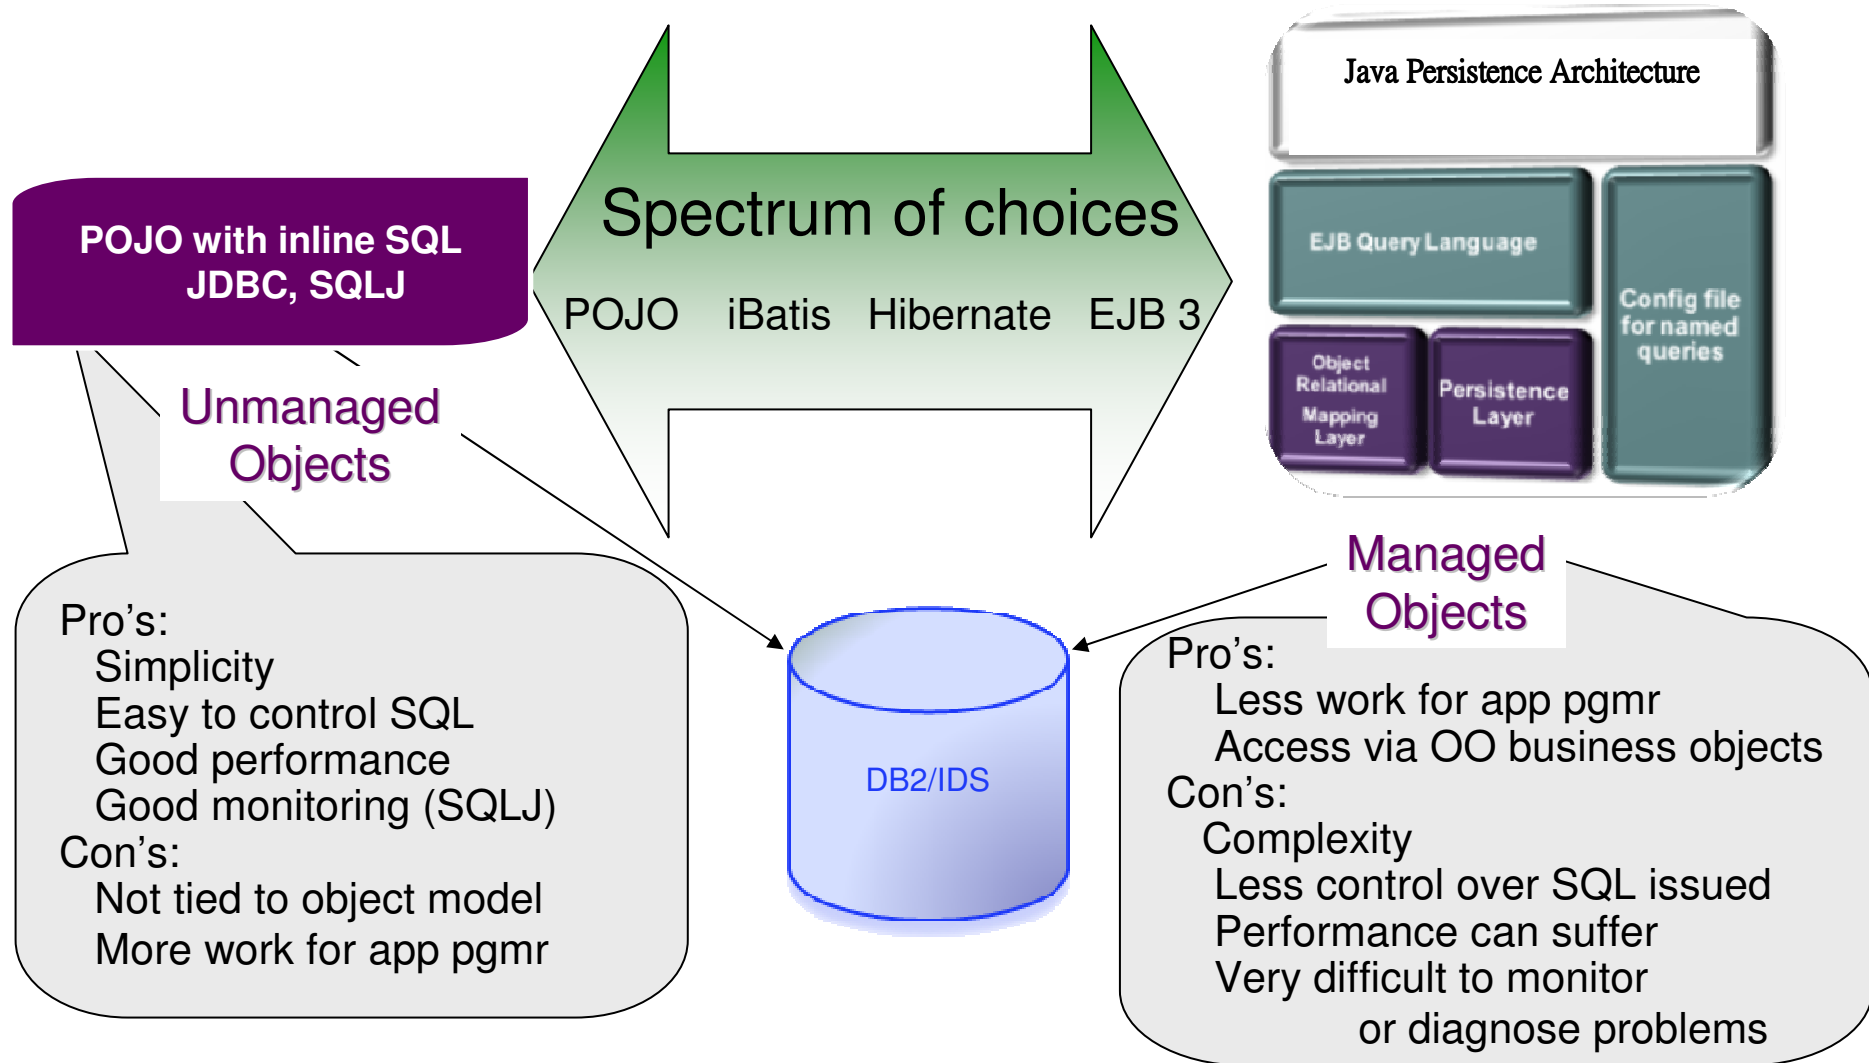
IBM Optim Development Studio and pureQuery Runtime

IBM Optim Development Studio is an integrated database development environment that speeds application design, development, and deployment while increasing data access performance and manageability.

- **Ease JAVA coding**
 - pureQuery code assistance
- **Improve predictability and manageability with static SQL**
 - Switch dynamic to static SQL without changing code
 - Eliminate SQL Injection Risk by approving which statements can be executed
 - Replace existing SQL with more efficient SQL without changing the source code
- **Visualization of JAVA request to SQL code**
 - SQL Outline correlates SQL to JAVA code and the associated objects
 - Provide developers with: No. of executions, elapsed time, min./max executions
- **Impact Analysis**
 - Display lines of code associated to a DB2 object to understand the impact of an object change.
- **PureQuery for JAVA, Open JPA and .NET**

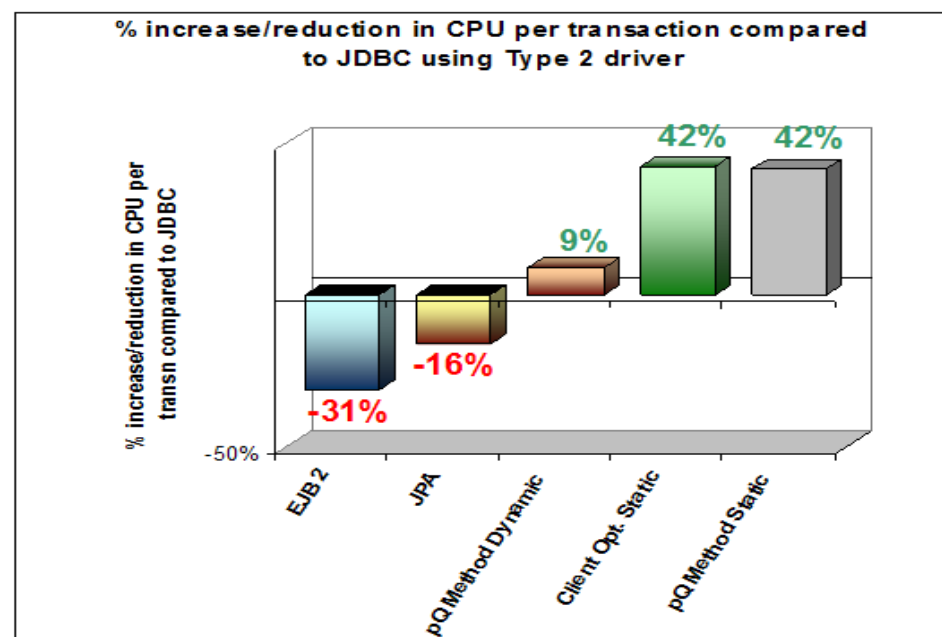
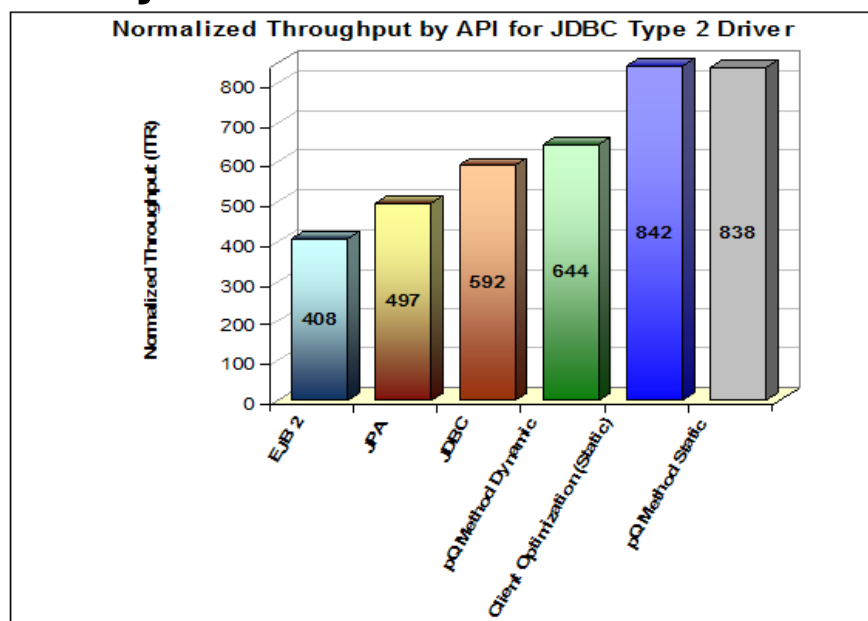


Optim Studio pureQuery – Best of Breed



Data Studio pureQuery Runtime for z/OS

- In-house testing shows double-digit reduction in CPU costs over dynamic JDBC

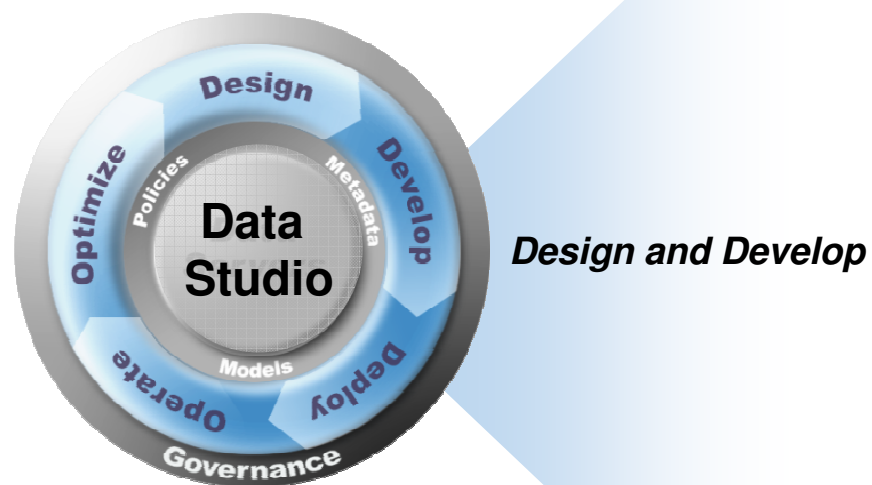


- IRWW – an OLTP workload, Type 2 driver (local call)
- Cache hit ratio between 70 and 85%
- 42% reduction in CPU per transaction over dynamic JDBC

What is pureQuery



pureQuery is a high-performance, data access platform to simplify developing, managing, securing, and optimizing Java data access for new and existing applications.



pureQuery Components:

- Development tools
 - Integrated development environment with Java and SQL support delivered with IBM Data Studio Developer
- Simple and intuitive API
 - Enables SQL access to databases or in-memory Java objects
- pureQuery Runtime
 - Flexible deployment options with static SQL support with IBM Data Studio pureQuery Runtime for z/OS or for Linux, UNIX, and Windows

pureQuery for IDS & DB2

- **For existing and new JDBC Applications**
 - **Capture**
 - See all SQL and Java for every database table & column
 - Find and Fix problems quickly
 - **SQL Injection Prevention and Rogue Query Control**
 - Lock down dynamic SQL to only approved SQL
 - **Tuned Query Replacement**
 - Replace SQL with optimized SQL without changing application
 - **Static SQL**
 - Lock access path and get performance and reliability of Static SQL against DB2
- **For new applications**
 - **Java Best Practices**
 - queryFirst, updateMany
 - Heterogenous Batch

Retrieve a single row from Database

pureQuery API's:

```
Employee my_emp = db.queryFirst("SELECT Name, HomeAddress, HomePhone
FROM Employee WHERE Name=?", Employee.class, my_emp);
```

-or-

```
Employee my_emp = getEmployee(name);
```

XML file or Java annotation
SELECT * FROM EMPLOYEE
WHERE NAME=?1;

SQLJ:

```
#sql [con] { SELECT NAME, ADDRESS, PHONE_NUM INTO :name, :addr, :phone FROM EMP
WHERE NAME=:name };
```

```
new Employee my_emp;
my_emp.setName(name);
my_emp.setHomeAddress(addr);
my_emp.setHomePhone(phone);
```

JDBC:

```
java.sql.PreparedStatement ps = con.prepareStatement(
"SELECT NAME, ADDRESS, PHONE_NUM FROM EMP
WHERE NAME=?");
```

```
ps.setString(1, name);
java.sql.ResultSet names = ps.executeQuery();
names.next();
new Employee my_emp;
my_emp.setName(names.getString(1));
my_emp.setHomeAddress(names.getString(2));
my_emp.setHomePhone(names.getString(3));
names.close();
```

Table	Column	Type
EMP	NAME	CHAR(64)
EMP	ADDRESS	CHAR(128)
EMP	PHONE_NUM	CHAR(10)

ad2

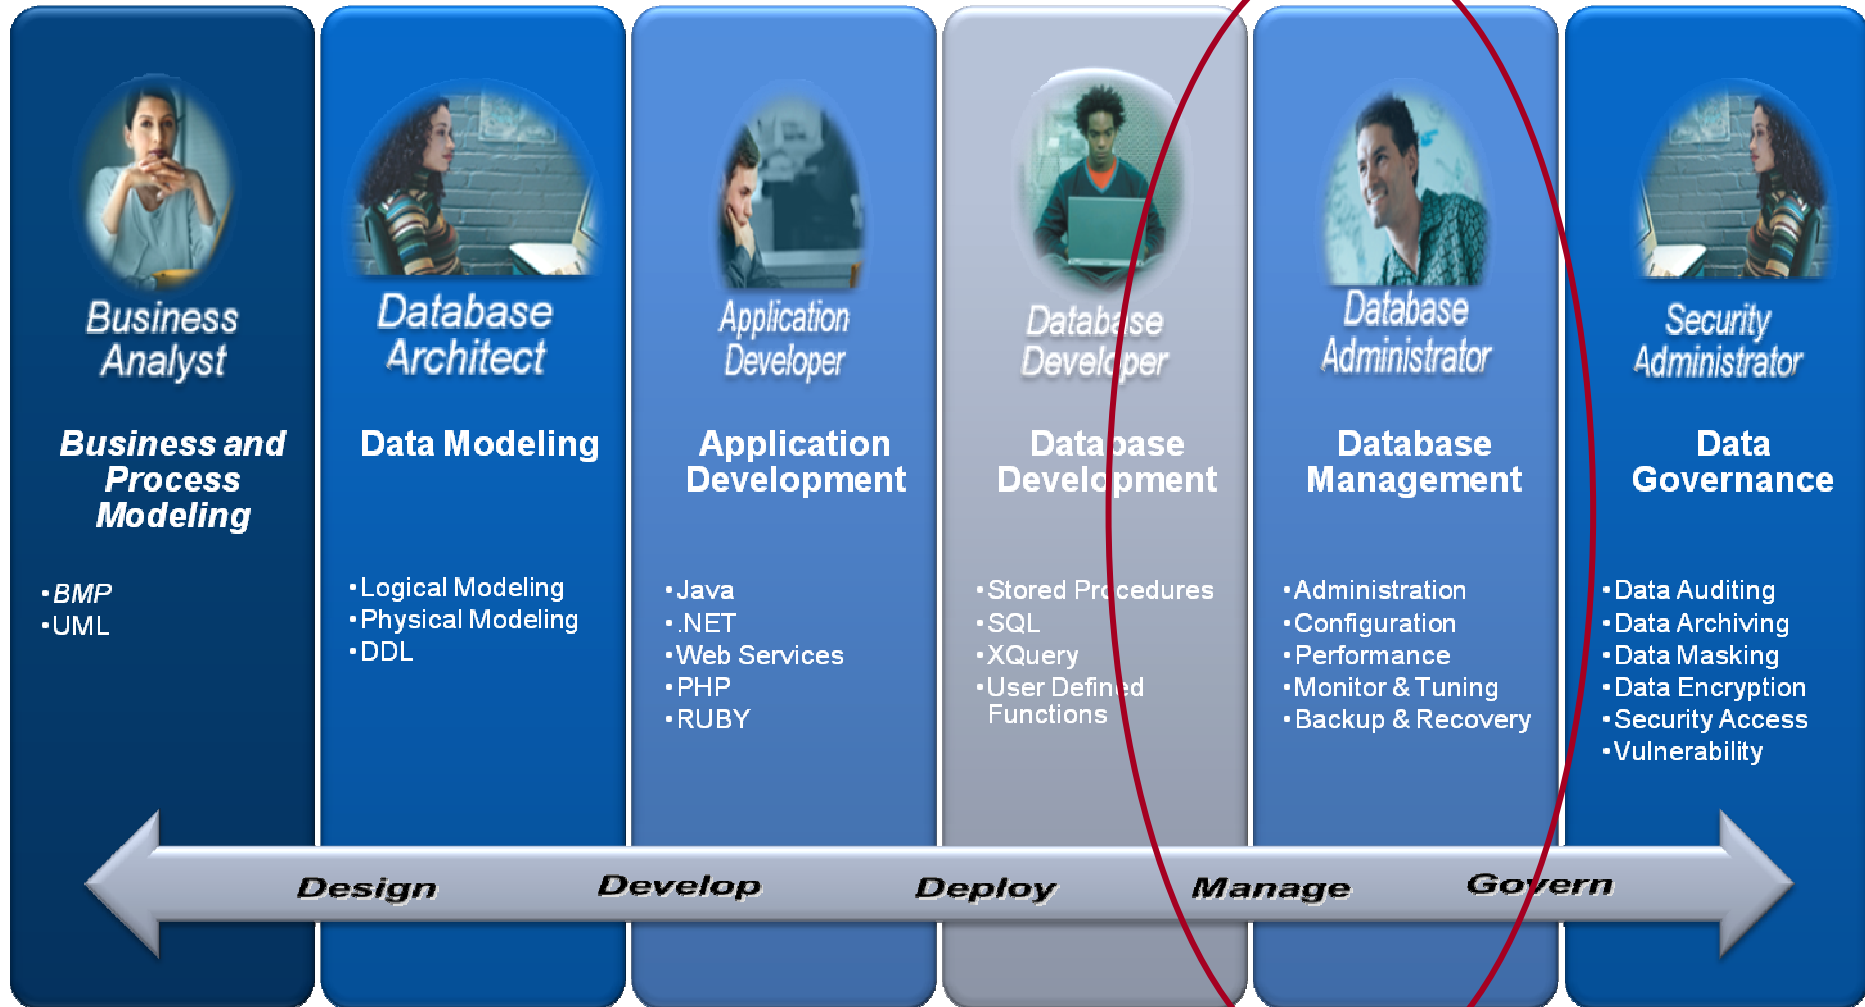
```
class Employee
{ public String Name;
  public String HomeAddress;
  public String HomePhone;
  ...
}
```

投影片 14

ad2

change the color
Anshul Dawra, 2008/10/23

Users of the IBM Data Studio



Optim Database Administrator

Future add-ons

Optim

Additional add-ons

Infosphere Data Architect

- Physical Database Model
- Logical Database Model
- Enforce Enterprise Standard for Data quality and consistency
- Integration with Industry Models
- Generate DDL for Various Data Servers

Optim Database Administrator

- Database Change Management
- Extended Alters
- Advanced Data Migration
- Enhanced Scripting
- Advanced Compare Multiple Objects
- Dependency Support
- Model Driven Change

ClearQuest / CVS

High Performance Unload

- Fast Unload
- Extraction from DB2 full backups
- Repartition feature

Data Studio base (no-charge capabilities)

Data Studio Administration Console

- At-a-Glance Health and Availability
- Problem Determination
- Replication Monitor
- Recommendation
- Limited Historical Information

Optim Development Studio

- ER Diagramming
- Integrated Query Editor
- SQL & Java Routine + Debugger
- XML, XML Schema Editors
- Data Web Services
- Object, Data and Security Management
- Visual Explain
- Export / Import Data
- Generate DDL

Optim Database Administrator

- Enhanced Database Object Management
- Instance Management
- Database Management
- Schema Management
- Table Management
- Configuration Management
- Backup / Restore Database
- Run DB2 Commands

Optim Database Administrator Features

- **Manage and Administer systems**
 - Manage and quickly access database objects
 - Manage database connections
 - Manage DB2 LUW server at instance level
 - Perform preventive maintenance by backing up and restoring databases or table spaces.
 - View / edit configuration parameters

- **Database Change Management**
 - Initiate database object modification from data source explorer
 - External Data Management
 - Auto generate undo
 - Auto generate delta DDL
 - Auto generated maintenance commands
 - Enhance audibility with integrated reporting and documentation of changes and their impact

Optim Database Administration Features

- **Easily find your database objects**
- **Group over a list of connections**
- **Stop or start your DB2 instances or databases**
- **Backup, restore or recover database or table spaces**
- **Unload and load data into tables**
- **Manage user privileges**

Scenario

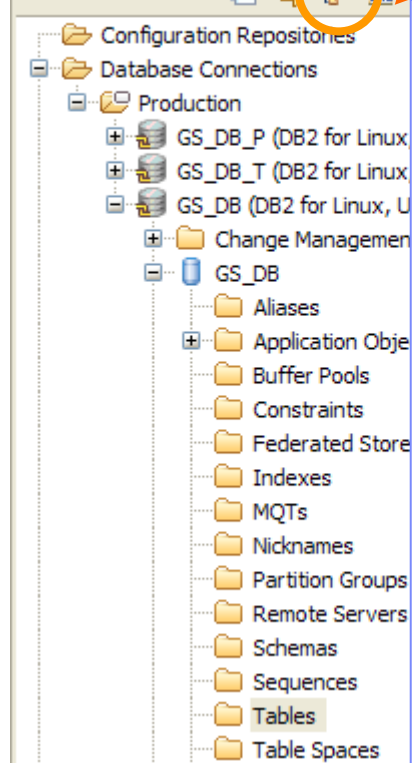
DBA has 100+ schemas and each schema has 20+ tables. He needs to only edit one column in the PROD table.

How is he going to quickly find that table?

Easily Find Database Objects

Flat presentation

One Click



Object List

Sort by
clicking on
the columns

Schema	Name
GOSALES	PRODUCT
GOSALES	ORDER_METHOD
GOSALESCT	PTNR_CNTCT
GOSALESCT	PTNR_ACTY
GOSALESCT	CUST_STATE_TAX
GOSALESCT	CUST_ORD_DETL
GOSALESCT	CUST_ORD
GOSALESCT	CUST_INV
GOSALESCT	CUST_CRDTCRD
GOSALESCT	CUST
SYSIBM	SYSXSROBJECTS
SYSIBM	SYSXSROBJECT...
SYSIBM	SYSXSROBJECT...
SYSIBM	SYSXSROBJECT...
SYSIBM	SYSXMLSTRINGS
SYSIBM	SYSXMLPATHS
SYSIBM	SYSXDBMAPSH...
SYSIBM	SYSXDBMAPGR...
SYSIBM	SYSWRAPPERS
SYSIBM	SYSWRAPOPTI...
SYSIBM	SYSWORKLOADS
SYSIBM	SYSWORKLOAD...

Scenario


**DBA normally has to manage 10-15 database...but
he current project only requires him to use a
subset of the databases**

***How is he going to quickly organize the
databases?***

Manage Database Working Sets

Group Database Together

Configure a connection



Identify a set of database connections you will be working with

New Working Set


Configure Working Sets

Enter a name for your working set as

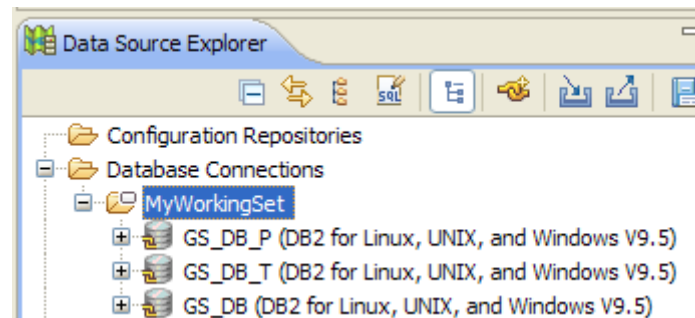
Working set name:

MyWorkingSet

Database Connections



Filter and group over a list of connections to create a working set



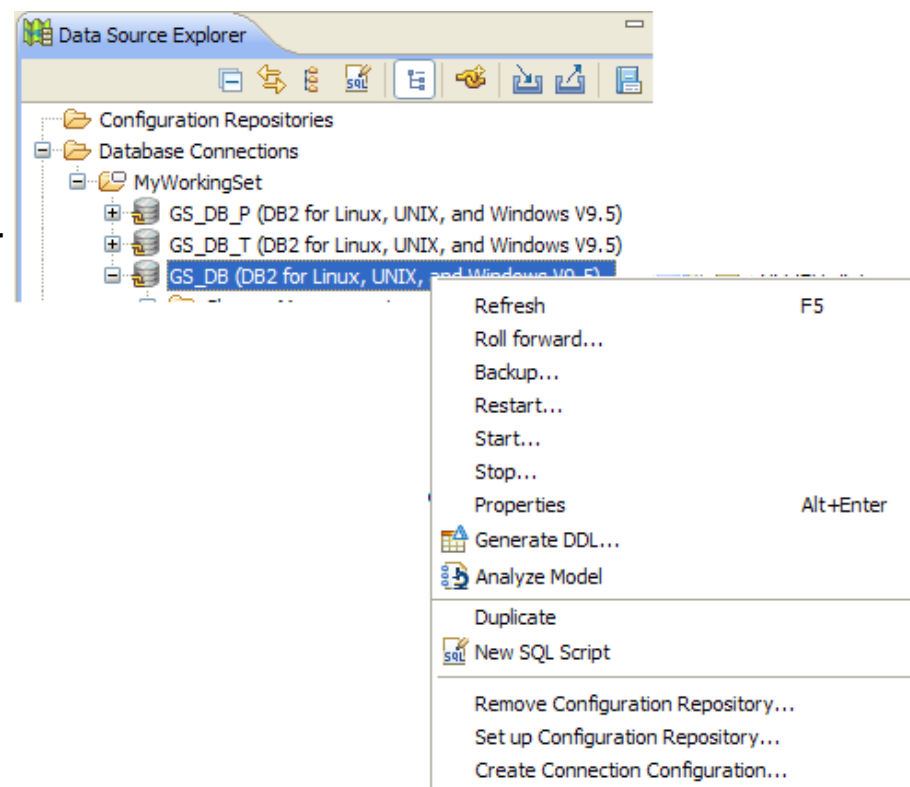
Hostname	Server Instance	Name	Database Alias
localhost	DB2_02	GS_DB_P	GS_DB_P
localhost	DB2_02	GS_DB	GS_DB
localhost	DB2_02	GS_DB_T	GS_DB_T

Scenario

What about my other day to day database administration task?

Simplified Database Administration

- **Using the Data Source Explorer**
 - Stop or start DB2 instances or database
 - Create or drop database
 - Backup, restore or recover database or table spaces
 - Unload and load data into tables
 - Reorganize tables and indexes
 - Manager user privileges



Optim Database Administrator Features

- **Manage and Administer systems**

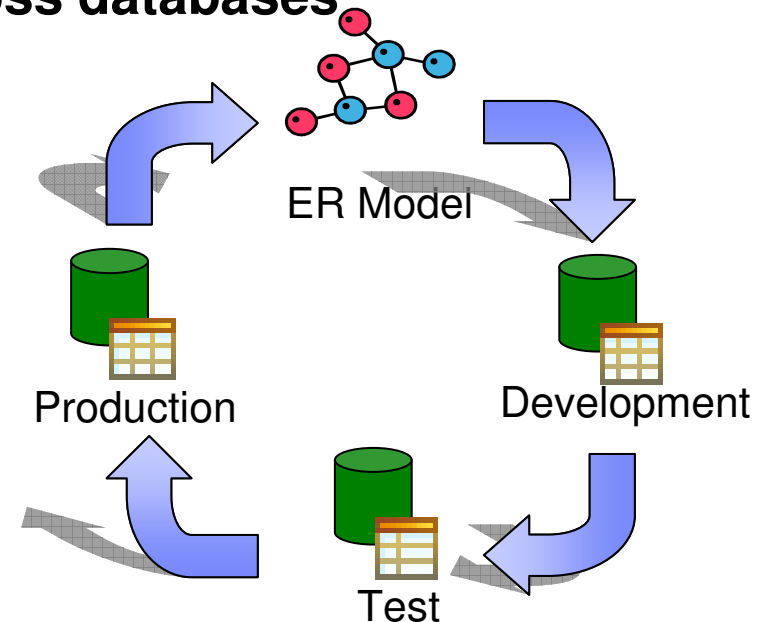
- Manage and quickly access database objects
- Manage database connections
- Manage DB2 LUW server at instance level
- Perform preventive maintenance by backing up and restoring databases or table spaces.
- View / edit configuration parameters

- **Database Change Management**

- Initiate database object modification from data source explorer
- External Data Management
- Auto generate undo
- Auto generate delta DDL
- Auto generated maintenance commands
- Enhance audibility with integrated reporting and documentation of changes and their impact


3 Database Change Scenarios

1. Change driven by data modeling
2. Compare and Synchronize from source to target
3. Simple object copy and paste across databases

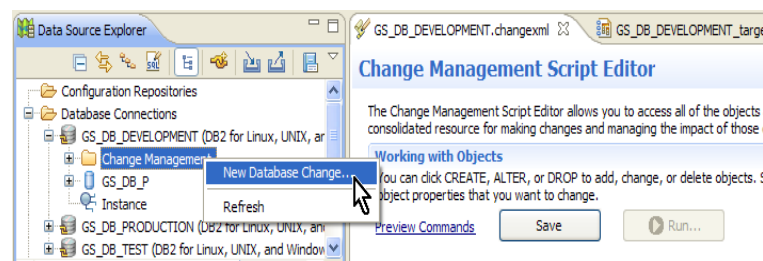


1 - Change driven by data modeling


- Business requirements modeled using Infosphere Data Architect
- Logical data model transformed to physical data model
- Physical data model drives Data Studio Administrator to deploy changes to target database



ER Model



Optim Database Administrator



Target Database

Change driven by data modeling

Select the physical database model as the source

Migrate Objects to Target Model

Migration Source

Select the source from the list

Type of source:

- Physical data model
- Database connection
- DDL script file

Select the source model

- .sqlxedito
- GSDB
- GOSA
- GOSA
- GSDB
- GS_DB_DE
- GS_DB_TE

Identify differences with target and select for deployment

Migrate Objects to Target Model

Review and Apply Changes

The Structural Compare tool allows you to identify differences between the source and target models. Use the down arrows to view the differences.

Structural Compare

Item
[-] Schema
indices
[+] Table
Table
Table
Table
Table
Table
Table
Table
Table
[+] Schema

View impacted objects, preview commands, customize commands and run

*GS_DB_TEST.changexml *GS_DB_TEST_target.dbm GS_DB_TEST_base.dbm

Change Management Script Editor

The Change Management Script Editor allows you to access all of the objects and actions that are relevant to a change. It is a consolidated resource for making changes and managing the impact of those changes.

Working with Objects
You can click CREATE, ALTER, or DROP to add, change, or delete objects. Select an object from the list to view object properties that you want to change.

[Preview Commands](#)

Objects to be Changed:
Select CREATE, ALTER, or DROP to change your objects.

CREATE
 ALTER
 DROP

- GOSALES.PRODUCT_BRAND
- GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_CODE
- GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_EN [VA
- GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_CODE,

<Column> PRODUCT_BRA

General Name:

Type Label:

Documentation

Annotation

Data Preservation

Customize DDL generated by Data Studio Administrator

Commands

The commands that will be issued against your database are listed below. Ensure that the commands are correct and click Run.

Format Open in SQLX Editor... **Customize...** Run... [U](#)

Customize your data preservation options

```
--<ScriptOptions statementTerminator=";">
CONNECT TO GSDB;
EXPORT TO ':\default_GOSALESC_T_CUST_ORD_DETL.dat' OF DEL SELECT * FROM GOSALESC_T.CUST_ORD_DETL;
EXPORT TO ':\default_GOSALESC_T_CUST_ORD.dat' OF DEL SELECT * FROM GOSALESC_T.CUST_ORD;
EX
CO
re
```

Specify the unload method for data preservation

Specify the default unload and reload command method and customize the command options.

Unload

EXPORT for DEL data format Customize

EXPORT for DEL data format

EXPORT for IXF data format

High Performance Unload

Verify generated commands and resolve errors

Select the table name

- GOSALESC_T
 - CUST_ORD_DETL
 - CUST_ORD
 - CUST

Details

Unload: EXPORT for DEL data format Change query

Reload: LOAD Change mappi

Data File Name: .\default_GOSALESC_T_CUST_ORD_DETL.dat Browse...

Commands

```
EXPORT TO ':\default_GOSALESC_T_CUST_ORD_DETL.dat' OF DEL SELECT ORD_DETL_CODE, ORD_NBR, ORD_SHIP_DATE, PROD_NBR,
PROD_PROM_CODE, PROD_QTY, PROD_UNIT_COST, PROD_UNIT_PRC, PROD_UNIT_SALE_PRC FROM GOSALESC_T.CUST_ORD_DETL
LOAD FROM ':\default_GOSALESC_T_CUST_ORD_DETL.dat' OF DEL MODIFIED BY IDENTITYOVERRIDE METHOD P (1, 2, 3, 4, 5, 6, 7, 8, 9)
INSERT INTO GOSALESC_T.CUST_ORD_DETL (ORD_DETL_CODE, ORD_NBR, ORD_SHIP_DATE, PROD_NBR, PROD_PROM_CODE,
PROD_QTY, PROD_UNIT_COST, PROD_UNIT_PRC, PROD_UNIT_SALE_PRC)
```

Default Que Auto Cast

Problems		
Unload Table	Reload Table	Problem
GOSALESC_T...	GOSALESC_T.CU...	Warning: An identity column was found. The DB2 LOAD command is recommended.
GOSALESC_T...	GOSALESC_T.CUST	Warning: An identity column was found. The DB2 LOAD command is recommended.

Reviews and Change Management

■ Reports

- Summary of changes
- Deployment
- Database unit test

Summary of Changes Report

Date (month/day/year): 10/31/2008, 01:29
 Machine URL:
 DB2 Instance Name: DB2

Changes to Database Objects:

- . For TABLE objects, 2 will be created, 2 will be dropped.
- . For PRIMARY KEY objects, 1 will be created.

TABLE

Schema	Name	CREATE	ALTER	DROP	RENAME	GRANT	REVOKE
GOLLUM	MYTAB			X			
GOLLUM	TAB2	X		X			
GOLLUM	EXAMPLETAB	X					

PRIMARY KEY

Schema	Name	CREATE	ALTER	DROP	GRANT	REVOKE
GOLLUM	TAB2_PK	X				

Data Preservation and Maintenance Commands:

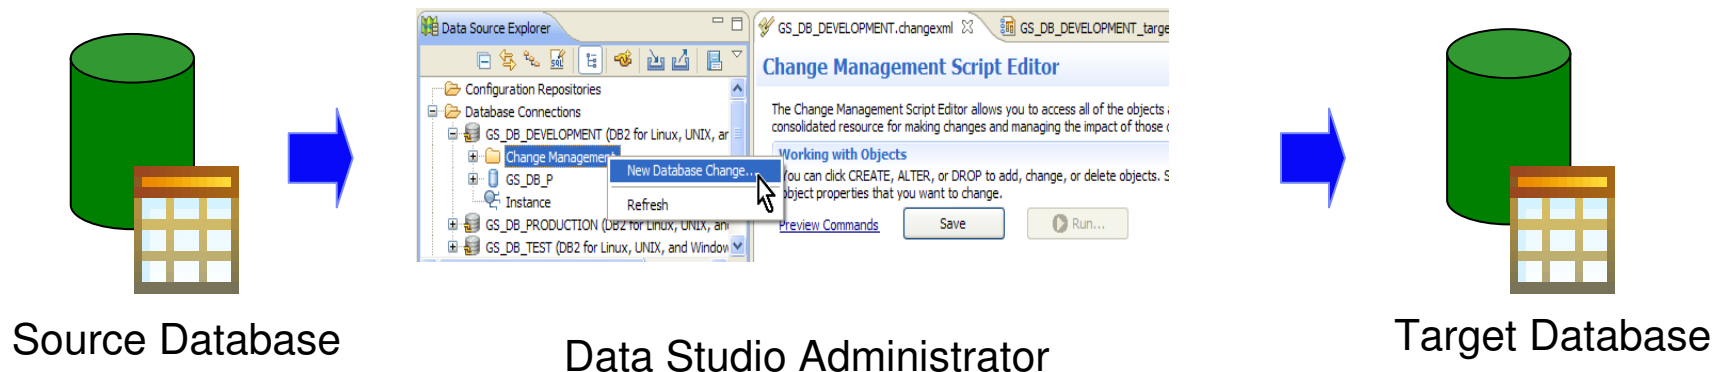
Schema	Name	UNLOAD	RELOAD	REORG	RUNSTATS	REBIND
GOLLUM	TAB2		X			
GOLLUM	MYTAB	X				

ERRORS:

- 1: Table GOLLUM.EXAMPLETAB does not reference a tablespace
- 2: Table cannot have an empty name

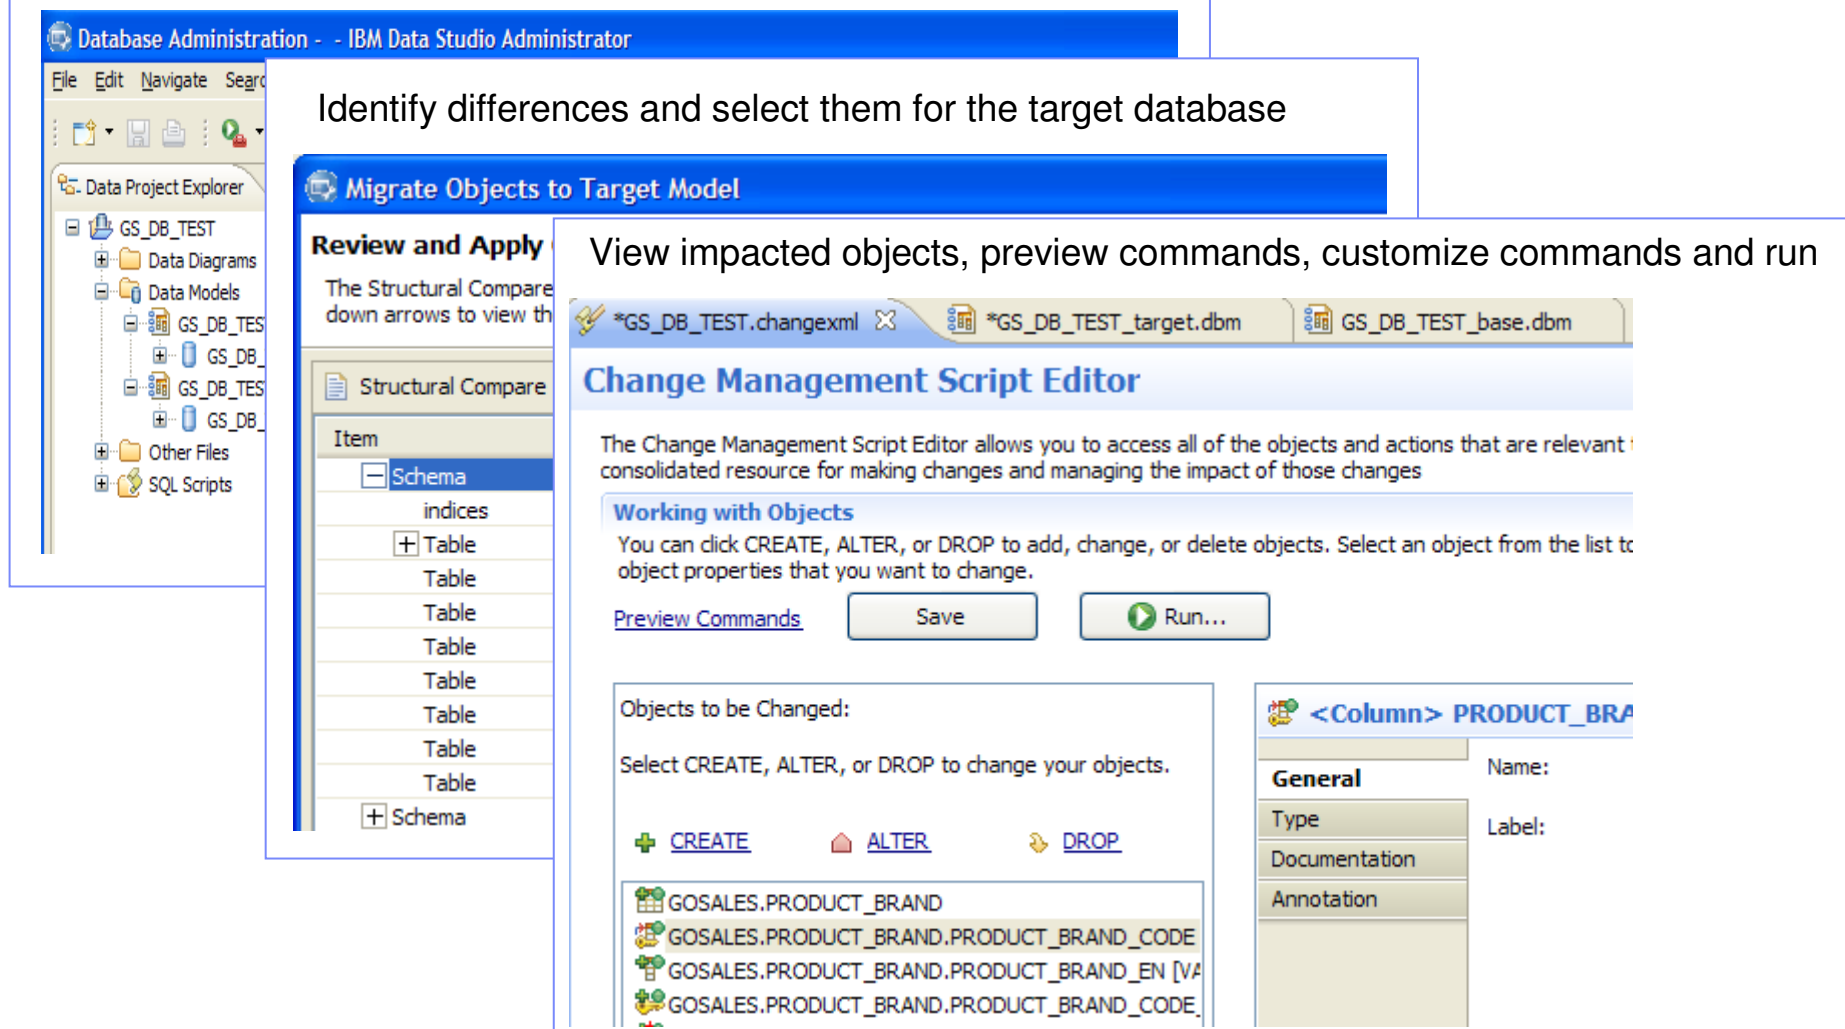
2 – Compare and Synchronize

- **Comprehensive solution for migration between database environments**



Compare and Synchronize Databases

Invoke compare and migrate objects



The screenshot displays the IBM Data Studio Administrator interface. On the left, the 'Data Project Explorer' shows a project named 'GS_DB_TEST' with sub-items for 'Data Diagrams', 'Data Models', and 'SQL Scripts'. The main window is titled 'Migrate Objects to Target Model' and is divided into two panes.

The left pane, titled 'Review and Apply', shows a 'Structural Compare' window. It lists items to be compared, including 'Schema', 'indices', and several 'Table' entries. The right pane, titled 'Change Management Script Editor', displays the results of the comparison. It shows a list of objects to be changed, including 'GOSALES.PRODUCT_BRAND' and its associated columns. The editor provides options to 'Preview Commands', 'Save', or 'Run...' the changes. A detailed view of a column, '<Column> PRODUCT_BRA', is shown on the right, with fields for 'Name' and 'Label'.

Identify differences and select them for the target database

View impacted objects, preview commands, customize commands and run

Database Administration - - IBM Data Studio Administrator

File Edit Navigate Search

Data Project Explorer

GS_DB_TEST

Data Diagrams

Data Models

GS_DB_TEST

GS_DB_TEST

GS_DB_TEST

GS_DB_TEST

Other Files

SQL Scripts

Migrate Objects to Target Model

Review and Apply

The Structural Compare down arrows to view th

Structural Compare

Item

Schema

indices

Table

Table

Table

Table

Table

Table

Table

Table

Table

Table

Schema

*GS_DB_TEST.changexml

*GS_DB_TEST_target.dbm

GS_DB_TEST_base.dbm

Change Management Script Editor

The Change Management Script Editor allows you to access all of the objects and actions that are relevant consolidated resource for making changes and managing the impact of those changes

Working with Objects

You can click CREATE, ALTER, or DROP to add, change, or delete objects. Select an object from the list to object properties that you want to change.

Preview Commands

Save

Run...

Objects to be Changed:

Select CREATE, ALTER, or DROP to change your objects.

CREATE ALTER DROP

GOSALES.PRODUCT_BRAND

GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_CODE

GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_EN [VA

GOSALES.PRODUCT_BRAND.PRODUCT_BRAND_CODE

<Column> PRODUCT_BRA

Name:

Label:

General

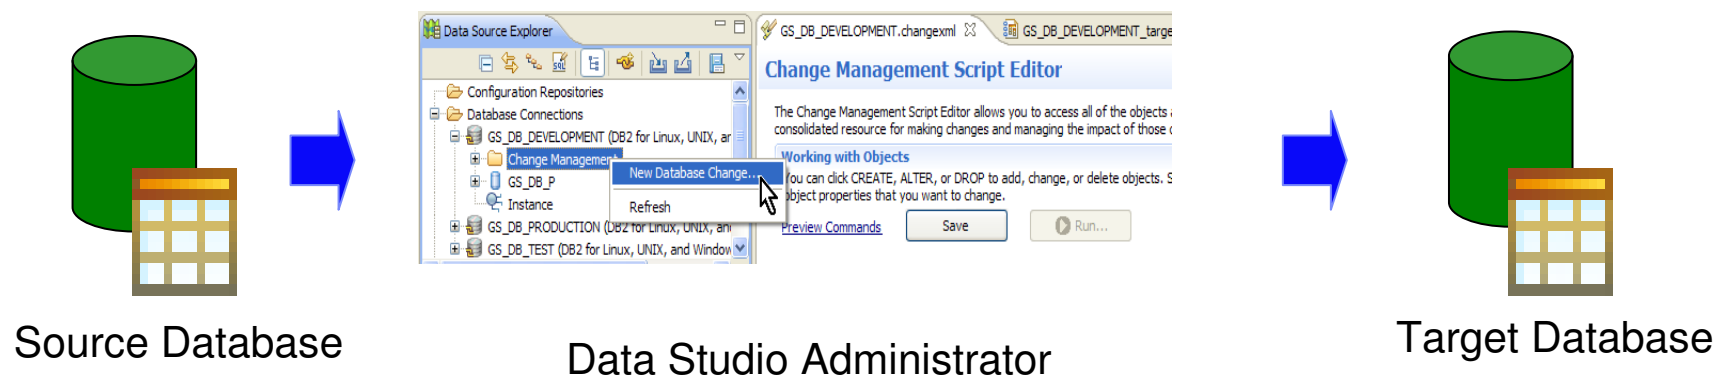
Type

Documentation

Annotation

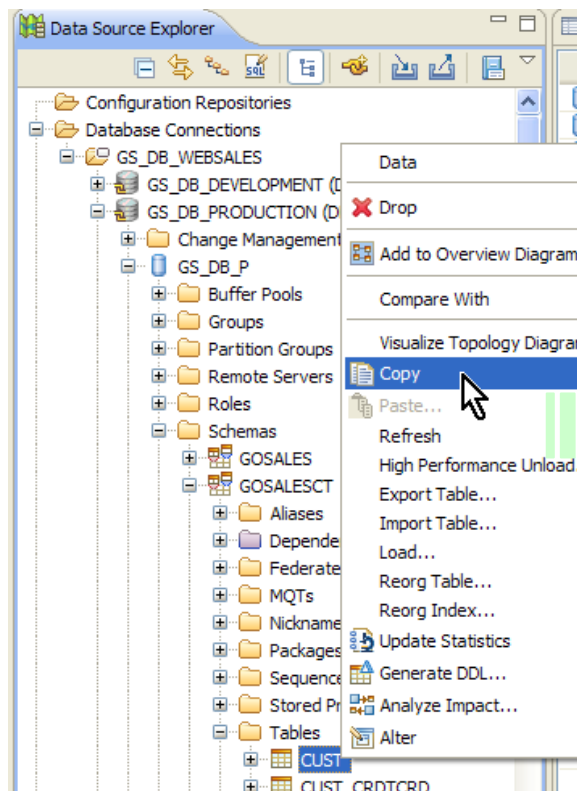
3 – Simple object copy and paste

- **Quick, easy and intuitive way to copy and paste or drag and drop objects from system to system**

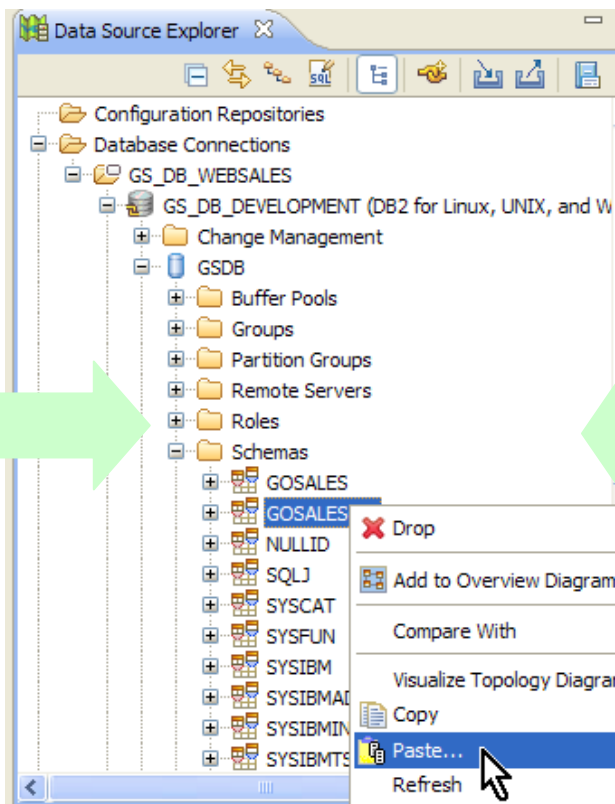


Simple Copy and Paste

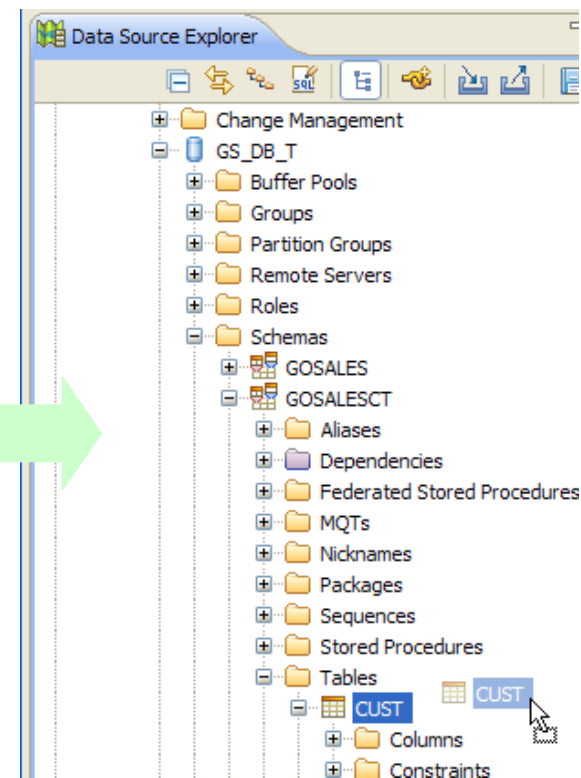
Copy the source object



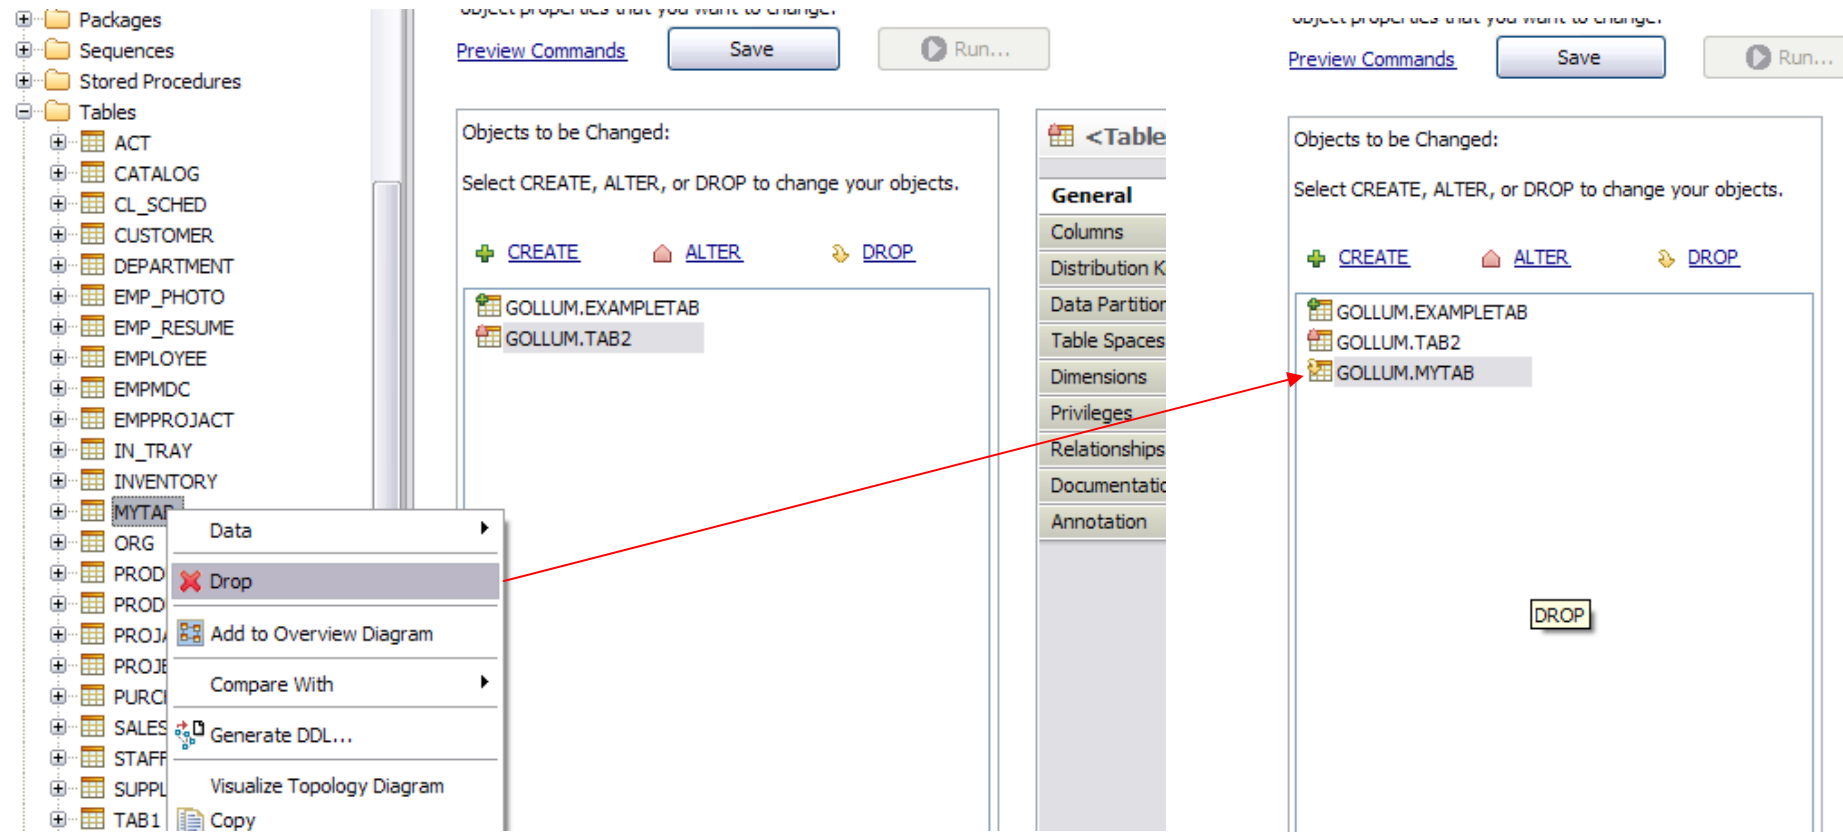
Paste the source object to the target database



Drag-and-drop objects



Drop Objects From Data Source Explorer



The screenshot illustrates the process of dropping a table from the Data Source Explorer in IBM Data Studio. It is divided into two panels showing the state before and after the action.

Left Panel (Before): The Data Source Explorer on the left shows a tree view of database objects under 'Tables'. A context menu is open over the 'MYTAB' object, with the 'Drop' option selected. The 'Objects to be Changed' dialog box is open, showing 'GOLLUM.EXAMPLETAB' and 'GOLLUM.TAB2' selected. The 'DROP' button is highlighted. The 'General' tab is selected in the right-hand pane.

Right Panel (After): The 'Objects to be Changed' dialog box now shows 'GOLLUM.MYTAB' selected. The 'DROP' button is now a floating button in the center of the dialog box, indicating the operation is being executed.

Impact Analysis

Select the object you want to do impact analysis on

Schema	Name	Row
GOSALESC	CUST	195
GOSALESC	CUST_CRDT_CHK	900
GOSALESC	CUST_CRDTCRD	53
GOSALESC	CUST_INV	3
GOSALESC	CUST_ORD	29

Context menu options: Alter, Drop, Copy, Analyze Impact..., Compare

Impact and dependency information can be viewed both in tabular and visual formats

The diagram shows the following dependencies:

- CUST_ORD_STATUS depends on CUST_ORD
- PTNR_ACTY_CUST_ORD_FK depends on CUST_ORD
- CUST_INV_CUST_ORD_FK depends on CUST_ORD
- CUST_ORD_DETL_CUST_ORD_FK depends on CUST_ORD
- P3291456 depends on CUST_ORD
- CUST_ORD impacts CUST_CRDTCRD
- CUST_ORD impacts CUST
- CUST_ORD impacts ORDER_METHOD

Dependent Object	Dependent O...	Impactor Object	Impactor Obj...	Relationship	Comment
GOSDB.GOSALESC...	Foreign Key	GOSDB.GOSALE...	Table	Reference	Customer ord...
GOSDB.GOSALESC...	Table	GOSDB.GOSALE...	Table	Reference	Customer cre...
GOSDB.GOSALESC...	Table	GOSDB.GOSALE...	Table	Reference	Customer inf...
GOSDB.GOSALESC...	Table	GOSDB.GOSALE...	Table	Reference	Different met...

Auto-Generated Undo Commands / Restart From Failure

Undo Commands

Undo Commands

You can run the undo commands that are listed below to reverse the changes that you have made to your database.

Format

Restart Commands

Run Undo...

```
--<ScriptOptions statementTerminator=";">
ALTER TABLE GOSALESCT.CUST_CRDT_CHK DROP PRIMARY KEY ;
ALTER TABLE GOSALESCT.CUST_CRDT_CHK DROP FOREIGN KEY CUST_CRDT_CHK_FK ;
DROP TABLE GOSALESCT.CUST_CRDT_CHK;
```

Run undo commands to reverse your changes

Restart From Failure

Undo Commands

You can run the undo commands that are listed below to reverse the changes that you have made to your database.

Format

Restart Commands

Run Undo...

```
--<ScriptOptions statementTerminator=";">
ALTER TABLE GOSALESCT.CUST DROP FOREIGN KEY CUST_CUST_STATE_TAX_FK ;
ALTER TABLE GOSALESCT.CUST_STATE_TAX DROP PRIMARY KEY ;
DROP TABLE GOSALESCT.CUST_STATE_TAX;
ALTER TABLE GOSALESCT.CUST_CRDTCRD DROP PRIMARY KEY ;
ALTER TABLE GOSALESCT.CUST_CRDTCRD DROP FOREIGN KEY CUST_CRDTCRD_CUST_FK ;
ALTER TABLE GOSALESCT.CUST DROP PRIMARY KEY ;
DROP TABLE GOSALESCT.CUST;
DROP TABLE GOSALESCT.CUST_CRDTCRD;
DROP TABLESPACE GOSALESCTTBLSPC;
```

Restart the commands

IBM Data Studio Administration Console

IBM Data Studio Administration Console provides a health and availability dashboard used to prevent downtime or performance issues before the business is impacted.

- Avoid outages and response time issues
 - Connect to and monitor multiple databases across different platforms from a single console
 - Configure thresholds for warnings and alerts for key performance indicators.
 - Collect and retain alert history for 72 hours
 - Use expert recommendations for problem solving
 - View the system and database parameters and statistics associated with a warning or alert



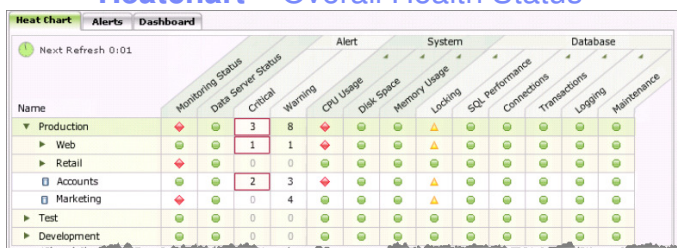
The screenshot shows a dashboard with tabs for 'Heat Chart', 'Alerts', and 'Dashboard'. Below the tabs, there is a 'Next Refresh 0:01' indicator. The main area is a grid with columns for 'Monitoring Status', 'Data Server Status', 'Alert', 'System', and 'Database'. The 'Alert' column is further divided into 'Critical', 'Warning', 'CPU Usage', 'Disk Space', 'Memory Usage', and 'Locking'. The 'System' column includes 'SQL Performance', 'Connections', and 'Transactions'. The 'Database' column includes 'Logging' and 'Maintenance'. The rows represent different databases: Production, Web, Retail, Accounts, Marketing, Test, and Development. The 'Production' row shows a red diamond icon in the 'Monitoring Status' column, a green circle in 'Data Server Status', a red box around the number '3' in the 'Critical' column, and a red diamond in the 'Warning' column. The 'Accounts' row shows a red box around the number '2' in the 'Critical' column and a red diamond in the 'Warning' column. Other rows show various combinations of green circles and yellow triangles.

Name	Monitoring Status		Alert		System		Database	
	Monitoring Status	Data Server Status	Critical	Warning	CPU Usage	Disk Space	Memory Usage	Locking
Production	Red Diamond	Green Circle	3	8	Red Diamond	Green Circle	Green Circle	Yellow Triangle
Web	Green Circle	Green Circle	1	1	Red Diamond	Green Circle	Green Circle	Yellow Triangle
Retail	Red Diamond	Green Circle	0	0	Green Circle	Green Circle	Green Circle	Green Circle
Accounts	Green Circle	Green Circle	2	3	Red Diamond	Green Circle	Green Circle	Yellow Triangle
Marketing	Red Diamond	Green Circle	0	4	Green Circle	Green Circle	Green Circle	Yellow Triangle
Test	Green Circle	Green Circle	0	0	Green Circle	Green Circle	Green Circle	Green Circle
Development	Green Circle	Green Circle	0	0	Green Circle	Green Circle	Green Circle	Green Circle

Administration Console Health Monitoring

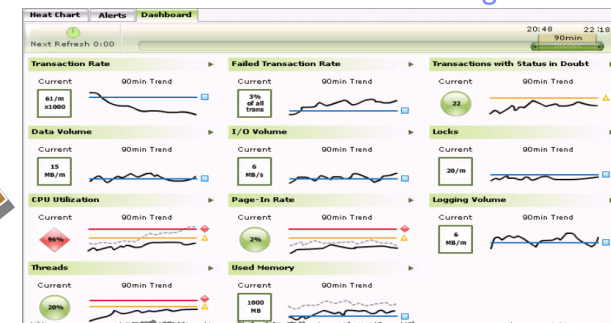
Problem Determination

Heatchart – Overall Health Status



Where are the most important hotspots that need my attention?

Dashboard – Adhoc Investigation



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



Administrator

Alert List – Historical Investigation

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

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

Recommendations – Root Cause Analysis

Table space TS1 in the TA databases off-line

At the time of the alert, the TS1 tablespace is off-line, 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 backonline.

Diagnosing the problem

- 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.

IBM Data Studio Performance Expert

Reduce downtime and quickly resolve performance issues with this comprehensive DB2 monitor.

Provides real-time and history data

DB2 Application monitoring

DB2 Engine monitoring

The screenshot shows the 'DB2 Performance Expert - System Overview' window. The main display area is divided into several sections:

- Sorts:**

Sorts per minute (sum)	46
PEDEMO	32
DB2PMLLOC	14
SAMPLE	0
Sort overflows per min. (sum)	0
Post threshold sorts per minute	0
Average sort time (max)	2ms
Average sorts per UOW (max)	0.13
- Buffer Pools:**

Buffer pool hit ratio (min)	49.8%
Async. read percentage (min)	0%
Avg. rows read per sel. row (max)	18.1
DB2PMLLOC	18.1
PEDEMO	1.34
SAMPLE	N/C
Avg. rows read per UOW (max)	136.19
- Locks:**

Escalations per minute (sum)	0
Timeouts per minute (sum)	0
Deadlocks per minute (sum)	0
Waits per minute (sum)	0
Avg. wait time per lock (max)	N/C
Avg. wait time per appl. & min (max)	2142ms
Avg. wait time per UOW (max)	250ms
Lock wait time per min. (sum)	60s
- Applications / Agents:**

Currently executing (sum)	2
Waiting for locks (sum)	1
Connected / max appls (sum)	20 / 26
PEDEMO	14 / 16
DB2PMLLOC	5 / 7
SAMPLE	1 / 3
Current agents / max agents	36 / 36
Stolen agents per minute	0
- Databases:**

Currently active	3
UOW per minute (sum)	638
Database files closed per minute (sum)	0
Package cache hit ratio (min)	45.6%
PEDEMO	45.6%
DB2PMLLOC	56.8%
SAMPLE	N/C
Catalog cache hit ratio (min)	63.5%
- Monitor Switches:**

Buffer Pool	On	Off
Lock	On	Off
Sort	On	Off
Statement	On	Off
Table	On	Off
Unit of Work	On	Off
Timestamp	On	Off

The bottom status bar shows: Server Status: Logon, DB2 System: Group, User ID: RMR, Exception: both, Trace Status: N/A, Session: 1, Operating S...: VMND0MVS, System Name: localhost, DB2: V8R2FP2, Server: V2.2, Host: localhost, Port: 0, Description: local DB2 in...

Performance Expert Advanced Features

Provides real-time and history data

DB2 Application monitoring

- Monitor resource consumption, SQL Activity, transactions
- Identify locking conflicts, heavy hitter applications and SQL statements
- Force applications

DB2 engine monitoring

- Provide metrics on the instance, database, table spaces, buffer pools, tables
- Provide DB and DBM configuration
- Determine potential and existing bottlenecks

Integrated OS monitoring

Performance mgmt of many partitions

- Single partition, aggregated and customized multi partition views
- Problem partition and skew detection

Proactive performance problems alerts

SQL Heavy Hitter Identification

SQL Tracing & Analysis

Performance History and Warehouse

- Post-mortem analysis, trend analysis

WLM Monitoring (statistics and activities)

Performance Expert - System Overview

DB2 Performance Expert - System Overview

Monitor Selected View Tools Window Help

Refresh rate 01:00

All Instances

- CM4WAYB
- DB2
 - DB2PMLOC
 - PEDEMO
 - SAMPLE
 - DB28_FROM_DB29
 - DB29_FROM_DB28
 - PEDEMO
 - DB2_01
 - MYAIX
 - PMON8M1
 - TURMALIN
- My Shortcuts

KPIs Data Views Exceptions

Sorts

Sorts per minute (sum)		46
PEDEMO		32
DB2PMLOC		14
SAMPLE		0
Sort overflows per min. (sum)		0
Post threshold sorts per minute		0
Average sort time (max)		2ms
Average sorts per UOW (max)		0.13

Applications / Agents

Currently executing (sum)		2
Waiting for locks (sum)		1
Connected / max appls (sum)		20 / 26
PEDEMO		14 / 16
DB2PMLOC		5 / 7
SAMPLE		1 / 3
Current agents / max agents		36 / 36
Stolen agents per minute		0

Buffer Pools

Buffer pool hit ratio (min)		49.8%
Async. read percentage (min)		0%
Avg. rows read per sel. row (max)		18.1
DB2PMLOC		18.1
PEDEMO		1.34
SAMPLE		N/C
Avg. rows read per UOW (max)		136.19

Databases

Currently active		3
UOW per minute (sum)		638
Database files closed per minute (sum)		0
Package cache hit ratio (min)		45.6%
PEDEMO		45.6%
DB2PMLOC		56.8%
SAMPLE		N/C
Catalog cache hit ratio (min)		63.5%

Locks

Escalations per minute (sum)		0
Timeouts per minute (sum)		0
Deadlocks per minute (sum)		0
Waits per minute (sum)		0
Avg. wait time per lock (max)		N/C
Avg. wait time per appl. & min (max)		2142ms
Avg. wait time per UOW (max)		250ms
Lock wait time per min. (sum)		60s

Monitor Switches

Buffer Pool	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Lock	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Sort	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Statement	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Table	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Unit of Work	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off
Timestamp	<input checked="" type="checkbox"/> On	<input type="checkbox"/> Off

Server Status Logon DB2 System Group User ID Exception Trace Status Session Operating S... System Name DB2 Server Host Port Description

DR2	RMR	hnh	N/A	1	WINDOWS	localhost	VRR2FP2	V2 2	localhost	0	Local DB2 in
-----	-----	-----	-----	---	---------	-----------	---------	------	-----------	---	--------------

Diagnose Application High Response Time

The sales_shopping_cart application has a very high response time, and most of this time is spent on the data server.

Double-click this application to drill-down into details.

Database: SALES

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Database Total (SALES)	300,000	1,661,1...	5.537	5.150	0.340	-	-
Sales Portal Applications	100%	100%	5.537	5.150	0.340	3%	17%
Sales Portal Application Servers	100%	100%	5.537	5.150	0.340	4%	14%

End-to-end time: SALES

Histogram: SALES

Sales Portal Applications

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Total	300,000	1,661,1...	5.537	5.150	0.340	3%	17%
sales_shopping_cart	69%	94%	7.549	7.210	0.288	5%	25%
sales_view_items	9%	1%	0.621	0.312	0.301	0%	0%
sales_purchase	22%	5%	1.238	0.670	0.519	0%	0%

End-to-end time: sales_shop...

Histogram: sales_shopping_cart

Sales Portal Application Server

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Total	300,000	1,661,1...	5.537	5.150	0.340	4%	17%
sales_portal1.jk_enterprise.com	48%	49%	5.350	4.880	0.321	4%	18%
sales_portal2.jk_enterprise.com	52%	51%	5.740	5.333	0.365	3%	16%

End-to-end time: sales_po...

Histogram: sales_portal1.jk_e...

Both application servers have about the same response time.

The blue line shows the average time, the grey areas indicate the maximum value.

Diagnose Time Spent

LIMETTE_59930_INSTANCE - End-to-End Details

End-to-End Details Selected View Tools Window Help

Data: **History** 05/28/2008 18:30:36 Since 05/28/2008 18:00:36

Aggregation: 1 minute Refresh: **Manual**

Main sales_shopping_cart

sales_shopping_cart

Legend

- Data server
- Sorting
- Network
- Driver processing
- Driver agent wait
- WS connection pool wait
- Application

Show data as

- Graphical View
- Text View

Distribution of end-to-end response time (s)

Transactions (x 1,000)

Physical reads (x 1,000,000)

Network bandwidth (MB/s)

Top SQL statements Show top 3 by Avg. end-to-end response

Statement text	Occurrences	End-to-end time
SELECT * FROM sales.customer AS cust, sales.order AS or...	51,000	20.143
SELECT a.schema, b.name FROM sysi...	37,000	1.408
DELETE FROM account WHERE aid = 3...	42,000	1.227

Top clients Show top 3 by Problems

Client	Problems
sales.portal1.jk-enterprise.com	18%
sales.portal2.jk-enterprise.com	16%

Buffer pool and sorting

- Buffer pool hit ratio (%)
- Buffer pool physical reads per min.
- Rows read per selected row (avg.)
- Sorts per minute
- Sort overflows

You see that most of the time is spent in the data server layer. Also note the high rate of physical reads.

This particular SQL statement has an extremely high response time. Double-click drill-down into details.

Diagnose SQL Statement

LIMETTE_59930_INSTANCE - End-to-End Details

End-to-End Details Selected View Tools Window Help

Data: **History** 05/28/2008 18:30:36 Since 05/28/2008 18:00:36

Aggregation: 1 minute Refresh: **Manual**

Main sales_shopping_cart **SQL Statement Information - SELECT * FROM**

SQL Statement Information

Statement text

```
SELECT * FROM
sales.customer AS cust,
sales.order AS order,
WHERE cust.custID = order.custID
AND order.itemID = items.itemID
AND order.status = "OPEN"
ORDER BY order.orderTS;
```

Statement details

Executions	51,000
Buffer pool hit ratio (%)	1.600
Rows read per selected row	200,000

	Total	Average
End-to-end response time (s)	1,020,000	20.000
CPU time (s)	255,000	5.000
Rows read	96,900,000,000	1,900,000
Rows written	0	0
Rows selected	255,000	5.000
Sorts	20,400	0.400
Sort overflows	0	0

Buffer pool

	Data	Index	XDA
Hit ratio (%)	1.600	0	0
Physical reads	95,625,000,000	0	0
Avg. Physical reads	1,875,000	0	0

Distribution of times

Data server other	91%
Data server sorting	3%
Network	5%
Driver	1%
Application	1%

Buttons: Explain, View Statement in New Window

This button lets you launch Visual Explain

The low buffer pool hit ratio, high number of rows read (per selected row) and table space scan indicated in Visual Explain suggest a missing index as the cause of this problem.

Diagnose Overall Response Time Problem

DB2 Performance Expert - System Overview

Monitor Selected View Tools Window Help

Data: Recent Refresh: Off

Time Period: Last 30 min. 18:00:36 - 18:30:36

End-to-End KPIs Data Server KPIs Data Views Exceptions

Database: SALES

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Database Total (SALES)	300,000	1,068,0...	3.350	0.345	0.318	-	-
Sales Portal Applications	100%	100%	3.350	0.345	0.318	3%	17%
Sales Portal Application Servers	100%	100%	3.350	0.345	0.318	3%	32%

End-to-end time: SALES Histogram: SALES

Sales Portal Applications

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Total	300,000	1,068,0...	3.560	0.345	0.318	3%	17%
sales_shopping_cart	69%	94%	3.694	0.210	0.288	5%	21%
sales_view_items	9%	1%	3.321	0.312	0.301	2%	12%
sales_purchase	22%	5%	3.238	0.470	0.419	3%	15%

End-to-end time: sales_shop... Histogram: sales_shopping_cart

Sales Portal Application Server

Workload Cluster	Transactions	Total End-to-End Time	Avg. End-to-End Time	Avg. Data Server Time	Avg. Network Time	Warnings	Problems
Total	300,000	1,068,0...	3.560	0.345	0.318	2%	15%
sales_portal1.jk_enterprise.com	48%	49%	6.468	0.359	0.271	3%	32%
sales_portal2.jk_enterprise.com	52%	51%	0.876	0.331	0.365	0%	0%

End-to-end time: sales_po... Histogram: sales_portal1.jk_e...

It seems that the first application server has a problem. Double-click to drill-down.

In this situation, all applications are equally affected, and the problem seems not to be in the data server.

Diagnose Application Server

LIMETTE_59930_INSTANCE - End-to-End Details

End-to-End Details Selected View Tools Window Help

Data: **History** 05/28/2008 18:30:36 Since 05/28/2008 18:00:36

Aggregation: 1 minute Refresh: **Manual**

Main sales.portal1.jk-enterprise.com

sales.portal1.jk-enterprise.com

Legend

- Data server
- Sorting
- Network
- Driver processing
- Driver agent wait
- WS connection pool wait
- Application

Show data as

- Graphical View
- Text View

Distribution of end-to-end response time (s)

Transactions (x 1,000)

Physical reads (x 1,000,000)

Network bandwidth (MB/s)

Top SQL statements Show top 3 by Avg. end-to-end response

Statement text	Occurrences	End-to-end time
SELECT * FROM sales.customer AS cust, sales.order AS or...	51,000	2.344
SELECT a.schema, b.name FROM sysi...	37,000	2.308
DELETE FROM account WHERE aid = 3...	42,000	2.227

Top clients Show top 3 by Problems

Client	Problems
sales.portal1.jk-enterprise.com	32%

Buffer pool and sorting

Buffer pool hit ratio (%)	43.400
Buffer pool physical reads per min.	1,900,000
Rows read per selected row (avg.)	11.520
Sorts per minute	8.647
Sort overflows	0

Most of the time is spent for „WAS connection pool wait“ time.

Double-click to drill-down and display detail information.

Diagnose Connection Pool

LIMETTE_59930_INSTANCE - End-to-End Details

End-to-End Details Selected View Tools Window Help

Data: **History** 05/28/2008 18:30:36 Since 05/28/2008 18:00:36

Aggregation: 1 minute Refresh: **Manual**

Main sales.portal1.jk-enterprise.com Client Information - sales.portal1.jk-enterprise.com

Client Information

- Problems (%) 32
- Warnings (%) 3
- Transactions per minute 300,000

Statement details

- Host name sales.portal1.jk-enterprise.com
- IP address 9.152.344.081
- Authentication ID YGH6E
- Driver level 1.0.3
- Connection start time 10/10/2007 06:43:23
- JVM version 1.5.1
- Operating system Microsoft Windows XP Profes...

System utilization

- CPU Usages (%) 56
- Memory usage (%) 81
- Pages swapped out per second 209
- Client up time 10/10/2007 06:40:52

Global transport pool

- Max. allowed transport objects 20
- Transactions rejected (%) 0
- Transactions slowed down (%) 0
- Avg. transaction wait time (s) 0
- Idle global transport pool hit ratio (%) 84
- Idle global transport pool size 15

Distribution of time (s)

- Data server 14%
- Sorting 1%
- Network 15%
- Driver processing 1%
- Driver agent wait 0%
- WAS connect. pool ... 67%
- Application 2%

Top applications

Name	CPU Usage (%)	Memory Usage (%)
db2pb.exe	16.000	14.200
javaw.exe	15.000	8.100
nlnotes.exe	11.000	2.500

Statement details

- Application server name salesnode1
- Connection pool size (max.) 17
- Connection pool size high water m... 17
- Current free connections 0
- Current used connections 17
- Used connections (avg.) 15.7
- Max. connection pool wait time (s) 4.8

Comparison with other clients

Name	Avg. Network Time (%)	CPU Usage (%)	Avg. Driver Wait Time	Avg. WAS Connection Pool Wait	Max. Allowed Connections	Max. Allowed Transport	Network Driver Level	Virtual Machine Version
sales.portal1.jk-enterprise.com	0.271	56.000	0.071	4.339	17.000	20.000	9.5.1	1.5.1.2
sales.portal2.jk-enterprise.com	0.365	62.000	0.082	0.723	20.000	20.000	9.5.1	1.5.1.2

These metrics indicate that the maximum number of allowed transactions is not sufficient...

... which becomes also evident when comparing the parameters and metrics of this client with other clients.

Data Studio Optimization Expert

IBM DB2 Optimization Expert tunes your SQL in context with your workload.

- **DB2 Optimization Expert offers a comprehensive set of expert advisors that can help improve system performance and lower total cost of ownership.**
 - Locates performance problems by being able to view query activity
 - Provides expert tuning advice to improve the performance of an SQL query
 - DB2 Optimization Expert provides profile-based monitoring and tuning for SQL statements that run as part of a workload on DB2 9 for z/OS or DB2 for z/OS Version 8 subsystems.
 - The latest release of DB2 Optimization includes support for DB2 data sharing environments and an improved user experience.

Query Advisor

Query Advisor Recommendations List

Select a recommendation to view more details. To implement the recommendation, change the SQL statement in the source from which it came.

No	Severity	Confidence	Line Number	Description	Suppress
1	Medium	Medium	5	The following predicate on column O_ORDE...	Suppress Rule
2	Low	Low	4	Consider rewriting the following predicate t...	Suppress Rule
3	Low	Low	2	Consider replacing the asterisk (*) or the lo...	Suppress Rule
4	Low	Low	2	Consider replacing the asterisk (*) or the lo...	Suppress Rule

SQL Text

```
SELECT *
FROM SYSADM.ORDER
, SYSADM.LINEITEM
WHERE ( SYSADM.ORDER.O_CLERK != 'abc'
AND ( SYSADM.ORDER.O_ORDERKEY + 3 ) = ( SYSADM.LINEITEM.L_DISCOUNT * 4 )
)
```

Selected Recommendation:

Description	Explanation
The following predicate on column O_ORDERKEY in table SYSADM.ORDER has better selectivity than other predicates in this query. Consider rewriting it so that it is either an indexable or a stage-1 predicate: (SYSADM.ORDER.O_ORDERKEY + 3) = (SYSADM.LINEITEM.L_DISCOUNT * 4) to filter out unnecessary rows earlier. Check the explanation for this warning for more details about possible impact and examples.	The specified SQL statement might perform faster if you rewrite the stage 2 predicate as an indexable predicate or as a stage 1 predicate. Stage 1 predicates are better than stage 2 predicates because they disqualify rows earlier and reduce the amount of processing that DB2 needs to perform during later stages of evaluation. Also, because processing of stage 2 predicates can take many CPU cycles, these predicates are generally slower than stage 1 predicates. Indexable predicates, which are predicates that match existing indexes, are all stage 1 predicates. (Stage 2 predicates are sometimes called residual predicates, and stage 1 predicates are sometimes called sargable predicates.) For more information about what makes a predicate stage 1 or stage 2, search for information on predicate properties in the Information Management Software for z/OS Solutions Information Center at http://publib.boulder.ibm.com/infocenter/imzic . Generally the predicate type depends on the structure of the predicate (for example, a range predicate, an IN-list predicate, a subquery predicate, and so on).

Annotations:

- Predicate that should be considered for re-write to get better performance:** Points to the predicate `(SYSADM.ORDER.O_ORDERKEY + 3) = (SYSADM.LINEITEM.L_DISCOUNT * 4)` in the SQL text.
- Re-write advice and details:** Points to the 'Description' and 'Explanation' tabs in the 'Selected Recommendation' section.

Stats Advisor

IBM DB2 Optimization Expert for z/OS

Project Tools Help

Project Navigator

- Welcome
- Configure Subsystems
- View Monitors
- View Queries
- View Workloads
- SimpleQuerywithnostats
 - Project
 - Query
 - Access Path Advisor
 - Access Plan Graph
 - Advisors
 - Query Advisor
 - Query Annotation
 - Report
- Statistics Advisor
- complexquerywithstats
 - Project
 - Query
 - Access Plan Graph
 - Advisors
 - Query Annotation
 - Statistics Advisor
- simpwithpartstats
 - Project
 - Query
 - Access Plan Graph
 - Advisors
 - Index Advisor
 - Query Annotation
 - Report
 - Statistics Advisor

Configure Subsystems View Queries View Workloads View Monitors complexquerywithstats SimpleQuerywithnostats simpwithpartstats

Recommendations

Existing statistics status - 2 tables need repair out of the 2 tables that were checked

Select a recommendation to view more details.

Number	Priority	Recommendation	Description
1	High	Run repairing RUNSTATS	Repair statistics problems for this query. Ga...
2	Low	Run complete RUNSTATS	Gather and recollect all of relevant statistic...

Selected Recommendations: No.2: Run complete RUNSTATS

View more details for this recommendation, run the statement, or copy the statement for later use.

RUNSTATS Control Statements	Description	Actions
<pre>RUNSTATS TABLESPACE DB4ORDER.TSORDER INDEX(SYSADM.UXO@CKOKODSP KEYCARD, SYSADM.UXO#CLOKOD KEYCARD, SYSADM.PXO@OKODCKSPOP KEYCARD) SHRLEVEL CHANGE REPORT YES RUNSTATS TABLESPACE DB4LINEI.TSLINEI INDEX(SYSADM.SXL@PKSKOKEPDSQN KEYCARD, SYSADM.PXL@OKSDRFSKEPDC KEYCARD) SHRLEVEL CHANGE REPORT YES</pre>	Gather and recollect all of relevant statistics for this query.	<input type="button" value="Details..."/> <input type="button" value="Run..."/> <input type="button" value="Copy"/> <input type="button" value="Save"/>

Project Query Query Annotation Path Advisor Advisors Statistics Advisor Details

Generating Query Report

start

Address

4:28 PM Wednesday 3/26/2008

100%

Infprint Manager ... Command Prompt Command Prompt Venkatesh Gopal - I... Lotus Sametim... db2oe.pdf - Adobe ... IBM Information Ma... OE presentation.ppt IBM DB2 Optimizati...

Annotation showed missing stats
Stats Advisor tells what stats to collect
Run the commands or Save them for later

Index Advisor

Index Advisor Recommendations

Customized and Recommended Indexes

Feature Details	Creator	Object Name	Columns	Estimated I
LINEITEM	DB2OE	LINEITEM_VIRT_IDX...	L_PARTKEY(ASC) ,L_SUPPKEY(ASC) ,L_COMMENT(ASC)	12030.9101
Index				

Existing indexes

Feature Details	Object Name	Columns
LINEITEM		
Index	PXL@OKSDRFSKEPDC	L_ORDERKEY(ASC) ,L_SHIPDATE(ASC) ,L_RETURNFLAG
Index	SXL@PKSKEPDSQN	L_PARTKEY(ASC) ,L_SUPPKEY(ASC) ,L_ORDERKEY(ASC)

DDL Details

```
CREATE INDEX "DB2OE"."LINEITEM_VIRT_IDX_1180338655265" ON
"SYSADM"."LINEITEM" ("L_PARTKEY" ASC, "L_SUPPKEY" ASC,
"L_COMMENT" ASC) FREEPAGE 0 PCTFREE 10;
```

Access Plan Graph

The screenshot displays the IBM DB2 Optimization Expert for z/OS interface. The main window shows an Access Plan Graph for a query. The graph consists of the following nodes:

- QUERY** (Yellow circle)
- QB1** (Blue circle, 3840000.0)
- NLJOIN** (Blue circle, 3840000.0)
- TBSCAN** (Green circle, 10000.0) for **LINEITEM** (10000/default)
- TBSCAN** (Green circle, 38401855) for **ORDER** (10000/default)

A callout box labeled "Type of scan" points to the TBSCAN nodes. The SQL Statement window shows the following query:

```
select * from ORDER,LINEITEM where
ORDER.O_Orderkey+3=LINEITEM.L_DISCOUNT*4 and
ORDER.O_CLERK <> 'abc';
```

The Node Descriptor window shows the following details:

- Table Scan: TBSCAN
- rscan
 - Stage1_Predicates
 - Stage2_Predicates

The Attribute explanation window shows the following details:

- Name
- Input Cardinality
- Scanned Rows
- Stage 1 Predicates
- SYSADM.ORDER.O_CLERK <> 'abc'
- Stage 1 Returned Rows
- Stage 2 Predicates
- (SYSADM.ORDER.O_ORDERKEY+3)=(SYSADM.LINEITEM.L_DISCOUNT*4)
- Stage 2 Returned Rows
- Output Cardinality
- Stage 1 Columns
- Prefetch
- Page Range

The bottom of the window shows the Project Navigator, Search Node, Bookmarks and History, and Graph/Plan Table tabs. The system tray at the bottom indicates the time is 8:48 PM on Wednesday, 3/26/2008, with 99% battery life.

IBM Data Studio

Getting started

Downloads

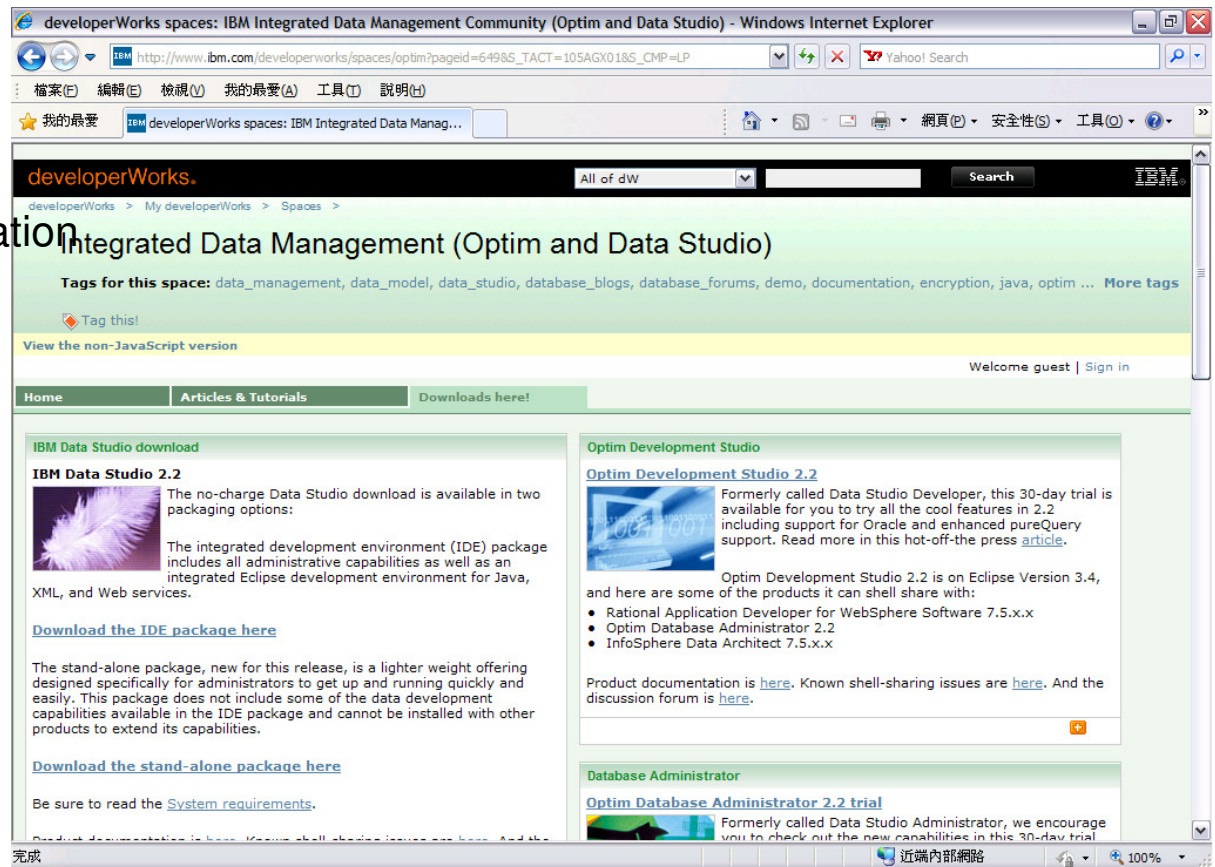
- IBM Data Studio
- IBM Data Studio Administration Console

Community

- developerWorks
 - forums
 - zone and space
 - articles
 - tutorials

Documentation

Support

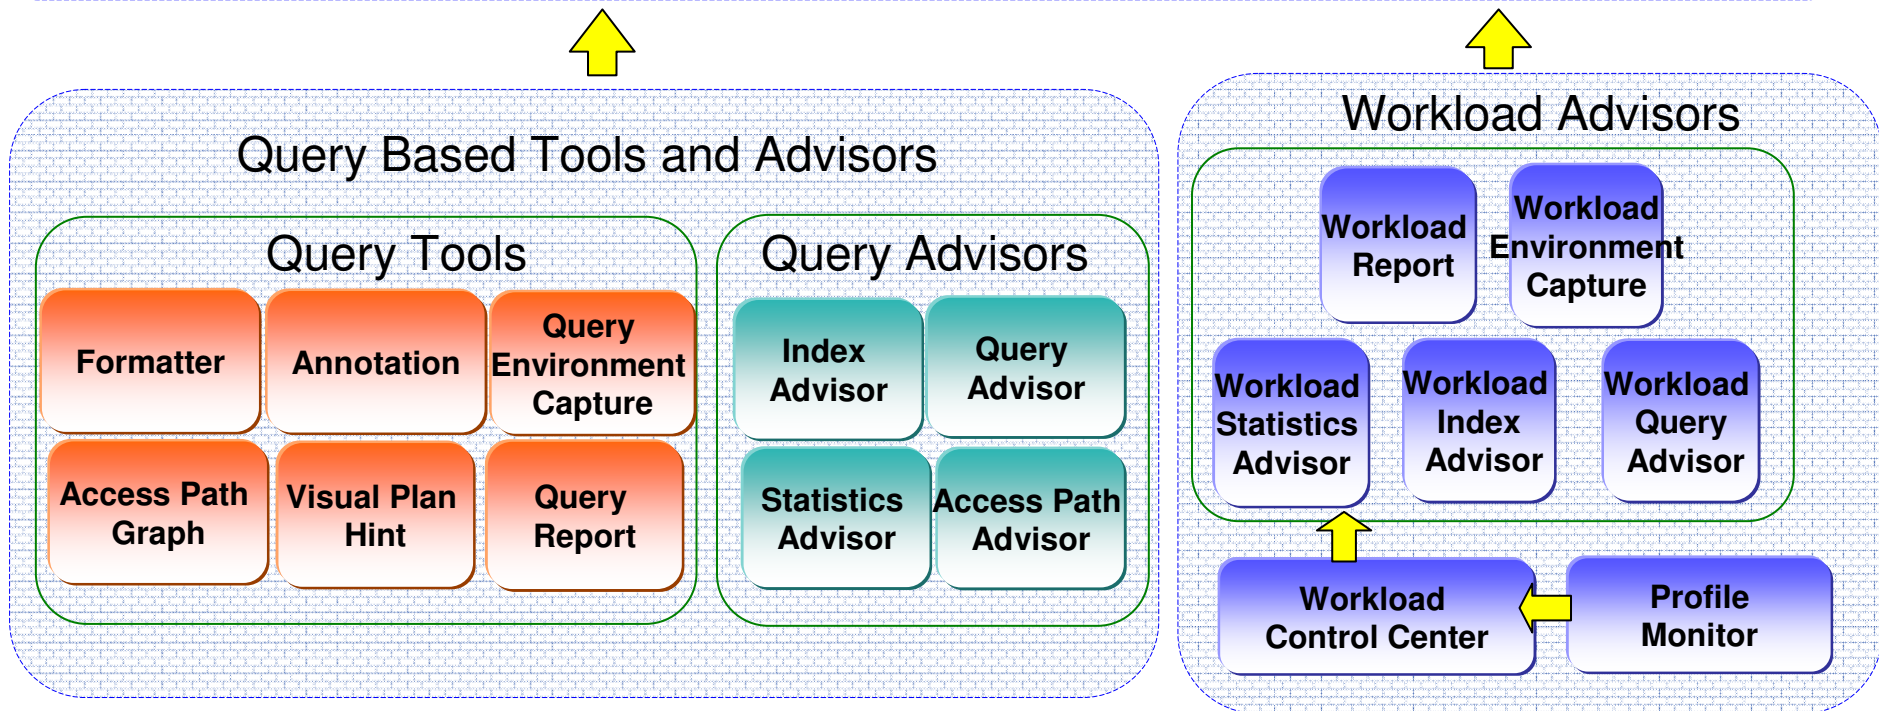




http://www.ibm.com/developerworks/spaces/optim?pageid=649&S_TACT=105AGX01&S_CMP=LP

THANK
YOU



Optimization Expert Key Features at a Glance



Optim Consolidation – Product Names

Existing Name	Rebranded Name
IBM Data Studio	IBM Data Studio
IBM InfoSphere Data Architect	IBM InfoSphere Data Architect
IBM Data Studio Developer	IBM Optim Development Studio
IBM Data Studio pureQuery Runtime	IBM Optim pureQuery Runtime for Linux, UNIX, and Windows
IBM Data Studio pureQuery Runtime for z/OS	IBM Optim pureQuery Runtime for z/OS
IBM Data Studio Administrator	IBM Optim Database Administrator
IBM DB2 Performance Expert	IBM Optim Performance Manager for DB2 on Linux, UNIX, and Windows
IBM DB2 Performance Expert Extended Insight Feature	IBM Optim Performance Manager Extended Insight for DB2 on Linux, UNIX, and Windows
IBM DB2 Optimization Expert	IBM Optim Query Tuner for DB2 on z/OS
	IBM Optim Query Tuner for DB2 on Linux, UNIX, and Windows
	IBM Optim Query Workload Tuner for DB2 on z/OS
	IBM Optim Query Workload Tuner for DB2 on Linux, UNIX, and Windows
IBM DB2 High Performance Unload	IBM Optim High Performance Unload for DB2 on Linux, UNIX, and Windows

DB2 Developer Workbench vs. Data Studio

Developer Workbench

- SQL Query Editor
 - SQLJ Editor
 - SQL Builder
 - XQuery Builder
 - SQL Routine Debugger
 - Java Routine Debugger
 - XML Editor
 - XML Schema Editor
 - Data Management
 - Visual Explain
 - Project Management

Data Studio

- Integrated Query Editor – SQL + XQuery
- SQLJ Editor
- SQL Builder
- XQuery Builder
- SQL Routine Debugger
- Java Routine Debugger
- XML Editor
- XML Schema Editor
- Data Management
- Visual Explain
- Project Management

***Data Studio is a full replacement of
DB2 Developer Workbench
plus much more***

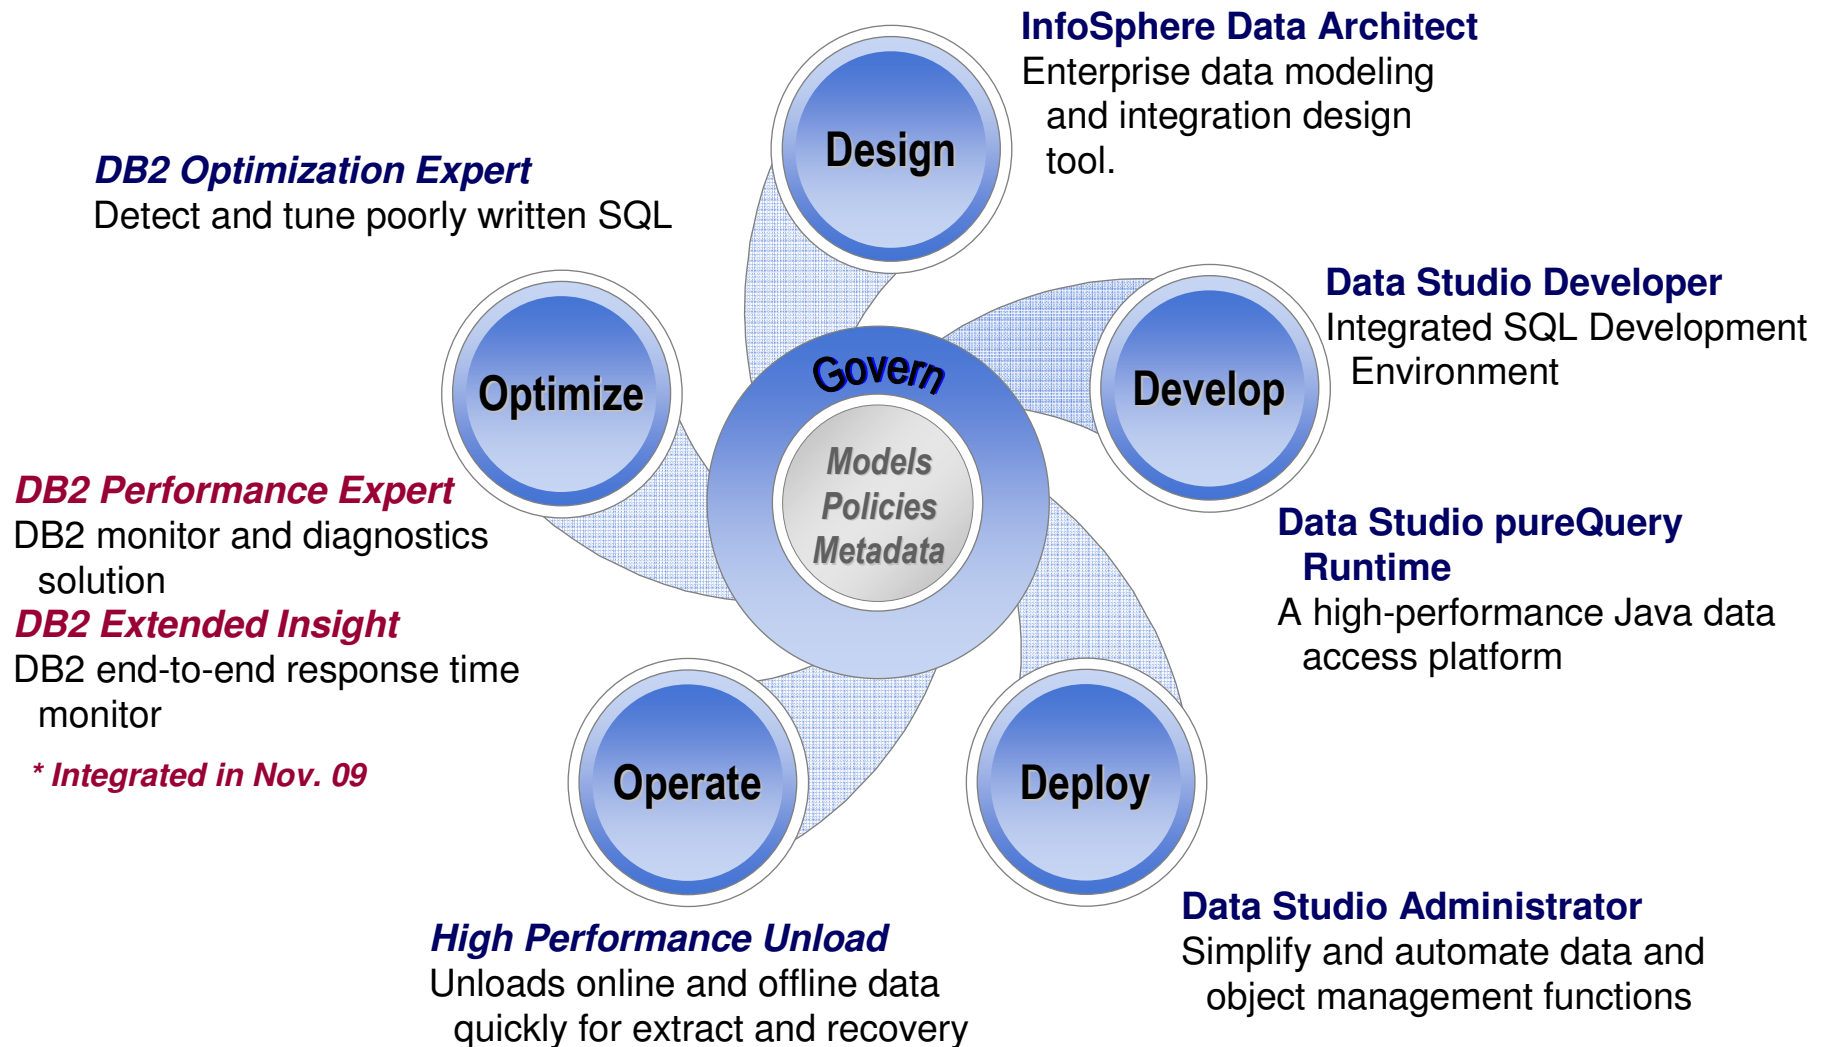
- ***DB2 for Linux, Unix, Windows v8.x, v9.1.x, v9.5***
- ***DB2 for z/OS v7, v8, v9***
- ***DB2 for i5/OS v5r2, v5r3, v5r4***
- ***Informix Dynamic Server (IDS) v9.x, v10.x, v11, v11.5***

■ ***Data Studio v2.1***

- ER Diagramming
- Data Distribution Viewer
- Object Management
- Privilege Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support
- Health Monitoring DB2 for LUW 9.5 and DB2 z/OS v9
- ***Instance and database commands***
- ***Utility support***
- ***Data management support***
- ***View and force applications***

Data Studio Offerings

Managing the value of your data throughout its lifetime



IBM Data Studio 關鍵價值及其整合元件



提高生產力

Java, XML, SOA 開發
作業自動化
團隊協同合作

提高服務品質

效能最佳化
增強的安全性
提升品質

提升效率


有效利用CPU性能
降低錯誤率
問題隔離

IBM InfoSphere Data Architect

A collaborative data design tool to understand information assets and their relationships, model data and integration designs, and enforce enterprise standards for data quality and consistency

IBM Optim Database Administrator

An administration environment to reduce application outages by automating and simplifying complex DB2 structural changes



IBM Optim Development Studio

An integrated development environment for rapidly creating and testing database and pureQuery applications and services.

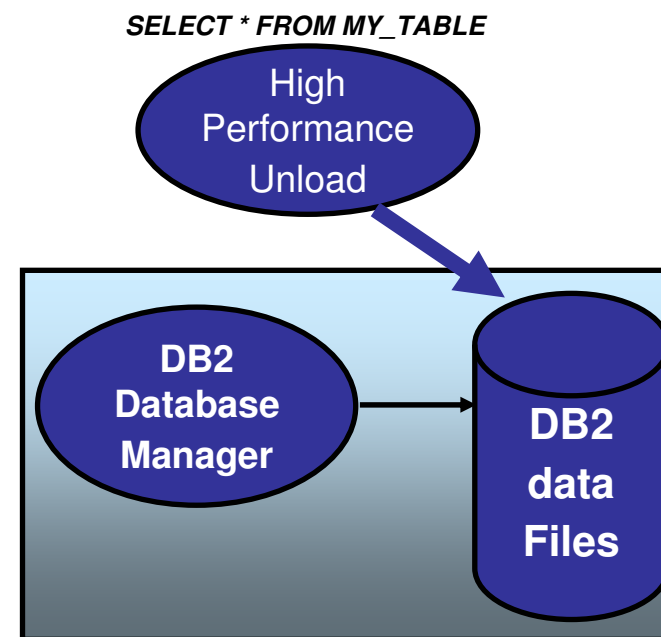
IBM Optim pureQuery Runtime

A high-performance Java data access platform to improve performance, security, and manageability of Java connections to databases.

DB2 High Performance Unload

DB2 High Performance Unload helps meet service agreements and eases application upgrades by extracting data quickly and efficiently.

- **Fast unload capability**
 - Often 4 to 6 times faster than export
 - Parallel processing for higher speeds
- **Flexibility to meet all use cases**
 - Full Partitioned DB2 support
 - Subset of SELECT syntax used to filter columns and/or rows
 - Many automatic data type conversions & formats available
 - Offline or Online Backups
- **Repartitioning feature**
 - Built-in “DB2 Splitter”
 - Unload and split in a single operation



■ **Data Studio v2.1**