



Integration Data Management Scenarios

Rafael A. Coss
Solution Architect for Optim
rcoss@us.ibm.com

Jay Bruce
DB2 Tools Architect/Strategist
jmbruce@us.ibm.com

Twitter: <http://twitter.com/racoss>
Blog: <http://datalifecycle.blogspot.com/>
LinkedIn <http://www.linkedin.com/pub/rafael-coss/9/311/268>
Facebook <http://www.facebook.com/racoss>

September 4, 2009

© 2009 IBM Corporation

Agenda

- **Eclipse and 3270 integration via Rational Developer for z/OS**

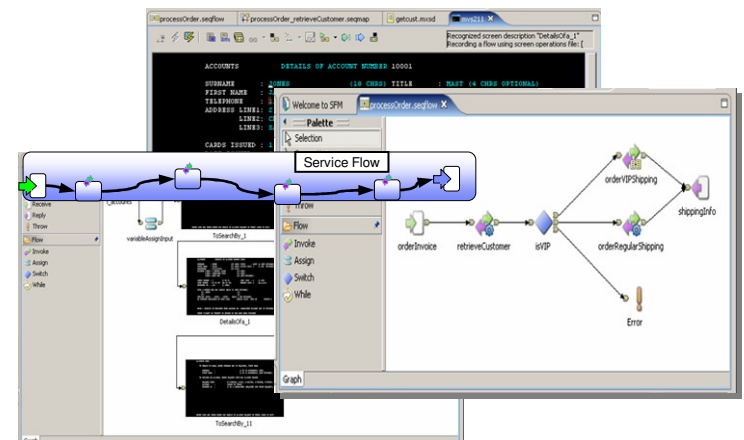
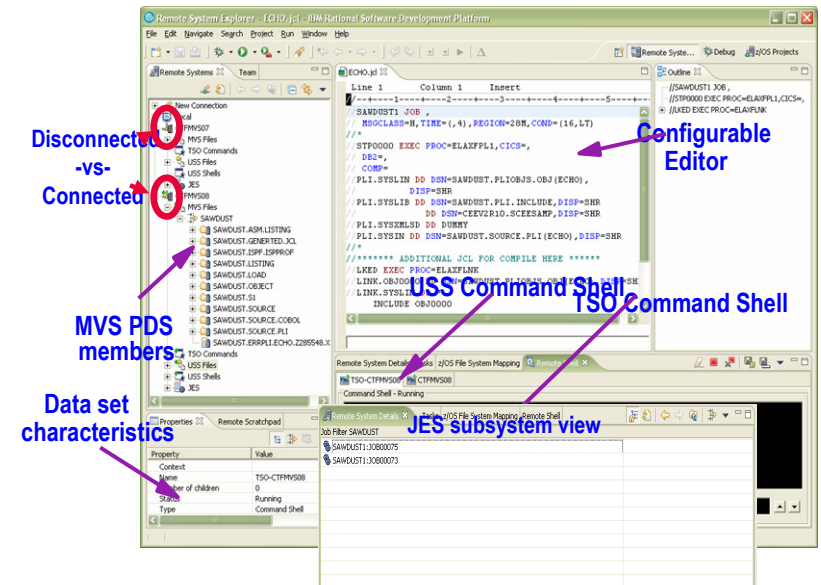
- **Integration Scenario**
 - Optim Development Studio/pureQuery

 - OMEGAMON

 - Optim Query Tuner

Rational Developer for System z (RDz) overview

- ▶ What is RDz
 - Eclipse-based IDE speeding modern composite (SOA) application development
- ▶ RDz supports Enterprise Modernization
 - Links WAS and core system z processing
 - Supports common IDE for COBOL, PL/I, C, C++, HLASM, Java, EGL and web services
 - Transforms UML to COBOL source code
 - Provides interactive access to z/OS for development, debug, job generation, submission, monitoring, command execution
 - Supports new and existing runtimes (CICS, IMS, Batch, USS, DB2 SP, WAS)
- ▶ RDz supports SOA
 - Enables CICS and IMS applications for web services and SOA
 - Supports for J2EE, JCA, XML, web services



IBM Rational Developer for System z V7.5 capability

JES and PD Tools

- Debug zOS applications from the workstation as they execute live in the remote runtime
- Read/Write/Update VSAM datasets via integration with IBM File Manager
- Access IBM Fault analyzer reports for analyzing ABENDS and associating back to source code
- Interact with the Job Entry Subsystem (JES) to submit jobs, monitor jobs, and review job output

Traditional Development

Development Environment

- Connect to z/OS systems
- Work with z/OS resources like COBOL, PL/I, C, C++, JCL, assembler, etc.
- Perform dataset management actions like allocating datasets and migrating datasets
- Perform typical edit, compile, and debug tasks on remote z/OS resources from the workstation
- Create, build, and catalog DB2 stored procedures on zOS
- Compile and test programs locally to ensure correctness

Screen design

- Visually create, modify, build, and deploy BMS maps sets or MFS/IMS maps remotely or on the local workstation

Code Generation

- Generate CRUD DB2 program code from UML, which can also be easily integrated into web service applications

IBM Rational Developer for System z

Host Tooling Integration

[JES, FA, FM, Debug Tool]

zOS Application Development

[COBOL, PL/I, C/C++, JCL, Screens, Stored Procedures, etc]

Enterprise Service Tools

[Web Services For CICS/IMS]

Mainframe / Runtime Integration

Eclipse Framework

zOS Web Service and Flow Creation

- Implements SOA and Web Services
- SOA access to CICS V3.2 and IMS V10 COBOL and PLI applications
- Bottom-up/Top-down/meet-in-the-middle COBOL to XML mapping support
- Integrated COBOL and PL/I XML converters, XML schemas, and WSDL generation
- Service Flow Modeler to build/deploy service flows out of your existing Commarea, Channel, MQ, and Terminal CICS applications.

Mainframe / System z Runtime Support



- Access to host SCMs such as SCLM
- Framework for writing/deploying custom SCM integration code
- SCM process integration (JCL, TSO commands, Green Screen applications) via HATS and Menu Manager
- CICS Explorer with Application Deployment Manager

Eclipse Platform and Java Development

- Plug-in integration framework
- Java Development (useful for System z Java development)
- Distributed team integration
- Database access/search tools

Access source code...


- RDz offers integration into a variety of Source Code Management (SCM) tools as well as a framework for creating SCM integration on your own
- Variety of vendors supply plug-ins to RDz to provide easy access to processes and source code controlled by their products



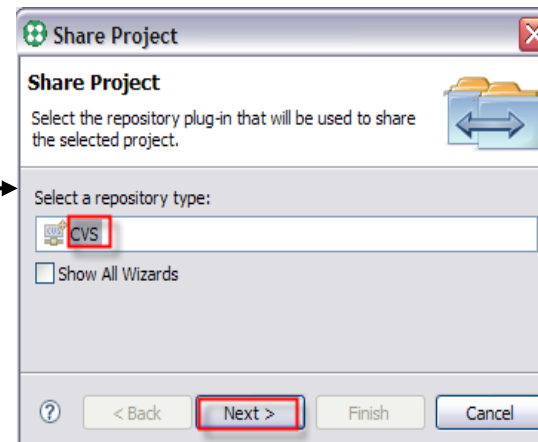
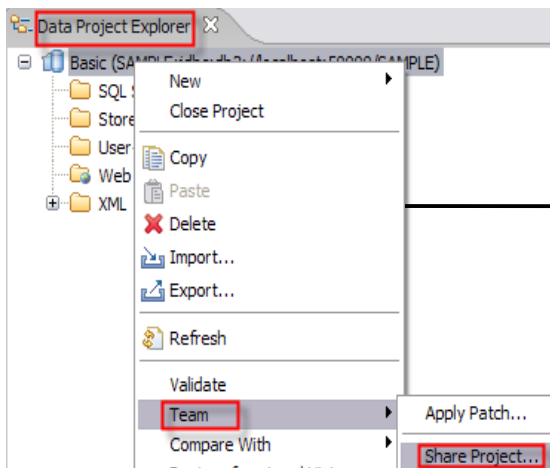
- IBM Supplied
- Vendor Supplied

Source Code Control *Project Management*

Project Explorer




Teaming options to share projects



RDz-based development

- Common development environment for COBOL, PL/I, C/C++, & Java
- Simplified development with more information at your fingertips



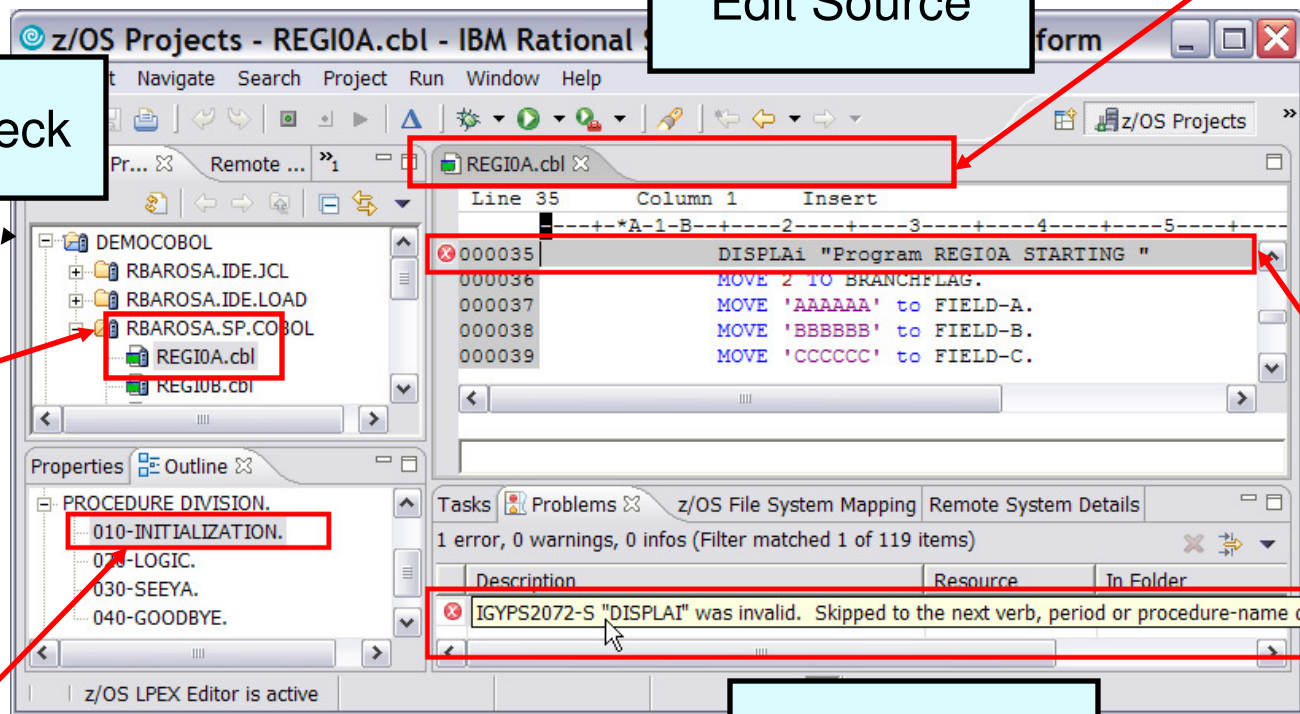
Open and edit multiple source and JCL members simultaneously

Edit Source

Syntax Check

Submit jobs, access job output, or open source members with a single click

Outline view presents COBOL structure



Statement in error indicated in source

Double-Click on the Error

Error list in Problems view

Navigate datasets or jobs live on zOS

The screenshot displays the IBM TSO/ISPF interface. The main terminal window shows a welcome message and a menu of available applications. The file browser on the right shows the directory structure for the remote system 'zserveros.demos.ibm.com', including 'MVS Files' and 'My Data Sets'.

Terminal Output:

```

DDDDDDDD EEEEEEEE MMMM MMMM 00000000
DD DD EE MM MM MM MM 00 00
DD DD EE MM M M MM 00 00 NNNN EEEE TT
DD DD EEEEEEE MM MMM MM 00 00 NN NN EE EE TTTT
DD DD EE MM M MM 00 00 NN NN EEEEE TT
DD DD EE MM MM 00 00 NN NN EE TT
DDDDDDDD EEEEEEEE MM MM 000000 NN NN EEEE TT

WELCOME TO IBM
ZSERVER05
HELP DESK: 972-561-7700

YOUR IP ADDRESS : 32.97.110.60
YOUR TELNET PORT: 30721

-----
APPLICATIONS AVAILABLE
-----
TSO      CICSA   CICSB   CICSC
IMSA     IMSB     IBMSM
NETVIEW  CICS    CICS    CICS

SELECTION ==>
M[A] a
    
```

File Browser Structure:

- Remote Systems > zserveros.demos.ibm.com
 - z/OS UNIX Files
 - z/OS UNIX Shells
 - MVS Files
 - My Data Sets (USER04.*)
 - USER04.COL
 - My Search Queries
 - TSO Commands
 - JES
 - My Jobs
 - DDS0753
 - z/OS UNIX Files
 - z/OS UNIX Shells
 - MVS Files
 - My Data Sets
 - DDS0753.ISPF.ISPPROF
 - ISPEDRT
 - ISPPROF
 - ISPSPROF
 - ISRLLIST
 - ISRPLIST
 - PROC1
 - DDS0753.SPFLOG1.LIST
 - My Search Queries
 - TSO Commands
 - TSO-DDS0753



The screenshot displays the IBM Data Studio environment. On the left, the 'Packager' and 'Data Source Explorer' panes are visible. The 'Packager' pane shows a project structure for 'myJavaProj' with various source files and libraries. The 'Data Source Explorer' pane shows a database connection to 'EOSDB205' with a tree view of schemas, including 'GOSALES'. The main editor window shows a Java file 'Default.bindProps' with configuration options for 'defaultOptions' and '-bindOptions'. A terminal window titled 'DDSO753.hce' is open, displaying the 'ISPF Primary Option Menu' with various utility options and system information. The terminal also shows a copyright notice for IBM and a list of function keys (PF1-PF12) for navigation.

Validate bind on the host

```
# on a line that begins with defaultOptions, as in this example:
#   defaultOptions= -bindOptions "qualifier yourSchema"

# If you are developing a query application, provide bind properties
# for specifying the schema name:
#   com.yourSchema
# These bind properties are specified in the defaultOptions line

# The content assistance in this editor can show you the options that are available.
# Press CTRL+Space or your default keystroke for this assistance as you are specifying

defaultOptions= -bindOptions "DDSO753"
```

Current host connection profile is: /HostConnectProjectFiles/DDSO753.hce

```
Menu Utilities Compilers Options Status Help
ISPF Primary Option Menu

P-JOHN-D001
0 Settings Terminal and user parameters User ID . . : DDS0753
  View Display source data or listings Time . . . : 20:19
  Edit Create or change source data Terminal. . : 3278
  Utilities Perform utility functions Screen . . : 1
  Foreground Interactive language processing Language. . : ENGLISH
  Batch Submit job for language processing Appl ID . . : ISPF
  Command Enter TSO or Workstation commands TSO logon : DBAUSER
  Dialog Test Perform dialog testing TSO prefix: DDS0753
  IBM Products IBM program development products System ID : ESYSMVS
  MVS acct. : 12345678
  Release . : ISPF 6.0

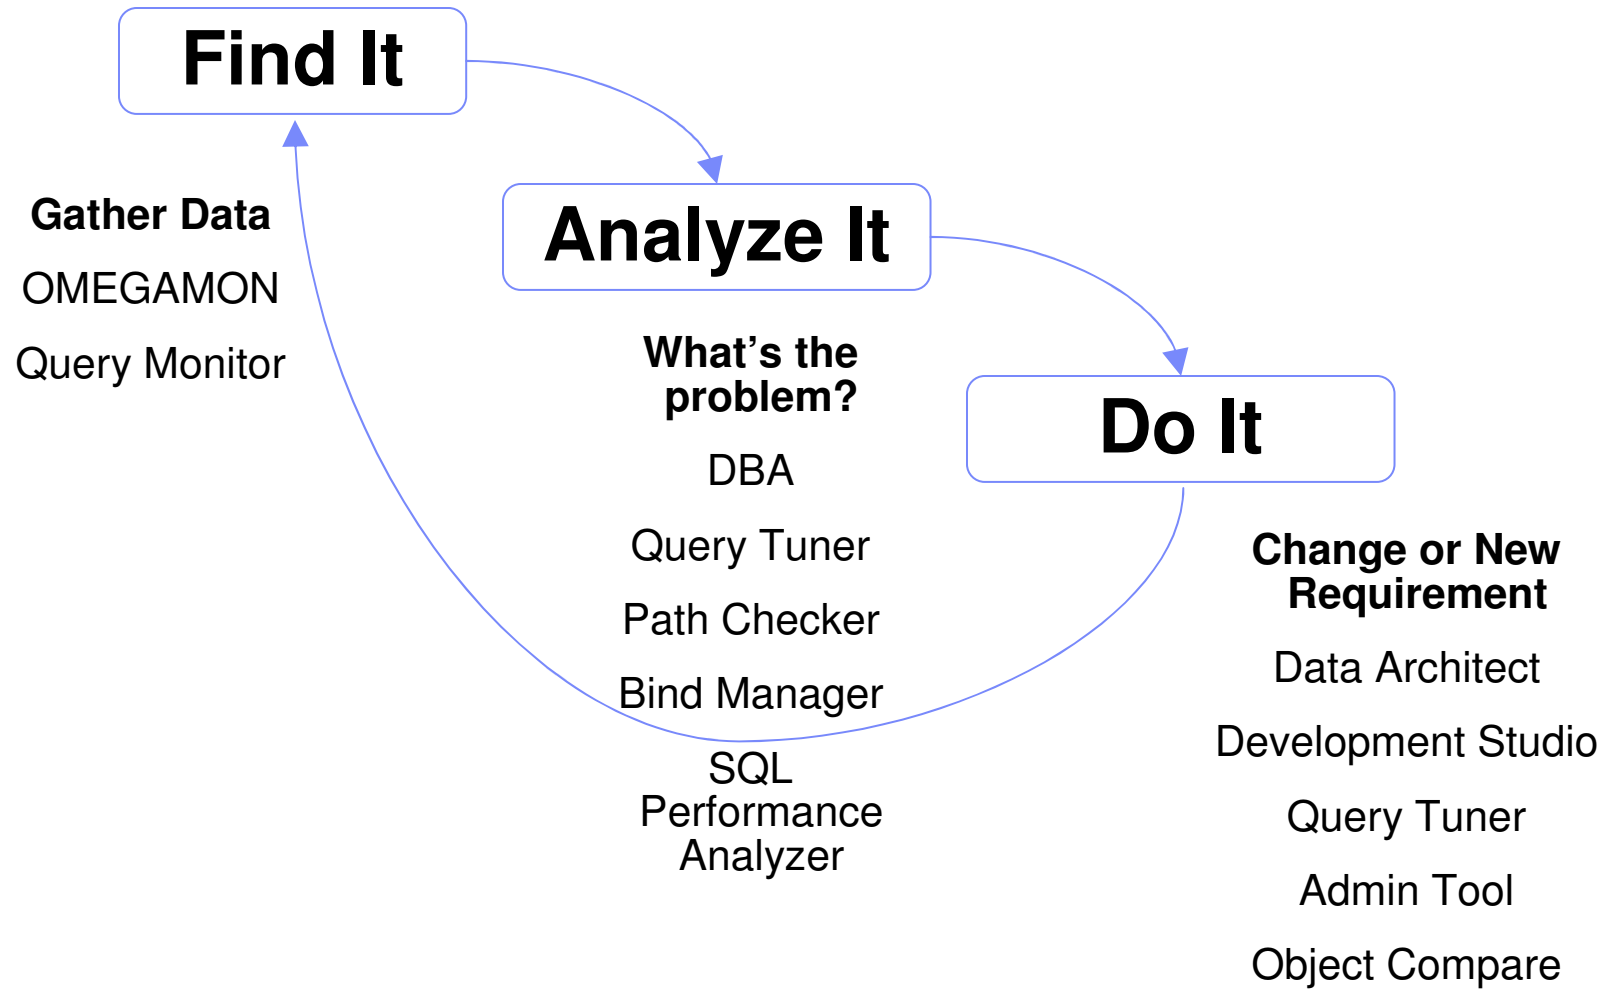
Licensed Materials - Property of IBM
5694-A01 Copyright IBM Corp. 1980, 2008.
All rights reserved.
US Government Users Restricted Rights -
Use, duplication or disclosure restricted
by GSA ADP Schedule Contract with IBM Corp.

Option ==>
F1=Help F2=Split F3=Exit F7=Backward F8=Forward F9=Swap
F10=Actions F12=Cancel
MA a 22/014
```

Host Properties Host Connection

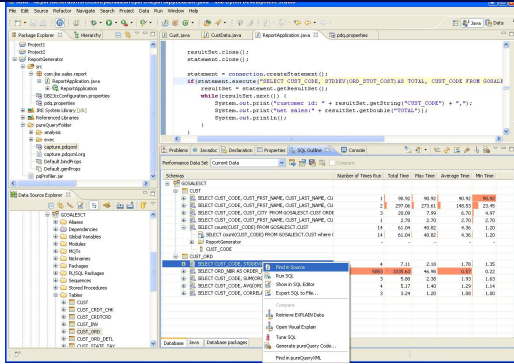
Performance Management

▪ A Closed Loop Approach

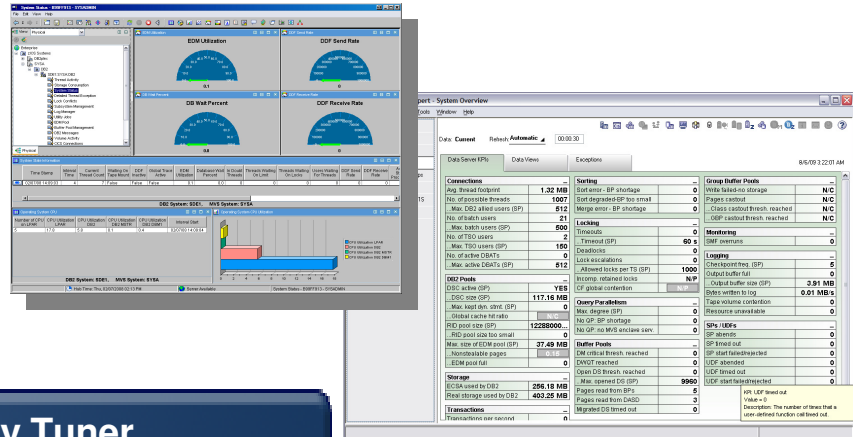


Prevent Problems, Improve Performance, Reduce costs

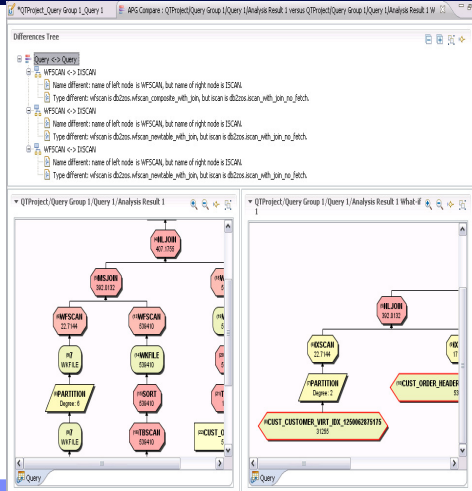
Optim Development Studio pureQuery



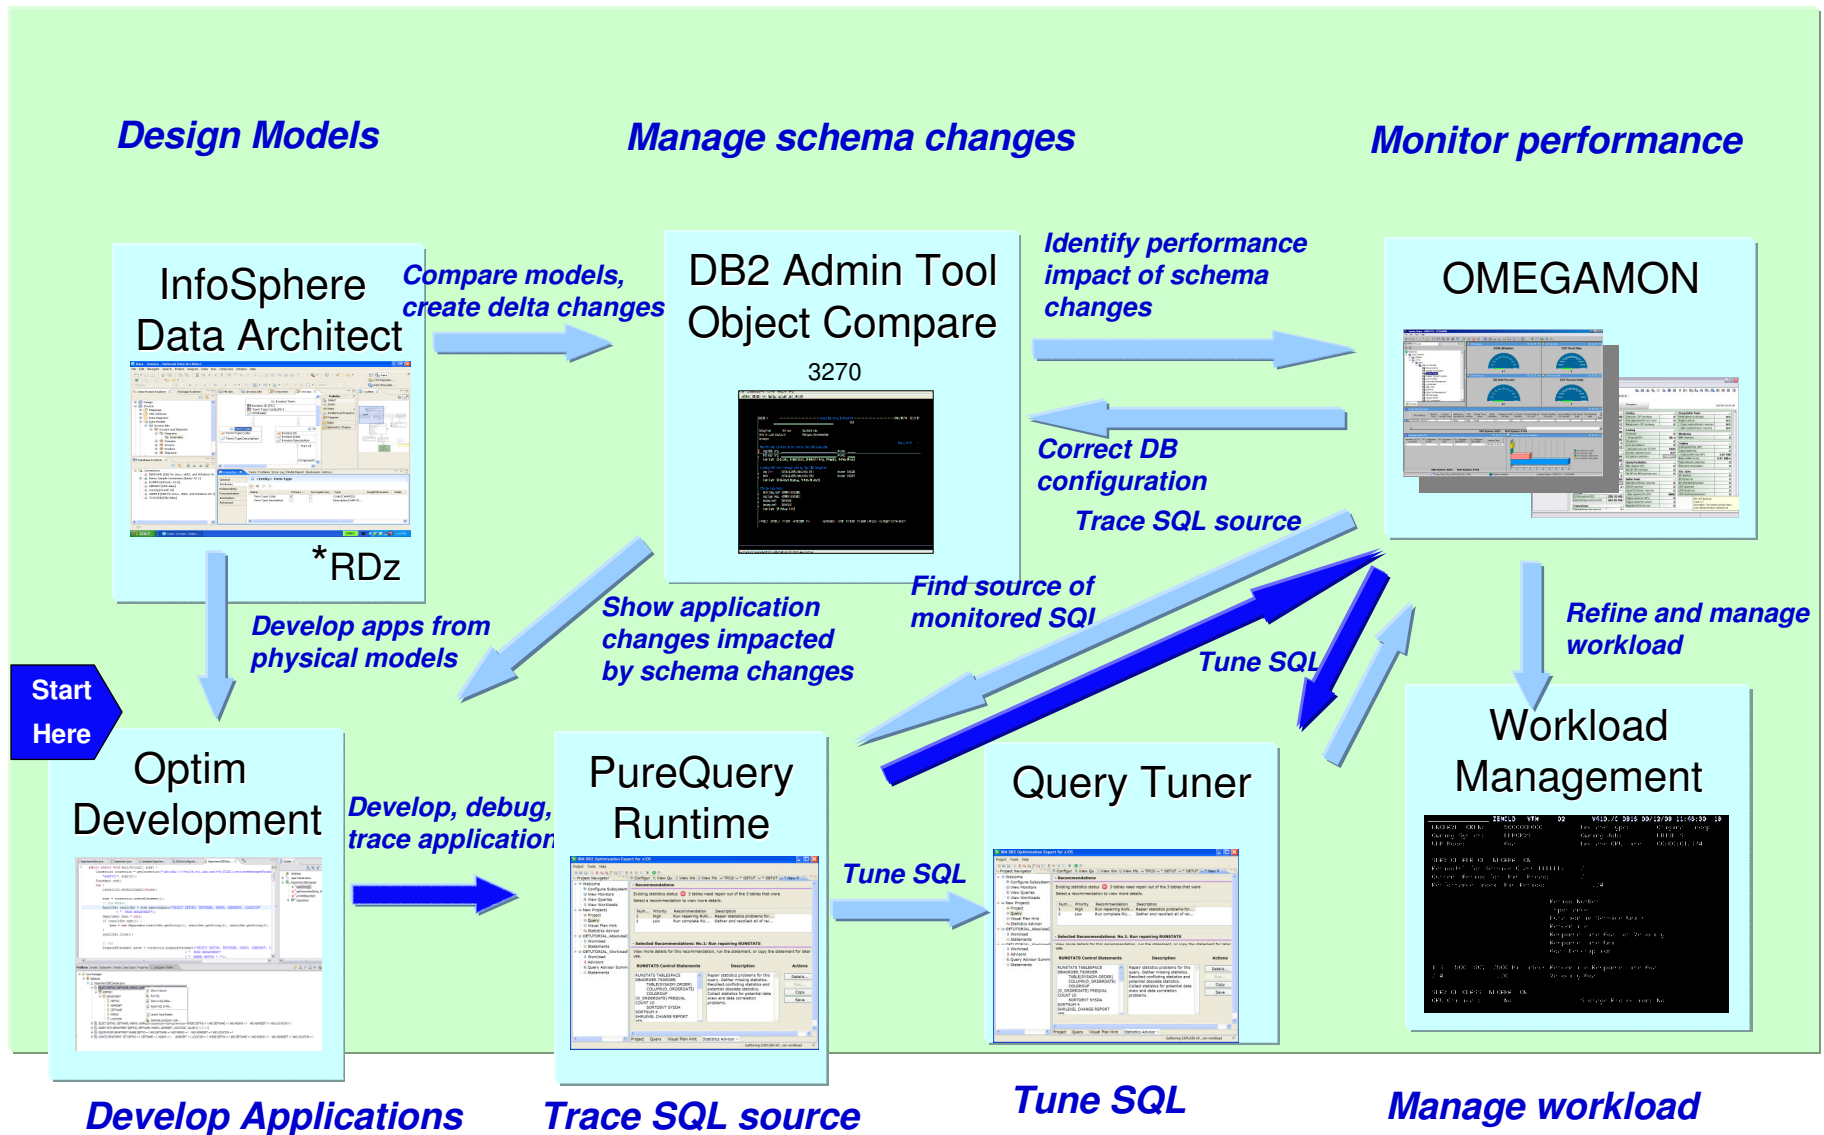
OMEGAMON



Optim Query Tuner



DB2 for z/OS IDM Scenario



COBOL & CICS Rocks!



“Our top story: Large Customer moves from COBOL to Java to become more agile. In other news, 15 DBA develop amnesia.”

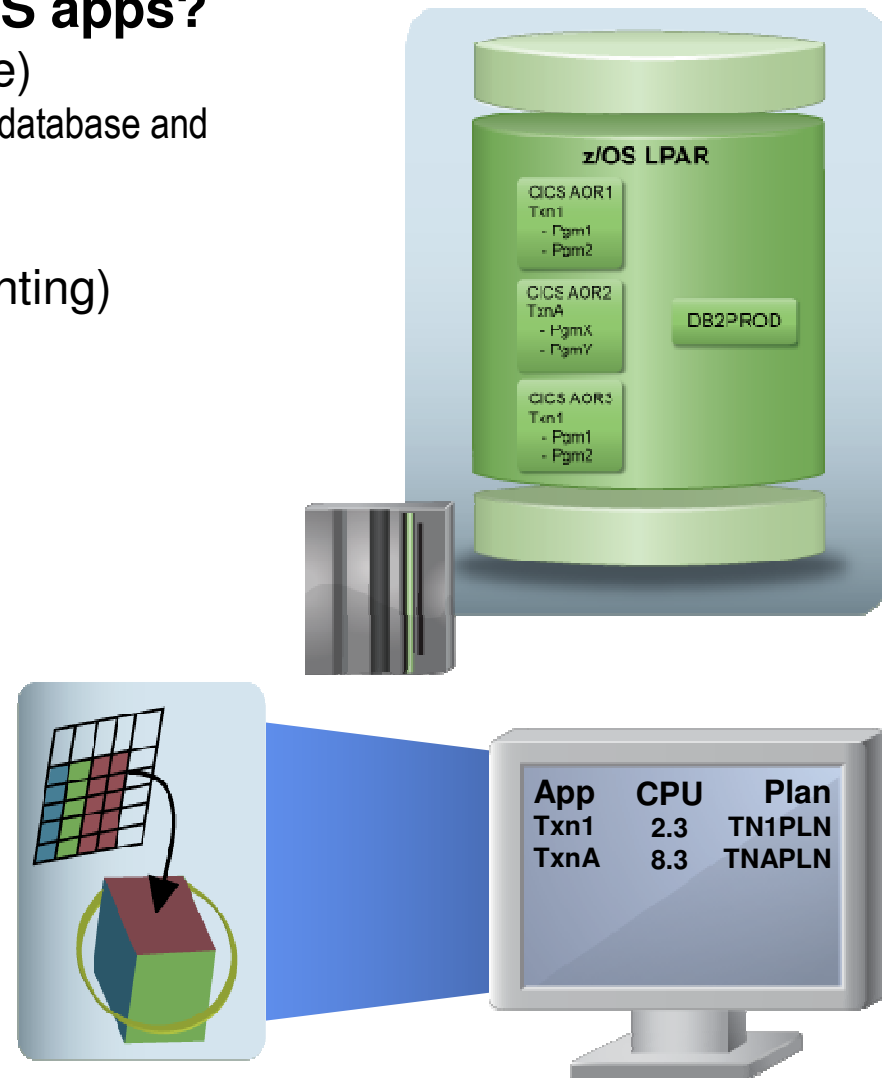
What's So Great About DB2 Accounting for CICS?

■ What's great about 30 yr old CICS apps?

- See end user ID (CICs thread reuse)
 - Take a user identify and push it down to the database and compare users to each other
- See address space (AOR name)
- See program (package level accounting)
 - See package data between programs or how a package has changes over time
- See TX (PLAN name)
 - Measure correlated performance

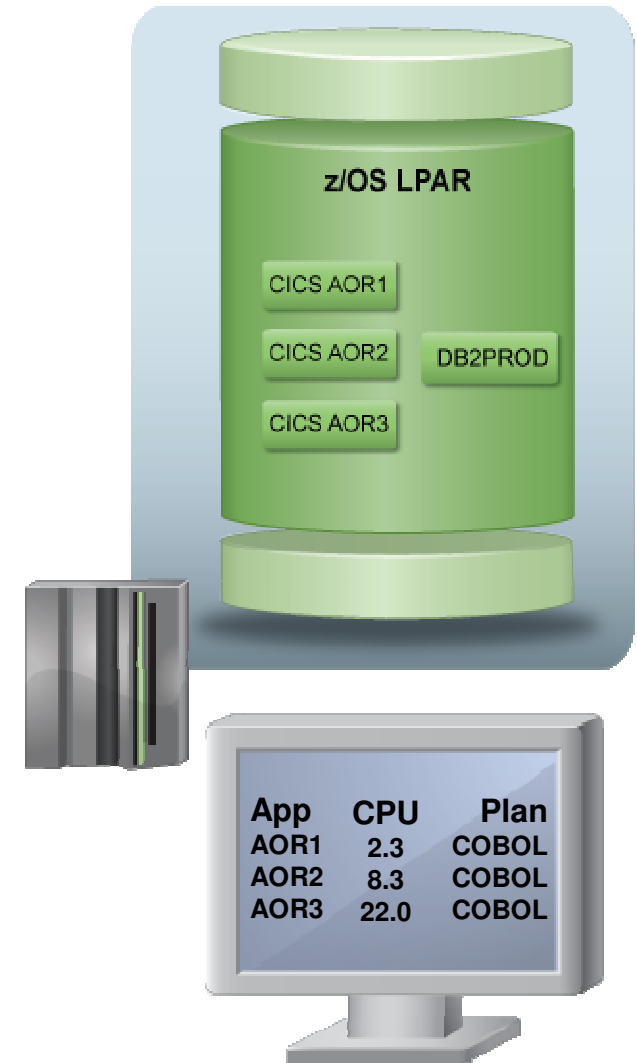
■ CICS give you performance monitoring and tuning which is really a giant Pivot table

- Rotate and analyze the program
- How much of my system does this application take?
- What called this poorly running query
- +++



What if We Treated CICS Like DB2?

- **What if we took away the key monitoring tokens you use for CICS and you only knew 3 things about the CICS workload?**
 - Only one plan name: COBOL
 - CICS AOR started task name
 - No end user ID - just user ID of the AOR that started task
- **How would you know?**
 - Which app is running?
 - Which user is running?
 - Which developer wrote the app?
 - What other SQL does this app issue?
 - When was the app last changed?
 - How has CPU changed over time?
 - +++



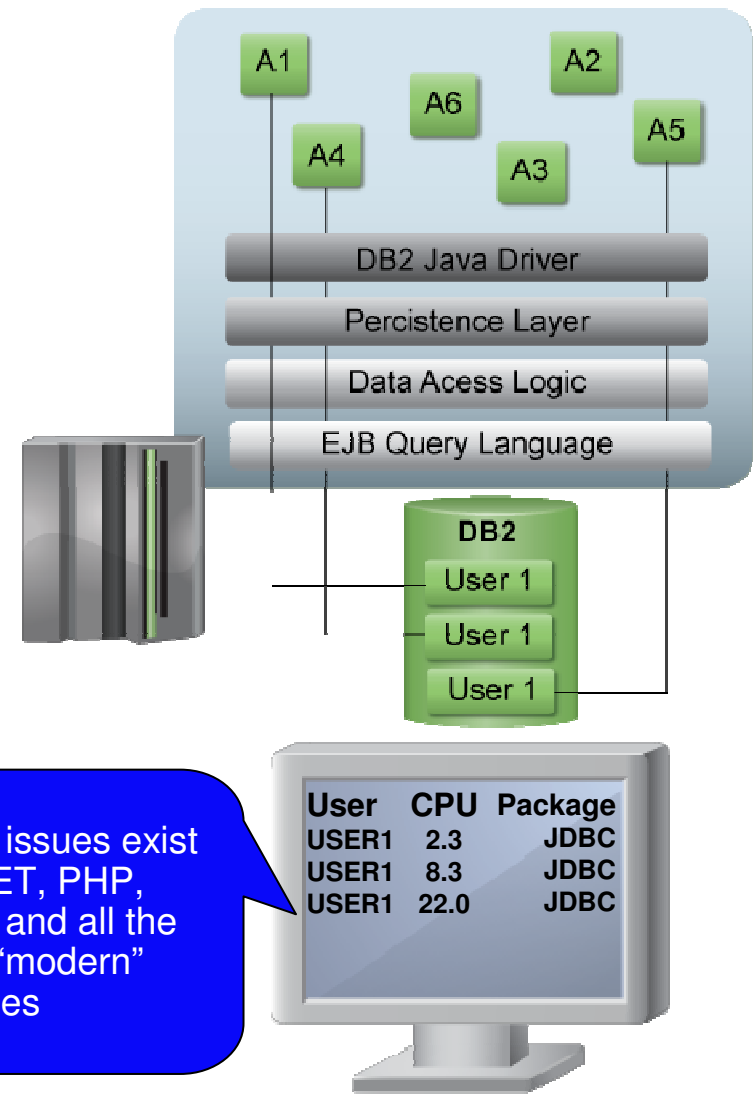
You'd Have Life Before Optim pureQuery

- **What is visible to the DBA?**
 - IP address of WAS application server
 - Connection pooling user ID for WAS
 - Application is running JDBC

- **What is not known by the DBA?**
 - Which application is running
 - What application developer wrote it?
 - What other SQL does this application issue?
 - When was the application last changed?
 - How has CPU changed over time?

- **We seemed to have forgotten a lot in this modern era:**

Enterprise Amnesia



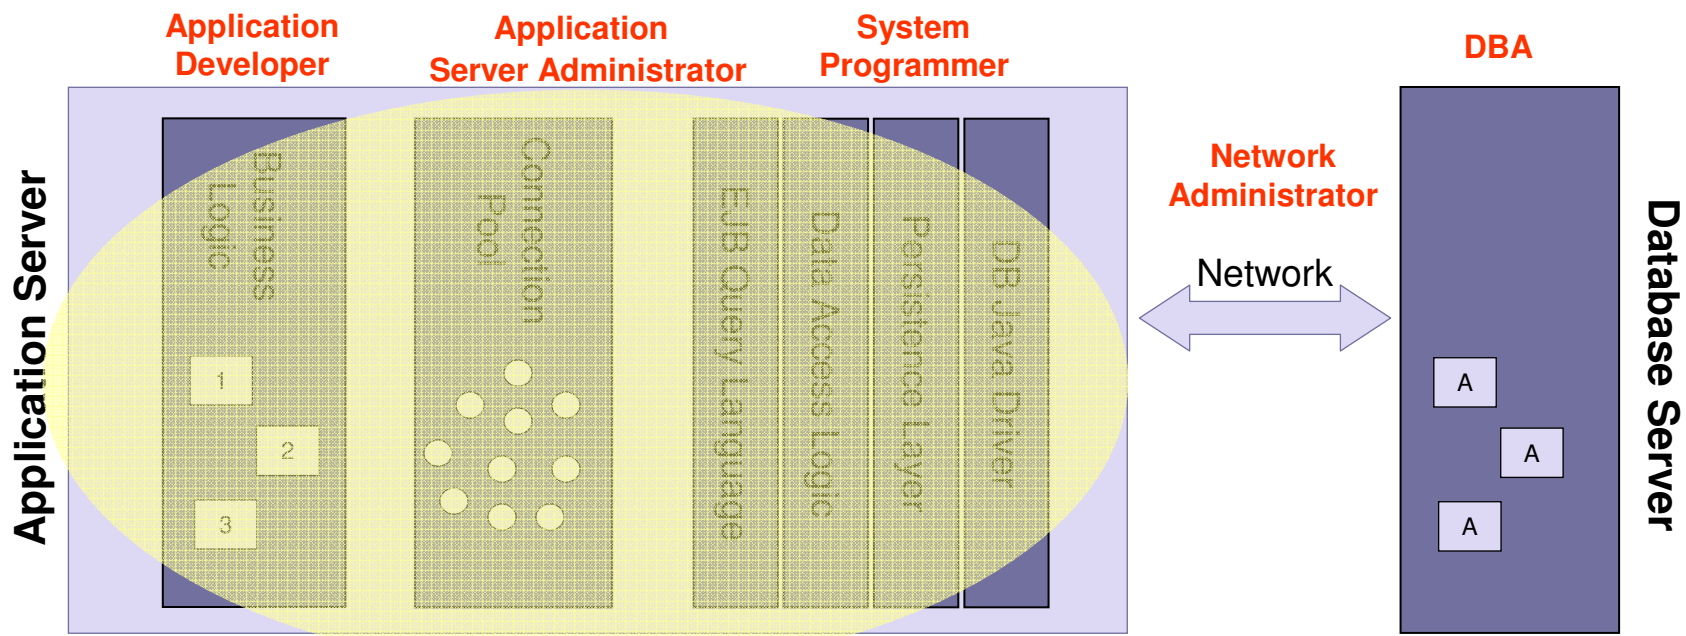
Enterprise Amnesia is a Well Known Problem

- This is a time-consuming and staff intensive processes
 - No holistic view
 - No resource consumption profile
 - No correlation of SQL to source application
 - Problem isolation requires cross-role collaboration

Microsoft SQL Server 2008 review

PC ADVISOR


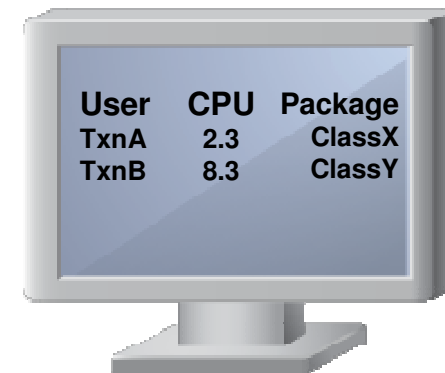
The recent expensive queries drill-down is almost completely pointless. Does it give you the most recent expensive queries? Yes, it does. Does it give you important stats on these queries such as CPU, number of reads and writes, average duration, and plan count? Absolutely. Does it tie these stats to a user so that you can tell who or what is performing these offending actions? No. What is a DBA supposed to do with that?



What's so Great About pureQuery Accounting?


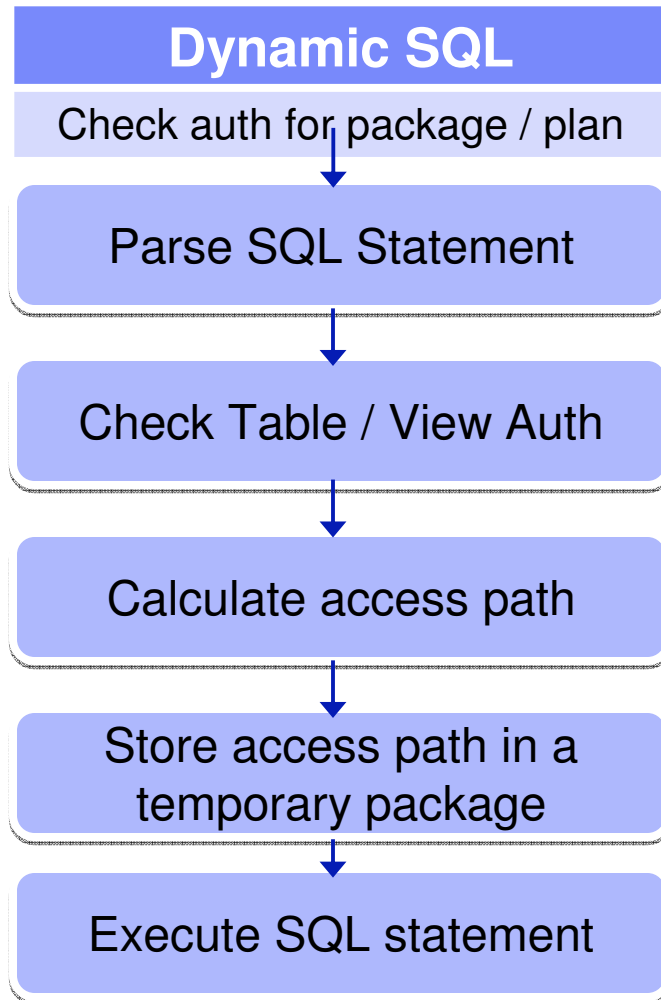
- **Data Studio and pureQuery provide the same granularity for reporting WebSphere's DB2 resources that we have with CICS:**
 - By transaction
 - Set Client Application name
 - By class name
 - Program/Package level accounting
 - By address space
 - IP address
 - By end user ID
 - DB2 trusted context and DB2 role

- **This flexibility makes it very easy to:**
 - Isolate performance problems
 - Perform capacity planning exercises
 - Analyze program changes for performance regression
 - Compare user's resource usage
 - +++

User	CPU	Package
TxnA	2.3	ClassX
TxnB	8.3	ClassY


SQL Execution - Dynamic vs. Static




pureQuery enables wide use of Static SQL



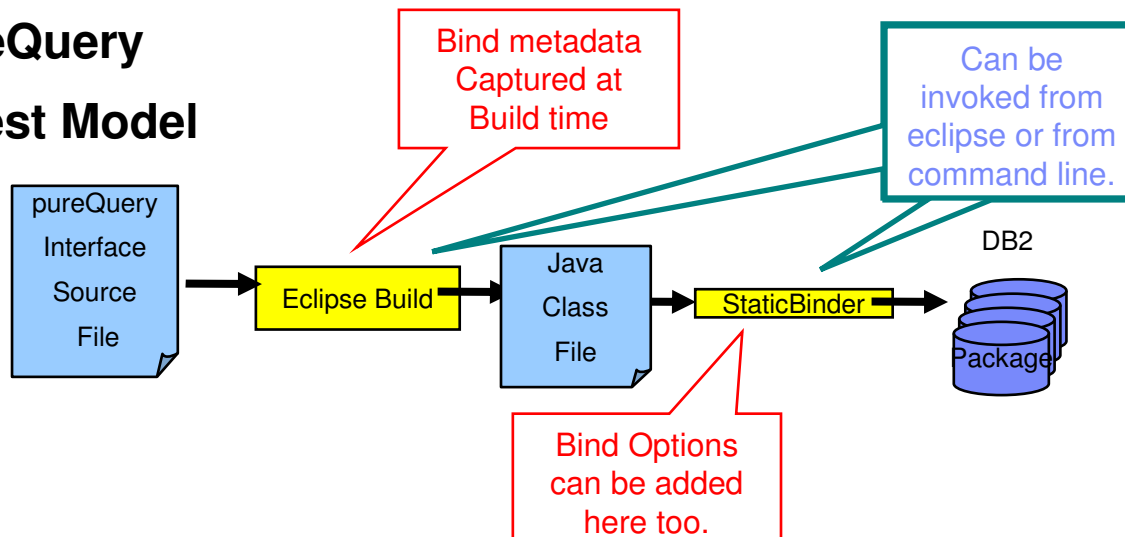
- Static SQL
 - Highest speed
 - Greatest reliability
- JDBC is basic access, uses Dynamic SQL
- SQLJ adds Static SQL
- pureQuery supports both Static SQL and Dynamic SQL
 - Code to dynamic SQL, turn on static SQL at deployment



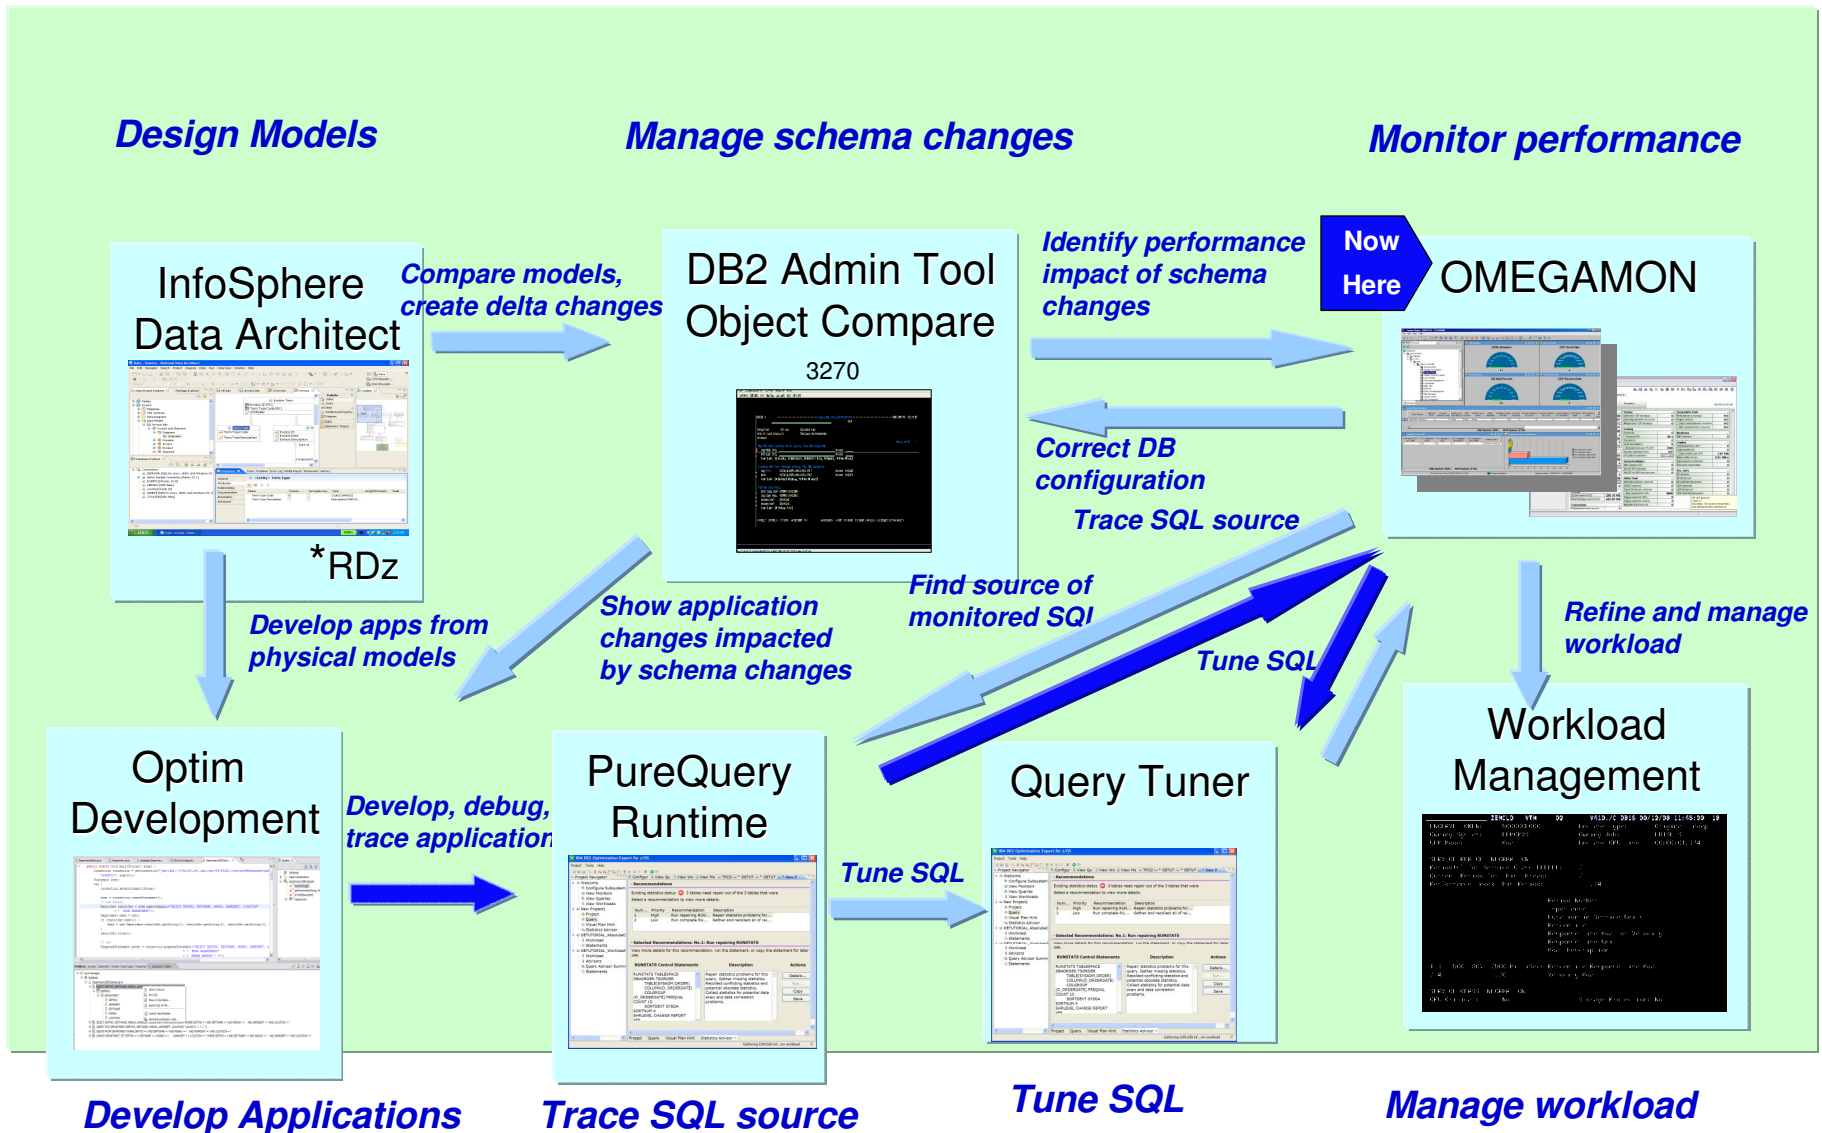
Static SQL Deployment Models - Comparison



pureQuery Simplest Model



DB2 for z/OS IDM Scenario



Unique Package Names Improves Problem Determination

- **Hard to identify unique Java applications**
 - Most dynamic Java applications use packages SYSLNx00
- **Unique package names link SQL to Java beans**
 - Similar to CICS transaction names to programs

Static
pureQuery
Java SQL

Dynamic
Java SQL

```

ZALLU   VTM   02   V410./C DB1S 09/12/08 11:29:22  2
> Help PF1   Back PF3   Up PF7   Down PF8   Sort PF10   Zoom PF11
> T.A       OMEGAVIEW PA2
>          THREAD ACTIVITY:  Enter a selection letter on the top line.
> *-ALL     B-TSO     C-CICS   D-IMS     E-BACKGROUND  F-DIST ALLIED
> G-DIST DBAC  H-UTIL   I-INACT  J-FILTER  K-FUNCTIONS  L-STORED PROC
> M-TRIGGERS N-SYSPLEX O-ENCLAVES P-WORKSTA
=====
>          ALL THREADS CONNECTED TO DB2
PTHDA                                         FLTR ON
+
+ Elapsed      * Package      CPU      Status      GetPg  Update  Commit  CORRID
+ -----      -
+ 00:00:13.6   PAW_OR_0    00.0%   IN-DB2      25     0       0   db2jcc_appli
+ 00:02:27.3   SYSLN200    00.0%   IN-DB2      897     0       0   db2jcc_appli
+ 00:02:52.3   SYSLN200    00.0%   IN-DB2     1025     0       0   db2jcc_appli
+ 00:03:05.8   SYSLN200    00.0%   IN-DB2     1324     0       0   db2jcc_appli
+ 00:02:32.7   SYSLN200    00.0%   IN-DB2     961     0       0   db2jcc_appli
+ 00:02:59.2   SYSLN200    00.0%   IN-DB2     1046     0       0   db2jcc_appli
=====

```

Problem Determination

Correlate Package and SQL with Captured Metadata

The screenshot shows the IBM Data Studio Developer interface. The main console window displays the following text:

```

-----
              ZSQL      VTM      02      V410./C DB1S 08/11/08 15:48:30  5
-----
> A-THREAD DETAIL B-LOCK COUNTS C-LOCK WAITS      D-LOCKS OWNED  E-GLOBAL LOCKS
> *-CURRENT SQL  G-SQL COUNTS  H-DISTRIBUTED  I-BUFFER POOL  J-GROUP BP
> K-PACKAGES      L-RES LIMIT   M-PARALLEL TASKS N-UTILITY      O-OBJECTS
> P-CANCEL THREAD Q-DB2 CONSOLE  R-DSN ACTIVITY  S-APPL TRACE   T-ENCLAVE
> U-LONG NAMES    V-SQL PA

=====
>
              SQL CALL BEING EXECUTED
PLAN
+ Thread: Plan=DISTSERV Connid=SERVER Corrid=db2jcc_appli Authid=DBA031
+ Dist : Type=DATABASE ACCESS, Luwid=G9274097.H546.C2D458DA7274=15806
+
  SQL call is active, call information is as follows :
+
+ Thread Status = IN-DB2           SQL Request Type = STATIC
+ Total SQL Reas = 93             SQL Call Type = FETCH
+ SQL DBRM Name = Order_de        SQL Statement Number = 00001
+ Collection ID = NULLID

+
+ DECLARE DB2JCCCURSOR1 CURSOR FOR SELECT ORDER_DETAIL_CODE, ORDER_NUMBER,
+ SHIP_DATE, PRODUCT_NUMBER, PROMOTION_CODE, QUANTITY, UNIT_COST, UNIT_P
+ RICE, UNIT_SALE_PRICE FROM GOSALES.ORDER_DETAILS FOR READ ONLY
  
```

The Package Explorer on the left shows a project named 'ClientOptimization' with a folder 'pureQueryFolder' containing a file 'capture.pdqxml'. A blue arrow points from this file to the 'Thread Status' information in the console.

The Database Explorer at the bottom shows a connection to a DB2 database. The 'pureQuery Outline' window displays the SQL statement: 'SELECT ORDER_DETAIL_CODE, ORDER_NUMBER, SHIP_DATE, PRODUCT_NUMBER, PROMOTION_CODE, QUANTITY, UNIT_COST, L'. A blue arrow points from this SQL statement to the 'SQL Statement Number' in the console.

Workload Manger (WLM) Classification

- **DB2 T4 driver provides properties for the connection class to help WLM classify work BUT they are not often used**
 - Example: setClientUser, setClientApplicationInformation, setClientWorkStation and setClientAccountingInformation

- **pureQuery's unique package names can provide WLM classification without programmers needing to including properties**

```

ZENCLD  VTM  02  V410./C DB1S 09/12/08 11:46:39 55
+
+ CLASSIFICATION WORK QUALIFIERS
+ Subsystem Type:   DDF                Correlation:      DB2JCC_APPLI
+ Proc Name:       Trans Program Name:
+ UserId:          DBA031              Transaction Class:
+ Network ID:      Logical Unit Name:
+ Plan Name:       DISTSERV            Package Name:      PAW_OR_0
+ Connection:      SERVER               Collection:        NULLID
+ Function Name:   DB2_DRDA             Subsystem Name:   DB1S
+ Accounting Info: JCC03520IBM-3D70     Subsystem Parm:   DBA031
+ Perform:        Subsystem Priority:   N/A
+ Scheduling Env: Subsys Coll Name:    DSNSG
+ Process Name:    DB2JCC_APPLICATION
+
+ Performance Index Input Data for Velocity Goal
+ Total Usings:    3930                 Total Delays:     712
=====

```

Without Proper WLM Classification: Default DDF Service Class Used

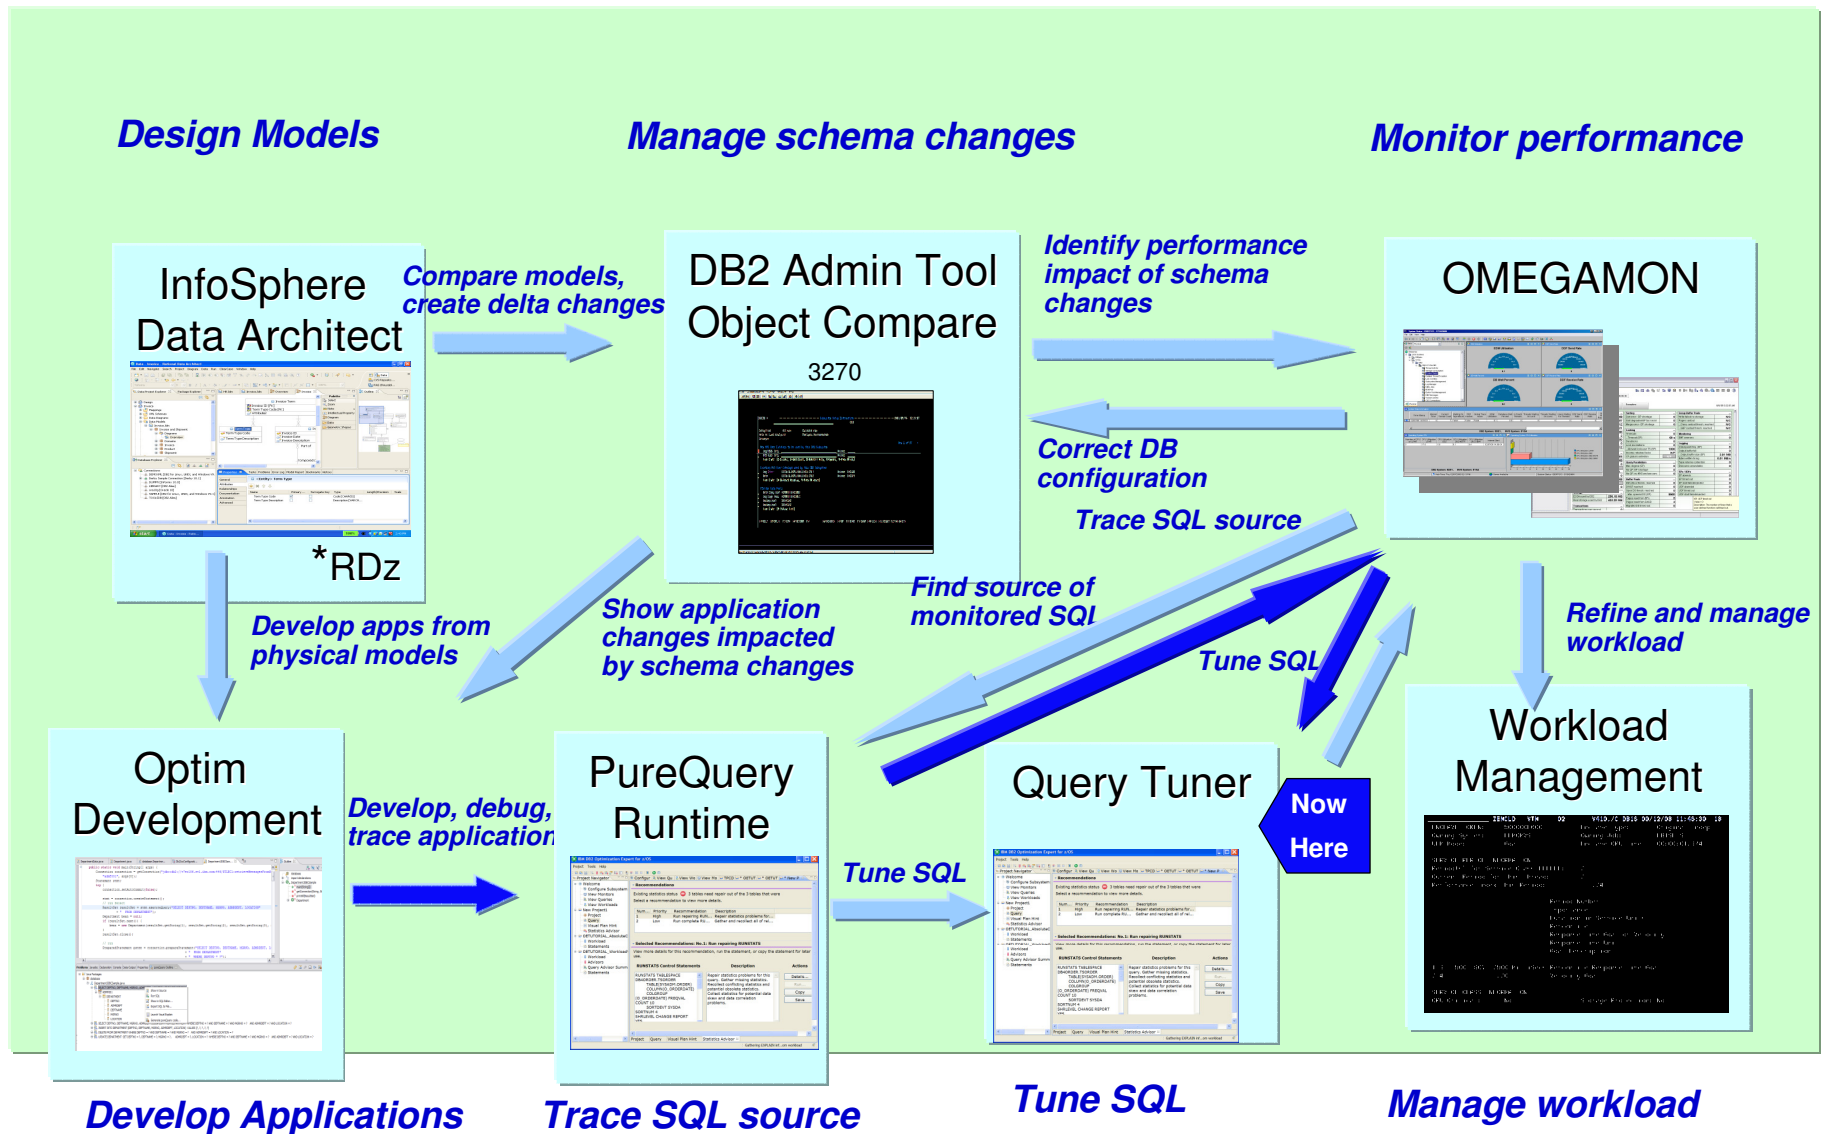
```

ZENCLD  VTM    02      V410./C DB1S 09/12/08 11:46:39 18
+ ENCLAVE TOKEN:      500000E600      Enclave Type:      Original Indep
+ Owing System:      DEMOMVS          Owing Job:         DB1SDIST
+ WLM Mode:          Goal              Enclave CPU Time   00:00:01.374
+
+ SERVICE PERIOD INFORMATION
+ Period(s) for Service Class DDFDEF:    2
+ Current Period for This Thread:        2
+ Performance Index This Period:         .24
+
+ +----- Period Number
+ | +----- Importance
+ | +----- Duration in Service Units
+ | +----- Percentile
+ | +----- Response Time Goal or Velocity
+ | +----- Response Time Unit
+ | +----- Goal Description
+ |-----|-----|-----|-----|
+ 1 3 500 80% 2500 MilliSecs Percentile Response Time Goal
+ 2 4 .20 Velocity Goal
+
+ SERVICE CLASS INFORMATION
+ CPU Critical:      No                  Storage Protection: No
+

```

Ouch! My Importance is only 4 and I'm using the DDFDEF service class!

DB2 for z/OS IDM Scenario



Top Pains for the DBA



Maintenance

- Space Management
- Efficient Resource Management
- Proactive Performance

Pains

- Time Consuming & Tedious
- Hidden Inefficiencies
- Too late; Business Affected



Performance & Tuning

- React to Production Problems
- Diagnose Response Time Issues
- Tune SQL & System
- Determine Best Practices

Pains

- Interruptions
- Difficult to Diagnose
- Broad Expertise Required
- Wasted Time

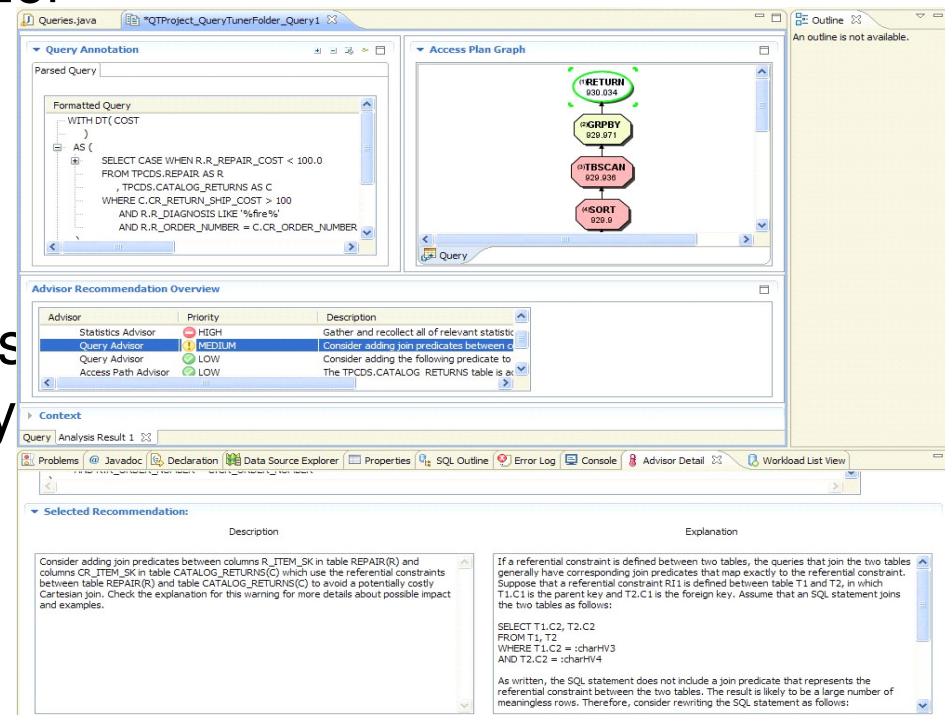
Get Expert Advice for Tuning DB2 Queries

■ Improve query performance

- Improve query structure
- Improve statistics quality
- Create indexes to reduce scans
- Mitigate problematic optimizer choices

■ Reduce costs

- Simplify analysis for DBAs
- Eliminate irrelevant statistics
- Improve database efficiency



The screenshot displays the IBM Optim Query Tuner interface. The top-left pane shows the 'Parsed Query' and 'Formatted Query' for a SQL query. The top-right pane shows the 'Access Plan Graph' with nodes for RETURN, GRPBY, TBSKAN, and SORT. The middle pane shows the 'Advisor Recommendation Overview' table.


Advisor	Priority	Description
Statistics Advisor	HIGH	Get and recollect all of relevant statistics
Query Advisor	MEDIUM	Consider adding join predicates between columns R_ITEM_SK in table REPAIR(R) and columns CR_ITEM_SK in table CATALOG_RETURNS(C) which use the referential constraints between table REPAIR(R) and table CATALOG_RETURNS(C) to avoid a potentially costly Cartesian join. Check the explanation for this warning for more details about possible impact and examples.
Query Advisor	LOW	Consider adding the following predicate to the WHERE clause of the query: AND R.R_ORDER_NUMBER = C.CR_ORDER_NUMBER
Access Path Advisor	LOW	The TPCDS.CATALOG_RETURNS table is accessed using a full table scan.

The bottom pane shows the 'Selected Recommendation' with a detailed description and explanation, including a sample SQL query:

```
SELECT T1.C2, T2.C2
FROM T1, T2
WHERE T1.C2 = :char#V3
AND T2.C2 = :char#V4
```

The explanation states: "If a referential constraint is defined between two tables, the queries that join the two tables generally have corresponding join predicates that map exactly to the referential constraint. Suppose that a referential constraint R11 is defined between table T1 and T2, in which T1.C1 is the parent key and T2.C1 is the foreign key. Assume that an SQL statement joins the two tables as follows: ... As written, the SQL statement does not include a join predicate that represents the referential constraint between the two tables. The result is likely to be a large number of meaningless rows. Therefore, consider rewriting the SQL statement as follows: ..."

Identify High Cost Queries and Workloads



Execute advisors

The screenshot displays the IBM Optim Query Tuner Client interface with several key panels highlighted:

- Statistics Quality & Collection:** Shows a `RUNSTATS` control statement:


```
RUNSTATS ON TABLE 'INVEK' 'EMBL OVFP'
WITH DISTRIBUTION ON COLUMNS ('LASTNAME' NUM_FREQVALUES 15
NUM_QUANTILES 25)
ALLOW WRITE ACCESS
```
- Query Design:** Displays a **Workload Query Advisor Recommendations Summary** table:

Statements Sorted by	Number
Statements Analyzed Successfully	22
Statements with Warnings	4
Number of High Severity Warnings	0
Number of Medium Severity Warnings	0
Number of Low Severity Warnings	7
Statements with High Severity Warnings	0
Statements with Medium Severity Warnings	0
Statements with Low Severity Warnings	4
- Query efficiency:** Shows performance metrics:
 - Estimated performance improvement: 97.8 %
 - Disk space required (DASD space): 37.92 MB
- DDL Details:** A pop-up window showing SQL commands to create indexes:


```
CREATE INDEX 'DBCODE' 'CUST_ORDER_HEADER_VIRT_IDX_1250594768821' ON
'03SALESCT' 'CUST_ORDER_HEADER' ('CUST_CODE' ASC, 'CUST_ORDER_NUMBER' ASC,
'CUST_ORDER_DATE' ASC) NOT PARDED FREEPAGE 0 PCTFREE 10;

CREATE INDEX 'DBCODE' 'CUST_CUSTOMER_VIRT_IDX_1250594768826' ON
'03SALESCT' 'CUST_CUSTOMER' ('CUST_CITY' ASC, 'CUST_CODE' ASC,
'CUST_LAST_NAME' ASC) NOT PARDED FREEPAGE 0 PCTFREE 10;

CREATE INDEX 'DBCODE' 'CUST_ORDER_DETAIL_VIRT_IDX_1250594768829' ON
'03SALESCT' 'CUST_ORDER_DETAIL' ('CUST_ORDER_NUMBER' ASC, 'PRODUCT_NUMBER'
ASC, 'CUST_QUANTITY' ASC) NOT PARDED FREEPAGE 0 PCTFREE 10;
```

Visualize Queries and Costs

- **View formatted queries with statistical annotations**
 - Automate time-consuming tasks

- **Get recommendations on:**
 - Writing queries
 - Influencing access path selection
 - Designing indexes

Query Format and Annotation

Original Transformed

Annotation to display: All

Expand All Collapse All Customize Save Print Clear Highlights

Formatted Query	Annotation	Additional Information
SELECT CU.CUST_CODE		
, CU.CUST_LAST_NAME		
, COH.CUST_ORDER_NUMBER		
, DATE(COH.CUST_ORDER_DATE) AS CUST_ORDER_DATE		
, COD.PRODUCT_NUMBER		
, COD.CUST_QUANTITY		
FROM GOSALESC.T.CUST_CUSTOMER AS CU	CARDF=31,255 QUALIFIED_ROWS	
, GOSALESC.T.CUST_ORDER_HEADER AS COH	CARDF=539,410 QUALIFIED_ROWS	
, GOSALESC.T.CUST_ORDER_DETAIL AS COD	CARDF=560,273 QUALIFIED_ROWS	
WHERE (COH.CUST_CITY = 'Aberdeen'	COLCARDF=1,376 MAX_FREQ	
AND COH.CUST_ORDER_NUMBER = COD.CUST_ORDER_NUMBER	COLCARDF=539,410/491,520 MAX_F	
AND CU.CUST_CODE = COH.CUST_CODE	COLCARDF=31,255/28,672 MAX_F	
)		

Query Format and Annotation

Original **Transformed**

Annotation to display: All

Expand All Collapse All Customize Save Print Clear Highlights

Formatted Query	Annotation	EXT VIEW	Additional Information
SELECT CU.CUST_CODE			
, CU.CUST_LAST_NAME			
, COH.CUST_ORDER_NUMBER			
, DATE(COH.CUST_ORDER_DATE) AS CUST_ORDER_DATE			
, COD.PRODUCT_NUMBER			
, COD.CUST_QUANTITY			
FROM GOSALESC.T.CUST_CUSTOMER AS CU	CARDF=31,255 QUALIFIED_ROWS		
, GOSALESC.T.CUST_ORDER_HEADER AS COH	CARDF=539,410 QUALIFIED_ROWS		
, GOSALESC.T.CUST_ORDER_DETAIL AS COD	CARDF=560,273 QUALIFIED_ROWS		
WHERE (CU.CUST_CITY = 'Aberdeen'	COLCARDF=1,376 MAX_FREQ		
AND COH.CUST_ORDER_NUMBER = COD.CUST_ORDER_NUMBER	COLCARDF=539,410/491,520 MAX_F		
AND CAST(CU.CUST_CODE AS INTEGER) = COH.CUST_CODE	COLCARDF=28,672 MAX_FREQ		
)			

Improve Statistics Quality and Collection

- **Provides advice on**
 - Missing statistics
 - Problematic statistics
 - Out-of-date conflicts
 - Influential objects obsolete

- **Results**
 - Less CPU consumption
 - Superior access plan selection
 - Improved maintenance window throughput
 - Accurate estimated cost of the query

Provides visual indicator of priority

Advisor	Priority	Description
Statistics Advisor	HIGH	Gather and recollect all of relevant statistics for this query.
Statistics Advisor	MEDIUM	Determine the access path again at runtime. This query contains host variables, par

Generates RUNSTATS control statements

```

RUNSTATS ON TABLE "VIVEK"."EMPLOYEE"
WITH DISTRIBUTION ON COLUMNS ( "LASTNAME" NUM_FREQVALUES 15
NUM_QUANTILES 25 )
ALLOW WRITE ACCESS
  
```

Substantiates recommendation

Analysis end time: 2009-07-09 11:57:51.787

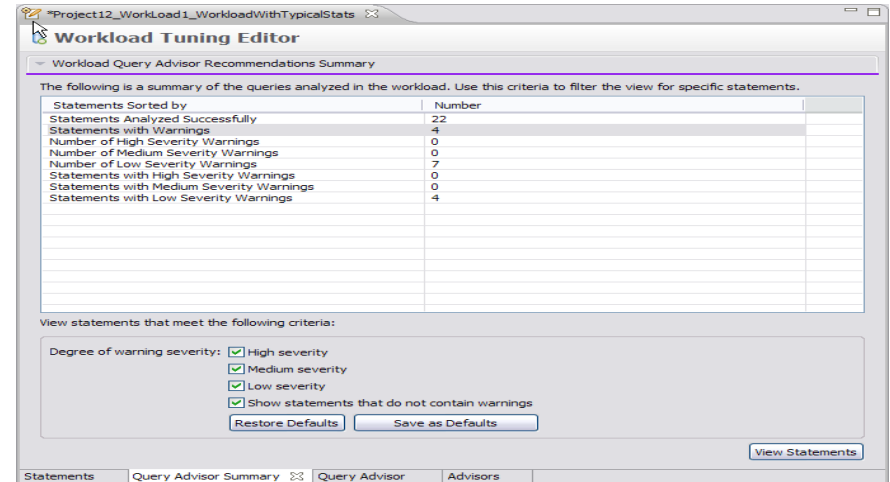
TABLE VIVEK.EMPLOYEE
Table type: Table
Cardinality: 42,0
Collection time: 2009-07-01 16:53:03.207
Statistics status: obsolete

Interesting columns:

LASTNAME
Cardinality: -1.0
Uniform statistics status: missing*
Frequency statistics status: missing*
Histogram statistics status: missing*
Possibly point skewed: Yes
Possibly range skewed: No
* Distribution statistics are gathered only when a sufficient lack of uniformity exists in the column values

Improve Query Design

- Query Advisor
- Get recommendations on:
 - Add missing join predicates for referential integrity
 - Use local predicates for joining tables or single table
 - Rewrite predicates as indexable predicates
 - Could use a matching index scan which is a more efficient access path



Workload Tuning Editor

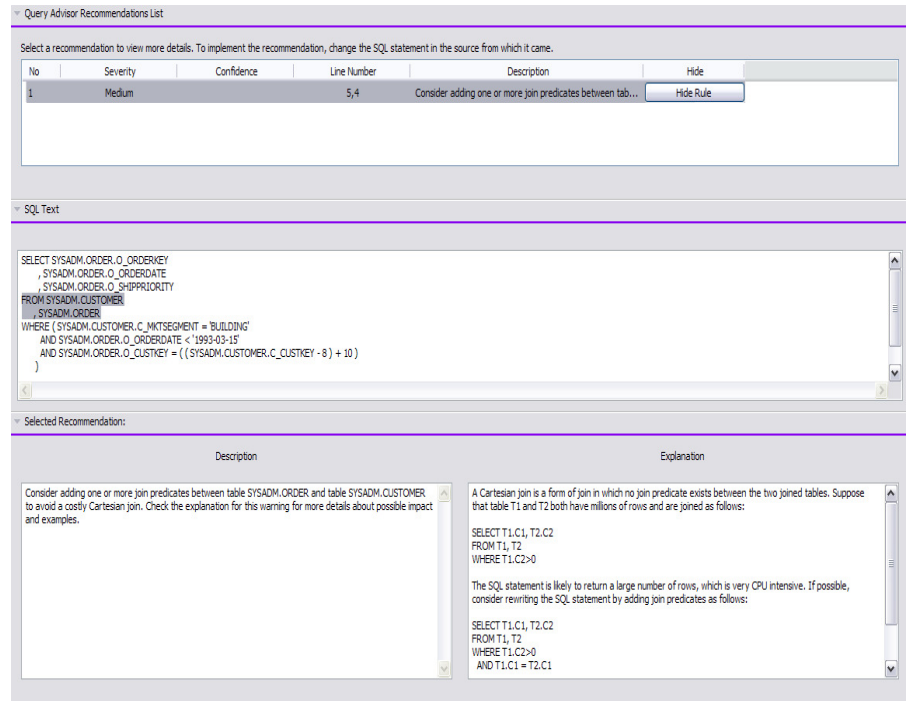
Workload Query Advisor Recommendations Summary

The following is a summary of the queries analyzed in the workload. Use this criteria to filter the view for specific statements.

Statements Sorted by	Number
Statements Analyzed Successfully	22
Statements with Warnings	4
Number of High Severity Warnings	0
Number of Medium Severity Warnings	0
Number of Low Severity Warnings	7
Statements with High Severity Warnings	0
Statements with Medium Severity Warnings	0
Statements with Low Severity Warnings	4

View statements that meet the following criteria:

Degree of warning severity: High severity
 Medium severity
 Low severity
 Show statements that do not contain warnings



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	Hide
1	Medium		5,4	Consider adding one or more join predicates between tab...	<input type="button" value="Hide Rule"/>

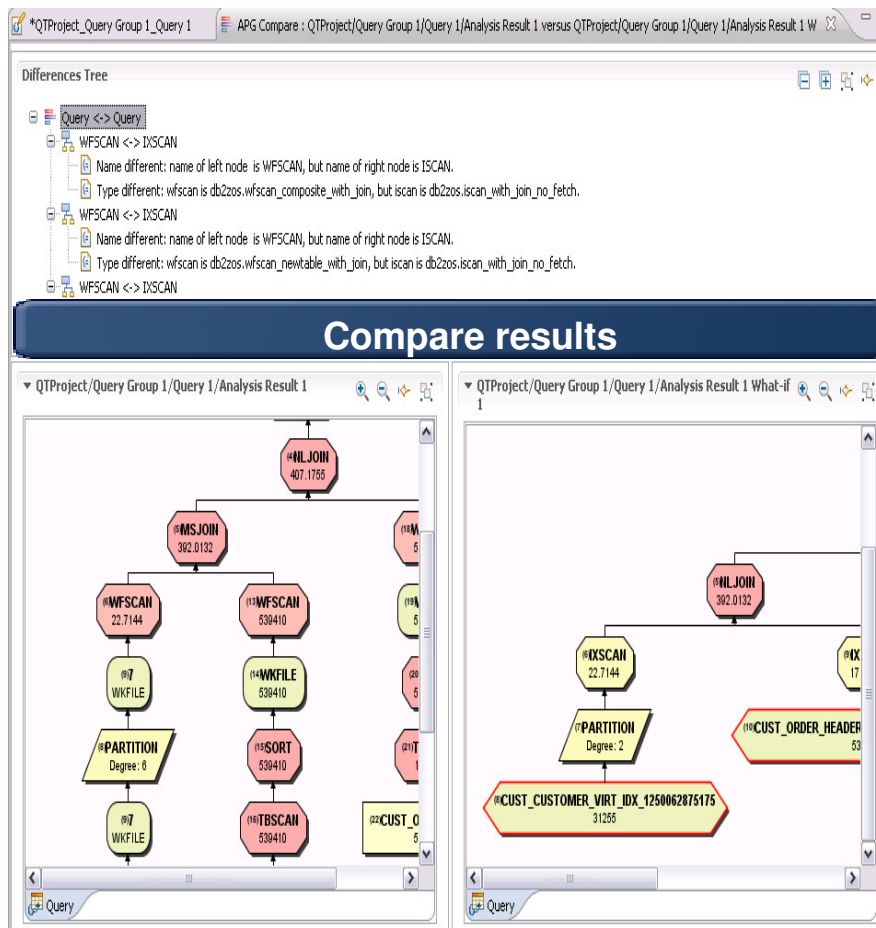
SQL Text

```
SELECT SYSADM.ORDER_O_ORDERKEY
, SYSADM.ORDER_O_ORDERDATE
, SYSADM.ORDER_O_SHIPRIORITY
FROM SYSADM.CUSTOMER
, SYSADM.ORDER
WHERE (SYSADM.CUSTOMER_C_MKTSEGMENT = 'BUILDING'
AND SYSADM.ORDER_O_ORDERDATE < '1993-03-15'
AND SYSADM.ORDER_O_CUSTKEY = ((SYSADM.CUSTOMER_C_CUSTKEY - 8) + 10)
)
```

Selected Recommendation:

Description	Explanation
Consider adding one or more join predicates between table SYSADM.ORDER and table SYSADM.CUSTOMER to avoid a costly Cartesian join. Check the explanation for this warning for more details about possible impact and examples.	A Cartesian join is a form of join in which no join predicate exists between the two joined tables. Suppose that table T1 and T2 both have millions of rows and are joined as follows: <pre>SELECT T1.C1, T2.C2 FROM T1, T2 WHERE T1.C2 > 0</pre> The SQL statement is likely to return a large number of rows, which is very CPU intensive. If possible, consider rewriting the SQL statement by adding join predicates as follows: <pre>SELECT T1.C1, T2.C2 FROM T1, T2 WHERE T1.C2 > 0 AND T1.C1 = T2.C1</pre>

Improve Access Path Stability



- **Why is it important?**
 - Reduces the risk of performance regression from minor changes in statistics or during maintenance

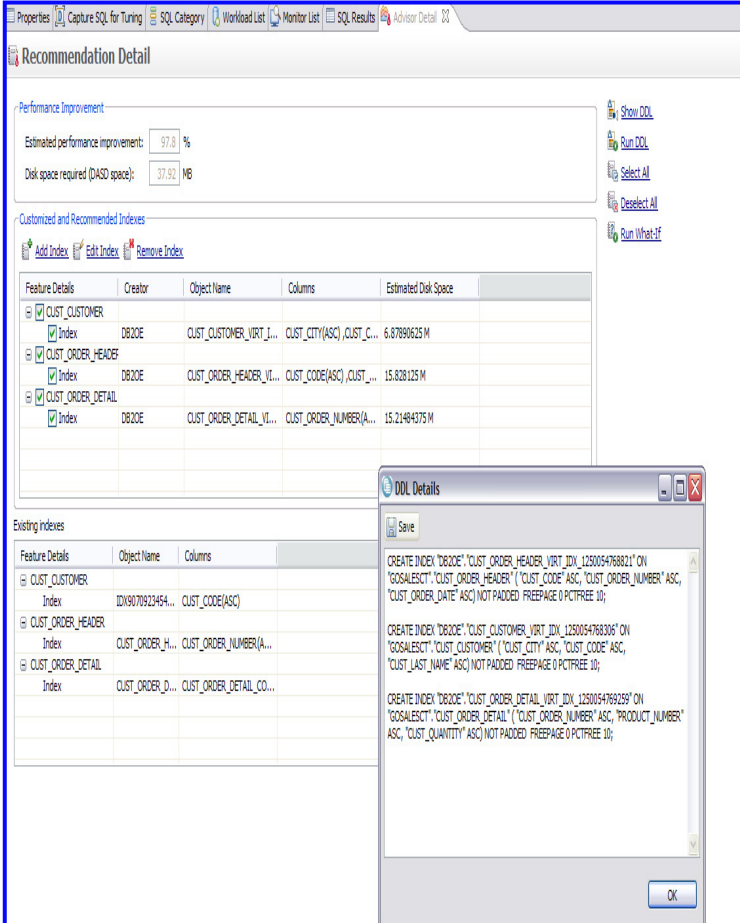
- **Key issues in access path selection**
 - Is the optimizer able to apply the filtering early?
 - Are there indexes that support an efficient path?
 - Do statistics allow distinction between the choices?

Improve query efficiency

- **Improve query efficiency**
 - Indexing foreign key in queries that do not have indexes defined
 - Identifying index filtering and screening
 - Support for index only access
 - Indexing to avoid sorts

- **Simplify use**
 - Consolidate indexes and provide a single recommendation

- **Test before deployment**
 - Utilizes virtual index capabilities built into the DB2 engine



The screenshot displays the IBM Optim SQL Tuning Advisor interface. The main window shows a 'Recommendation Detail' for a performance improvement of 97.8% with a disk space requirement of 37.92 MB. It lists three recommended indexes: CUST_CUSTOMER_VIRT_IDX, CUST_ORDER_HEADER_VIRT_IDX, and CUST_ORDER_DETAIL_VIRT_IDX. An 'Existing indexes' table is also visible, showing current indexes on the same tables. A 'DDL Details' dialog box is open, displaying the SQL DDL for creating the three recommended virtual indexes.

Feature Details	Creator	Object Name	Columns	Estimated Disk Space
<input checked="" type="checkbox"/> CUST_CUSTOMER				
<input checked="" type="checkbox"/> Index	DB2OE	CUST_CUSTOMER_VIRT_IDX	CUST_CITY(ASC), CUST_C...	6.8789625 M
<input checked="" type="checkbox"/> CUST_ORDER_HEADER				
<input checked="" type="checkbox"/> Index	DB2OE	CUST_ORDER_HEADER_VIRT_IDX	CUST_CODE(ASC), CUST...	15.828125 M
<input checked="" type="checkbox"/> CUST_ORDER_DETAIL				
<input checked="" type="checkbox"/> Index	DB2OE	CUST_ORDER_DETAIL_VIRT_IDX	CUST_ORDER_NUMBER(A...	15.21484375 M

Feature Details	Object Name	Columns
<input checked="" type="checkbox"/> CUST_CUSTOMER		
Index	ID19078923454...	CUST_CODE(ASC)
<input checked="" type="checkbox"/> CUST_ORDER_HEADER		
Index	CUST_ORDER_H...	CUST_ORDER_NUMBER(A...
<input checked="" type="checkbox"/> CUST_ORDER_DETAIL		
Index	CUST_ORDER_D...	CUST_ORDER_DETAIL_CO...

```

CREATE INDEX 'DB2OE'. 'CUST_ORDER_HEADER_VIRT_IDX_1250054768821' ON
'GOSALESCCT'. 'CUST_ORDER_HEADER' ( 'CUST_CODE' ASC, 'CUST_ORDER_NUMBER' ASC,
'CUST_ORDER_DATE' ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;

CREATE INDEX 'DB2OE'. 'CUST_CUSTOMER_VIRT_IDX_1250054768306' ON
'GOSALESCCT'. 'CUST_CUSTOMER' ( 'CUST_CITY' ASC, 'CUST_CODE' ASC,
'CUST_LAST_NAME' ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;

CREATE INDEX 'DB2OE'. 'CUST_ORDER_DETAIL_VIRT_IDX_1250054769259' ON
'GOSALESCCT'. 'CUST_ORDER_DETAIL' ( 'CUST_ORDER_NUMBER' ASC, 'PRODUCT_NUMBER'
ASC, 'CUST_QUANTITY' ASC) NOT PADDED FREEPAGE 0 PCTFREE 10;

```


Integrated Data Management DBA Summary

- **Optim provides has a rich ecosystem around the data life cycle**
 - **Enable alignment across the lifecycle**
 - Manage by intention
 - Share policies with downstream tasks
 - **Improve individual and team productivity**
 - Increase automation
 - Facilitate collaboration
 - Provide context-aware capabilities
 - **Optimize performance**
 - Embed best practices into solutions
 - Provide expert advice to increase skills
 - Improve resource utilization

Dúvidas?



Obrigado



- Por favor, entre em contato
- Rafael Coss
 - rcoss@us.ibm.com
 - Twitter: <http://twitter.com/racoss>
 - Blog: <http://datalifecycle.blogspot.com/>
 - LinkedIn <http://www.linkedin.com/pub/rafael-coss/9/311/268>
 - Facebook <http://www.facebook.com/racoss>
- Jay Bruce
 - jmbruce@us.ibm.com



Other Resources

Optim Community and Resources

- **Integrated Data Management community space on developerWorks – your portal to:**

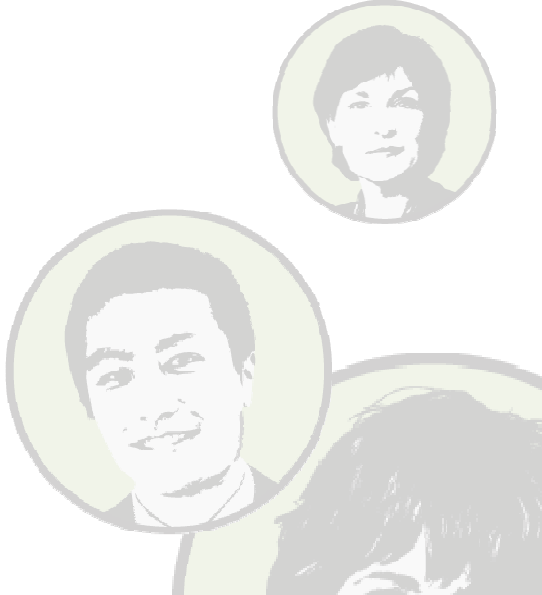
- Forums & blogs (PLEASE subscribe to Integrated Data Management Experts blog – lots of news comes through here)
- Downloads
- Technical articles, tutorials, product documentation
- News and events
- Videos (ChannelDB2 and youtube)
- Online Demos (developerWorks and marketing)
- Technical training
- Easy url: ibm.com/developerworks/spaces/optim

- **Optim Fan Facebook page**

- Log into Facebook and search on IBM Optim

- **Optim family page on developerWorks**

- ibm.com/developerworks/data/products/optim/



Kathy Zeidenstein sends out a weekly update with links to new articles, blogs, announcements, etc. Send her a note if you're interested.


krzeide@us.ibm.com

Integrated Data Management Community Space

Home
Articles & Tutorials
Downloads here!

Latest news

[Free virtual training!](#)



Basic we, D last year, technical briefings that are short enough to attend over a lunch hour (or a long coffee break). The goal is to get you educated on IDM products and provide you with access to product experts for your questions.

Our next virtual tech briefing is on August 20, 2009 at 1 PM Eastern. The topic is **Optim Development Studio 101** and will include a RAD expert as well as an Optim expert. [Register now!](#)

The schedule of upcoming briefings and links to replays can be found [here](#).

Forums

- [IBM Data Studio](#)
- [IBM Optim Development Studio and pureQuery Runtime](#)
- [IBM Optim Query Tuning Solution](#)
- [IBM Optim Database Administrator](#)
- [Optim Database Relationship Analyzer \(DRA\)](#)
- [Optim LUW](#)
- [Optim z/OS](#)
- [IBM InfoSphere Data Architect](#)
- [High Performance Unload](#)

Tags

Popular spaces tags | My spaces tags

0 2009 6000 a aix ajax and architecture articles as asc aws blogs brasil cascading champion cics **cloud** cloud_computing cloud_spa collaboration comunicatio community computing c data data_management da data_studio database datab database_forums db2 dedu demo design designer dev developers **developerv** development disk docum domino download downloa education emerging encryp energy engineering environ **green green_it** ibm in integration java managem planet report saas service **smarter_planet** soa s software **spaces** storage technology tools videos

View as **cloud** | list

About this space

What is IDM?

[Integrated Data Management: Managing data across its lifecycle](#)

Trying to figure it all out? Follow these easy steps:

1. Read the big picture article. Subscribe to the experts blog. See the Virtual tech briefing replay on **Data Studio becomes Optim: What it means for you.**
2. Check out the pureQuery platform web page and FAQ.
3. See the products in action with our [Day in the Life of a DBA](#) demo or our Optimizing Query Performance for z/OS demo or any of the other demos and videos on this page.
4. [Download the software.](#)

Performance

[Optimize your existing JDBC applications with pureQuery](#)

Now you can optimize both Java and .NET

Blogs

- [Integrated Data Management Experts](#)
- [Informix Application Development](#)
- [Fillmore Group blog](#)
- [Willie Favero on DB2 for z/OS](#)
- [Chris Eaton on DB2 for LUW](#)
- [Dave Beulke](#)

Demos

- [A Day in the Life of a DBA](#): Ever have one of those days? Jackie handles it well - find out how.
- [Introduction to InfoSphere Data Architect](#): 15 minutes of pure modeling fun with Oracle and DB2.
- [DB2 Performance Expert for LUW](#): Solving a runaway query problem.
- Query Tuning and client optimization for DB2 for z/OS: Two part demo: 1. Tune query using Optimization Expert. 2. Optimize existing JDBC apps using pureQuery.
- [Upgrading an application](#): The original integrated demo shows the Federal 1K team adding new features to an app and how it's easier

Click here for a comprehensive list of articles and tutorials.

Click here for links to all free and trial downloads.

Integrated Data Management Community Space

(more)

why you should care, for performance, manageability, and security reasons.

Be sure to check out the results of this [performance study](#) that compares static SQL using pureQuery with JDBC, EJB2 and OpenJPA.

You can get better performance using pureQuery for dynamic access as well. The WebSphere Customer Center team migrated from EJB 2 to pureQuery and got [great results](#).

Latest news: [Performance results using pureQuery with .NET applications](#)



Where we'll be

- [IOD NA \(October 25-20\)](#): Be there!
- [IDUG EMEA \(October 5-9\)](#): Curt Cotner goes to Rome.
- [Baltimore Washington DB2 Users Group \(September 9\)](#): For z/OS customers - Tom Howell on planning for upgrades and migrations and Data Privacy.



Learn More!

- [Basics of Data Studio Environment](#): Online course basics of IBM Data Studio Developer. NO CHARGE if you enroll in 2008.
- [Preparing for Stored Procedure Certification, Tutorial Series](#): Prep for Exam 735
- Developing Database Applications with IBM Data Studio and pureQuery: Three-day intensive course offered by InferData.
- [Using InfoSphere Data Architect](#): 2-day course.



Channel DB2 videos

- [How to Set Up, Configure, and Use IBM DB2 Performance Expert's Performance Warehouse](#)
- [DB2 Performance Expert Extended Insight Feature](#): Torsten walks you through the new capabilities for monitoring Java apps accessing DB2.
- NEW! Data Studio 2.2 Video: 5 short parts for those with short attention spans.
- [DB2 Performance Expert System Overview](#): Nathan walks you through the cool stuff in the PE dashboard.
- [Build an eforms solution with pureXML](#): Uses Data Studio to build the Web service.
- pureQuery explained: Rafael explains how pureQuery is different than other data access approaches.
- Sonali demonstrates Optim Development Studio 2.2: Oracle support and copy/paste are demo'd.
- NEW! Optim Query Tuning solution for Java developers: Tony Leung demonstrates the integration of Optim Development Studio and Optim Query Tuner.



Webcasts & podcasts

- [Integrated Data Management Solutions for Performance Optimization](#): You're going to be hearing a lot about integrated data management. Learn about an integrated solution to improving performance and manageability.
- [Data Studio for z/OS: From requirements to retirement, managing your data environment](#)
- Improve data security with Rational Data Architect V7.5
- [Monitoring and managing change for database administration](#)
- [Podcast: Curt Cotner with YL&A \(Cocktail Hour\)](#)
- [Podcast: Enhancing Java environments with pureQuery](#): Wow! You'll learn a lot in about the time it takes you to get to the grocery store.
- [Use pureQuery to improve QoS for WebSphere&DB2 apps](#)
- [Accelerating Solutions Delivery](#): No matter what your development methodology, it's always better to do faster and with better quality.
- NEW! Replay of Data Studio becomes Optim: What it means for you: First in IDM Virtual tech briefing series. Good overview of the major

leading-edge integrated data management capabilities, making it easier to manage data-centric tasks across the lifecycle spanning design, development, deployment, management, optimization, and governance. This space will serve as a central hub of information and conversation about Integrated Data Management products, technologies, and best practices.

Objective: To share experiences with IBM Data Studio users.

Audience: Data Architects
Database Administrators
Data Stewards
Database developers
Java developers

Group type: Public

Date created: 19 Sep 2007

► [Show member list](#)

Optim Fan Facebook page



Edit Page
Promote Page with an Ad
Add Fan Box to your site
More

Write something about IBM Optim.

Information

Founded:
The Data Studio and Optim portfolios are part of IBM's Integrated Data Management software initiative and help organizations manage application data with greater consistency across the enterprise.

Insights [See All](#)

6.1 Post Quality

1 Interactions This Week

Most Active Countries

United States	1
----------------------	---

Insights are visible to page admins only.

IBM Optim

Wall Info Photos Boxes +

What's on your mind?

Attach Share


IBM Optim Just Fans Settings

IBM Optim Next virtual tech briefing August 20: Optim Development Studio 101. Get a tour of Optim Development Studio from a product expert and learn how Optim Development Studio can extend the capabilities in Rational Application Developer to turbo-charge the development and optimization of data persistence layers. Register here: https://www.informationmanagementrequest.com/mk/get/OPTIM_LL1_REG

Reg Page
Source: www.informationmanagementrequest.com

July 24 at 11:36am · Comment · Like · Share

IBM Optim In case you didn't hear, yesterday IBM announced new releases and new products for z/OS. Optim Query Tuner for DB2 for z/OS, Optim Query Workload Tuner for DB2 for z/OS and Optim pureQuery Runtime for z/OS. You can find links to the announcements in Curt Cotner's recent blog here: http://www.ibm.com/developerworks/blogs/page/datastudioteam?entry=new_announcements_for_z_optim



IBM developerWorks : Blogs : Managing the data lifecycle: from design to deletion

Source: www.ibm.com


Our team of Integrated Data Management experts from the lab offer insights, advice, hints and tips, and lessons learned. Topics include everything from data modeling and database application development to database administration, performance monitoring, and more.


Blog (ibm.com/developerworks/mydeveloperworks/blogs/idm/)

developerWorks > My developerWorks > Connect


My developerWorks: Blogs

Browse Blogs | **My Blog** | My Updates

 [Subscribe to Blog Entries](#)

 [Subscribe to Blog Comments](#)

Tags



.net 2008 **administration**
 administrator agile
 announcement
 announcements appliance
 architect archive baumbach
 bhagavan bireley
bommireddipalli brodsky
 budget bui business_analyst
 change chen china cio
 client_optimization commands
 conference correlation COSS
cotner cpu data data_architect
 data_model data_privacy
 datapower **dba** ddf **demo**
 deploy **design** **develop**
 developer development
 download dsac eclipse egl
 erwin faq feedback fixpacks
 free fun glossary **nonal** group

Managing the data lifecycle

Blog Authors: IBM_Optim DataStudio KRZ RafaelCoss gopalv Anson_Forum
 vijayrb@us.ibm.com hollyann

New Entry

Settings

Sort by: Date | Title | Most Recommendations | Most Comments | Most Visits

Main | Next

Understanding SQL Replacement with pureQuery's Cli...

★ 2

IBM_Optim | 7 Aug | Tags: [purequery](#) [client_optimization](#) [bireley](#) | Comments (0)

Hi, I'm the lead architect for Optim pureQuery Runtime, and I want to start using this blog to help address questions that I get from people as they learn about or use pureQuery capabilities. In this first blog, I'll discuss the new SQL replacement capability.

By using client optimization, an administrator can modify the SQL from a captured application. The enhanced tooling to support this capability is described in Sonali's article, [What's new and cool in Optim Development Studio 2.2](#). The intended usage of this feature is to let a DBA make a change to an SQL statement without the need to edit and recompile an application. This could be useful, for example, in late night or weekend emergencies when an application can't easily be changed. It is also useful in cases where a third party application embeds or generates sub-optimal SQL and a change to the application is not possible without contacting the vendor. In any of these cases, you should aim to change the application directly at the first practical opportunity to use the improved SQL.

When I talk about this capability to people, there are two questions that frequently come up:

- What is the extent of the change I can make to the SQL?
- Isn't there a security risk to allow editing of the SQL? How can we control access?