

# Two for One: Optim Solutions Update and Optimizing DB2 for z/OS Query Performance

Bryan F. Smith [bfsmith@us.ibm.com](mailto:bfsmith@us.ibm.com)  
IBM

Tuesday, August 3, 2010: 1:30 PM-2:30 PM  
Hynes, Room 103



**SHARE** in Boston



# Abstract


- **Optim Solutions Update:** Are you wondering how Optim Solutions relates to your z/OS DB2 Tools investment? Do you want to see what the lab is doing to help the z/OS DBA? Join one of our chief architects to hear how Optim solutions can extend administrative function for DB2 for z/OS, both today and in the future. The speaker will cover data design, performance and query monitoring, performance and query optimization, and key roadmap items.
- **Optimizing DB2 for z/OS Query Performance:** There are lots of factors that affect query performance and no silver bullets. Investigating performance issues includes looking at indexes, statistics, query design, cache size, execution mode, and the list goes on. Join the speaker as he digs into some of these areas in detail and at how Optim Solutions can help get resolutions faster and with more flexibility.

# Agenda


- Introduction to Integrated Data Management (IDM)
- What can Data Studio be used for in DB2 for z/OS environments?
- How Java and .NET applications can be managed like CICS/IMS applications
- What results can be realized by query tuning
- Data movement options
- What value is there to extending the DB2 backup and recovery utilities?
- Other DB2 for z/OS Tooling
- What's cooking back at the lab
  - Performance Management
  - Configuration Management
  - Maintenance Management
- Wrap-up

# Data Management Must Drive Competitive Advantage


Survey: CIO's want to strengthen competitive advantage by better managing enterprise data



75% of CIO's believe they can strengthen their competitive advantage by better using and managing enterprise data.



78% of CIO's want to improve the way they use and manage their data.



...but **only 15%** believe that their data is currently comprehensively well managed.

- ***Data management must drive competitive advantage***
- Facilitate business alignment
- Accelerate delivery of enterprise-ready, data-driven applications
- Free up staff to drive business growth and optimization

# Introducing Integrated Data Management



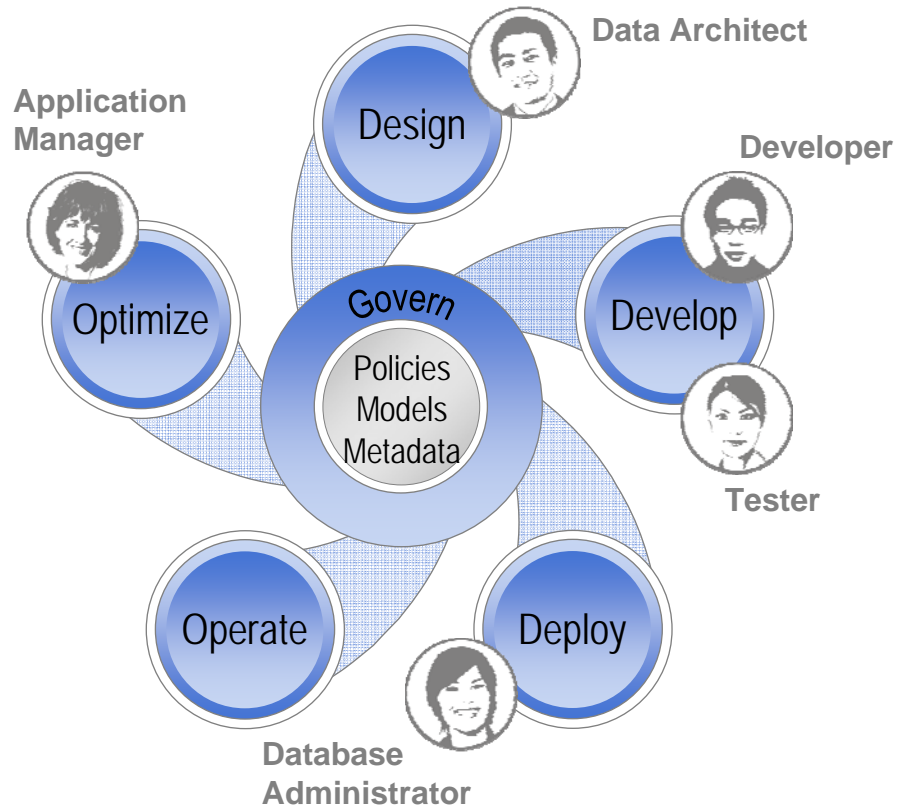
An integrated, modular environment to manage enterprise application data, and optimize data-driven applications, from requirements to retirement



- ***Grow the business, without growing costs***
  - Develop and deploy business critical applications faster
  - Mitigate compliance risks with model-driven data governance
  - Prevent runaway infrastructure spending
  - Improve performance – of work teams, databases, applications, and business units

**Integrated Data Management = Data Studio+Optim+DB2 for z/OS Tools**

# Integrated Data Management

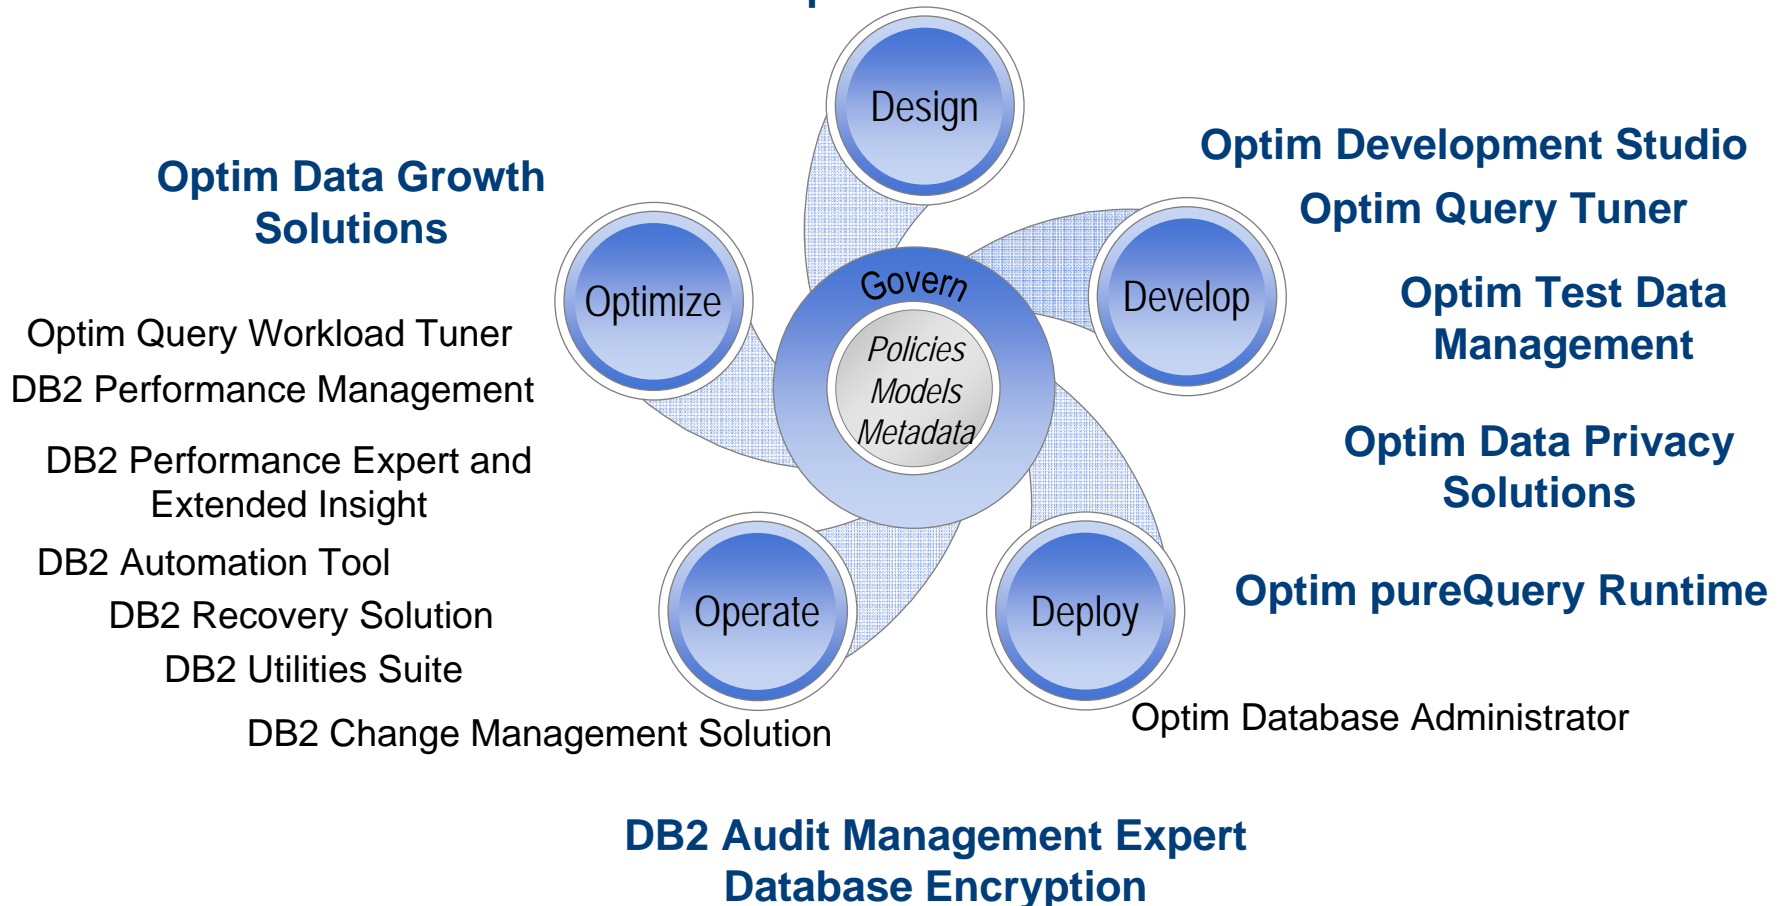


- **Deliver increasing value across the lifecycle**, *from requirements to retirement*
- **Facilitate collaboration and efficiency across roles**, *via shared artifacts automation and consistent interfaces*
- **Increase ability to meet service level agreements**, *improving problem isolation, performance optimization, capacity planning, and workload and impact analysis*
- **Comply with data security, privacy, and retention policies** *leveraging shared policy, services, and reporting infrastructure*

# Supporting Heterogeneous Environments



## InfoSphere Data Architect




# Integrated Data Management User Interface Directions



## Data Studio and Optim

## z/OS Tools

### Eclipse UI



### Develop and Deploy

- Design and Development
- Configuration
- Object Management
- Change Management
- Automation Planning
- Data Governance

### Web UI



### Operate

- Monitoring
- Automation Management
- Reporting
- Data Governance

### 3270



### z/OS Database Administration

- Rich z/OS function
- Application Management
- Database Administration
- Performance Management
- Data Governance



# DB2 for z/OS Tools Portfolio



## Application Management

- InfoSphere Data Architect
- Optim Development Studio
- DB2 Path Checker
- DB2 Bind Manager
- DB2 SQL Performance Analyzer
- DB2 Table Editor
- Optim pureQuery Runtime

## Utilities Management

- DB2 Utilities Suite
- DB2 Automation Tool
- DB2 Automation Toolkit SAP Edition
- DB2 Utilities Enhancement Tool
- DB2 High Performance Unload

## Business Intelligence

- IBM DataQuant
- IBM QMF
- DB2 Web Query Tool

## Database Administration

- DB2 Administration Tool
- DB2 Object Comparison Tool
- DB2 Administration Toolkit SAP Edition
- DB2 Storage Management Utility

## Performance Management

- OMEGAMON XE DB2 Performance Expert
- OMEGAMON XE DB2 Performance Monitor
- DB2 Buffer Pool Analyzer
- Optim Query [Workload] Tuner
- DB2 Query Monitor
- DB2 Performance Toolkit SAP Edition

## Information Integration

- WebSphere Classic Data Event Publisher
- WebSphere Classic Federation Server
- WebSphere Classic Replication Server
- WebSphere DataStage
- WebSphere Replication Server
- DataMirror

## Backup and Recovery

- Application Recovery Tool for IMS and DB2 Databases
- DB2 Archive Log Accelerator
- DB2 Change Accumulation Tool
- DB2 Cloning Tool
- DB2 Log Analysis Tool
- DB2 Object Restore Tool
- DB2 Recovery Expert

## Data Governance

- IBM Optim Data Growth
- IBM Optim Data Privacy
- IBM Optim Test Data Management
- DB2 Audit Management Expert
- IBM Database Encryption Expert
- Data Encryption for DB2 and IMS

# DB2 Developer Workbench vs. Data Studio



*before*

*now*

## **IBM DB2 Developer Workbench V9.1**

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

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

## **IBM Data Studio V2.2**

- 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
- ER Diagramming
- Data Distribution Viewer
- Object Management
- Browse & Update Statistics
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support

**No-charge**

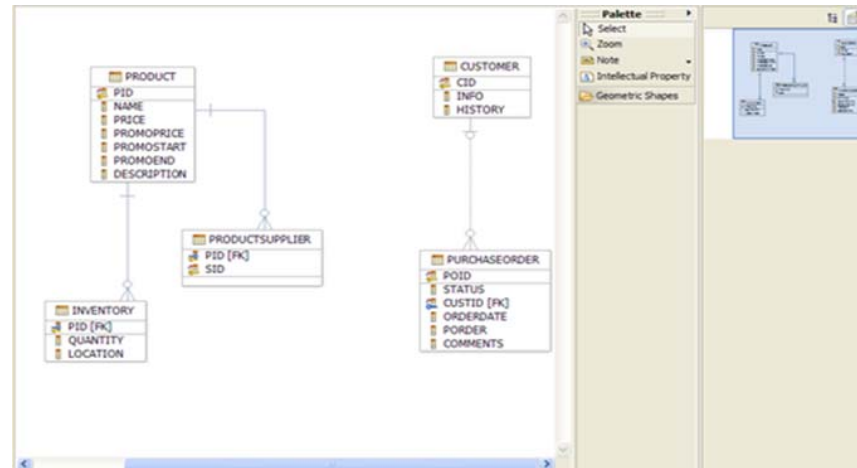
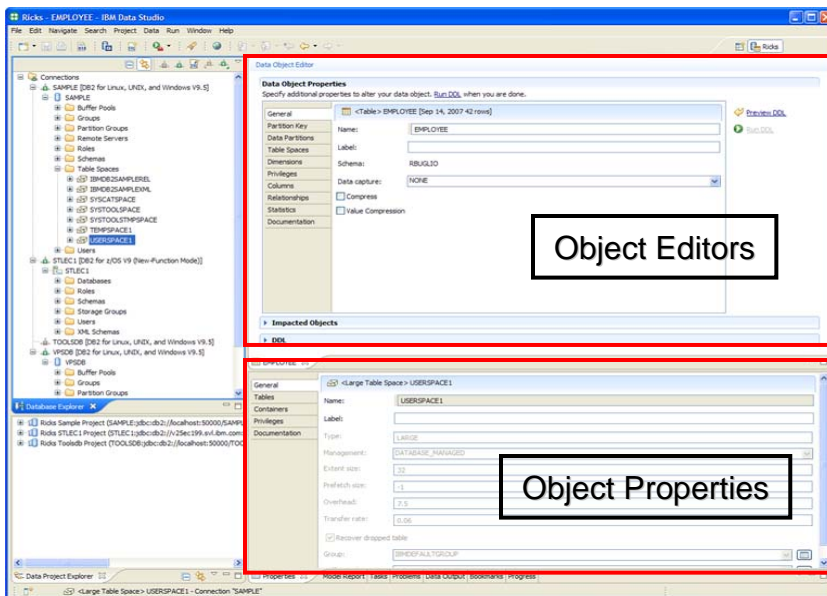
**Data Studio is no longer used as a brand**

# Data Studio: Data Modeling / App & DB Dev

## A Consistent and Productive work environment

Create, Alter, Drop, Browse and Filter database objects

ER Diagramming



Integrated Query Editor

- Express yourself with optimal queries
- Content assistance for database objects
- Rapid interactive end-user feedback
- Extensible templates
- Multiple SQL statement testing
- SQL assistance and XQuery assistance

```

-- <ScriptOptions statementTerminator=";" />
XQUERY declare default element namespace "http://posample.org";
for $city in fn:distinct-values(db2-fn:xmlcolumn('CUSTOMER.INFO')/custo
return
<city name ='($city)'+
(
for $cust in db2-fn:xmlcolumn('CUSTOMER.INFO')/customerinfo[addr
let $po := db2-fn:sqlquery("SELECT XMLELEMENT( NAME ""pos"",
(XMLCONCAT( XMLELEMENT(NAME ""custid"", c.custid),
XMLELEMENT(NAME ""order"", c.porder)))
FROM purchaseorder AS c")
let $id := $cust/@Cid,
$order := $po/pos[custid=$id]/order
return
<customer id='($id)'+

```

# Unleash SQL in your Java IDE

- SQL content assist

```
// Select GOSALEST.CUST by parameters
@Select(sql = "SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CUST_ADDR1,"
+ " CUST_ADDR2, CUST_CITY, CUST_POST_ZONE, CUST_CTRY_CODE, CUST_PHN_NBR,"
+ " CUST_INFO, CUST_EMAIL, CUST_GNDR_CODE, CUST_PROV_STATE"
+ " FROM GOSALEST.CUST"
+ " WHERE ")
Cust getCust(int cust_code)
```

Press CTRL-SPACE to invoke content assist or auto complete

CUST_ADDR1 - VARCHAR(128)
CUST_ADDR2 - VARCHAR(128)
CUST_CITY - VARCHAR(128)
CUST_CODE - INTEGER
CUST_CTRY_CODE - VARCHAR(128)
CUST_EMAIL - VARCHAR(128)
CUST_FRST_NAME - VARCHAR(128)
CUST_GNDR_CODE - DECIMAL(3, 0)
CUST_INFO - CLOB(32768)

Press 'Ctrl+Space' to show SQL Proposals

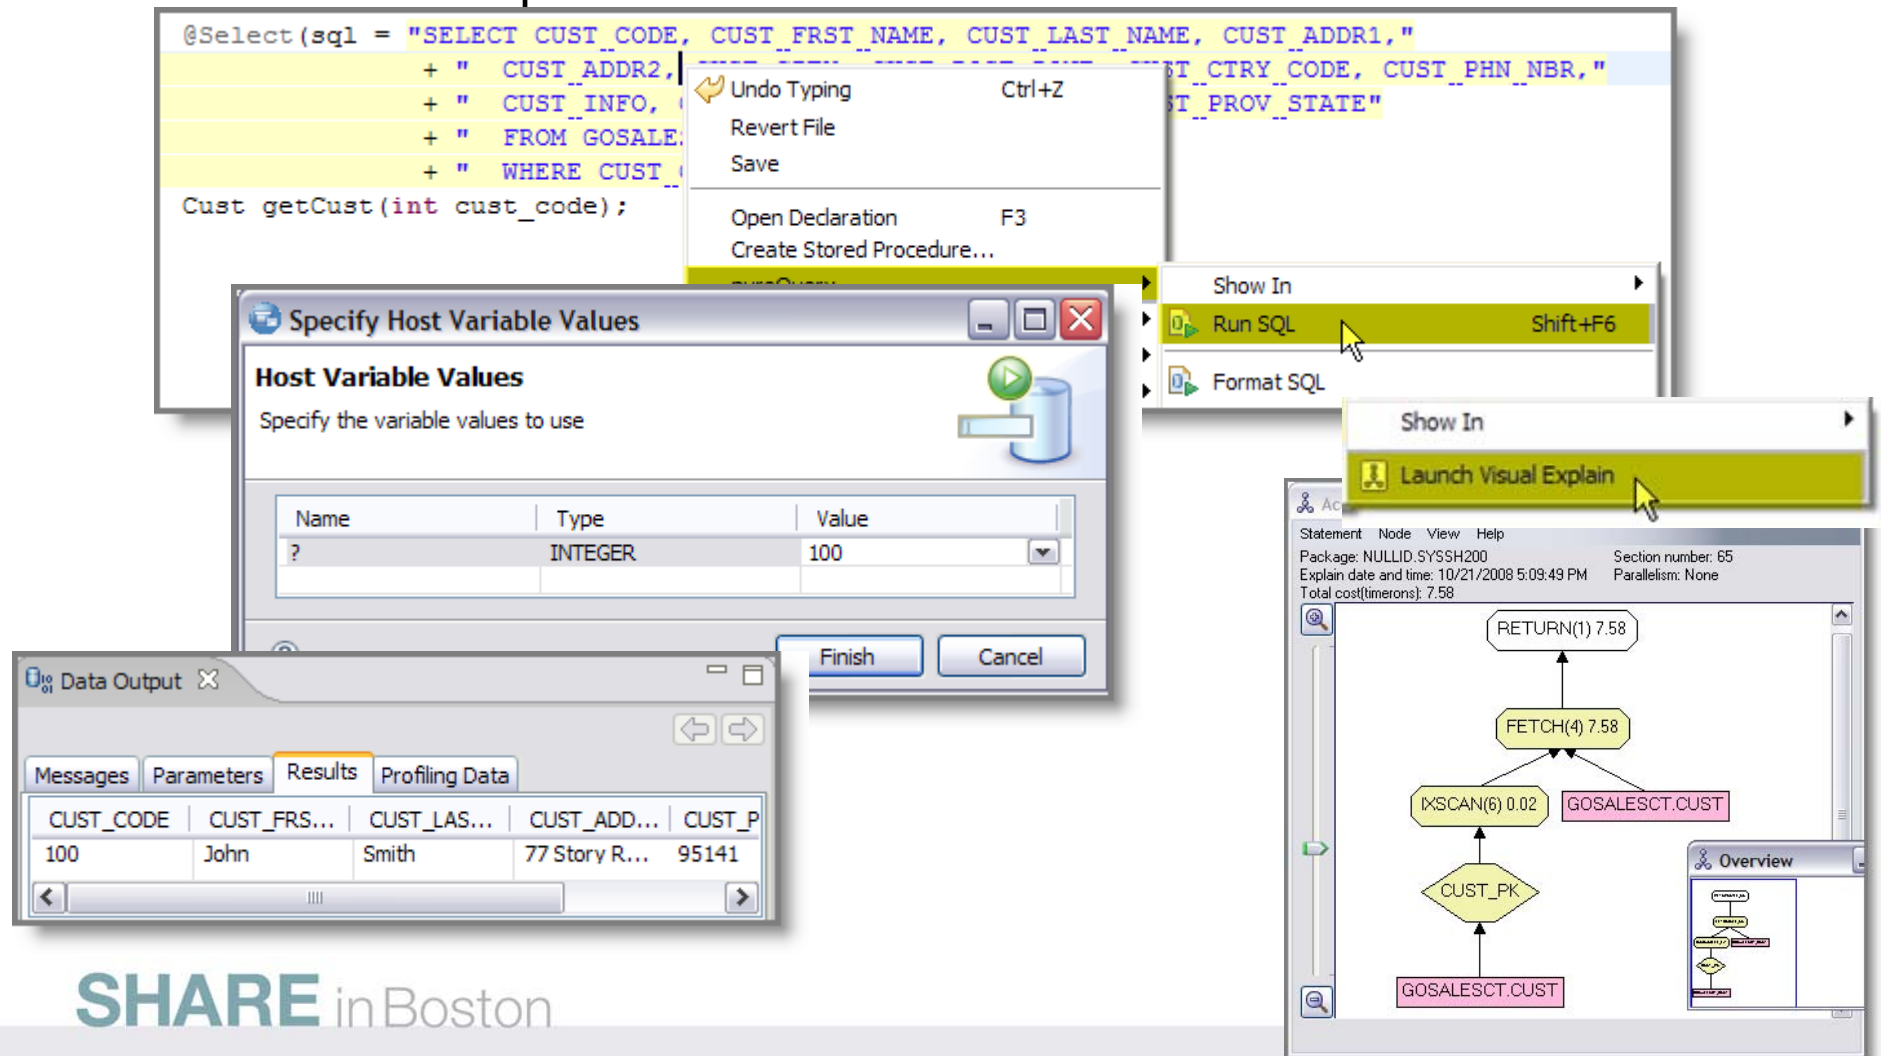
- SQL validation

```
// Select GOSALEST.CUST by parameters
@Select(sql = "SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CUST_ADDR1,"
+ " CUST_ADDR2, CUST_CITY, CUST_POST_ZONE, CUST_CTRY_CODE, CUST_PHN_NBR,"
+ " CUST_INFO, CUST_EMAIL, CUST_GNDR_CODE, CUST_PROV_STATE"
+ " FROM GOSALEST.CUST"
+ " WHERE CUST_CODE = ?")
Cust getCust(int cust_code)
```

Table "CUST" does not contain column "CUST\_COD".  
Press 'F2' for focus.

# Unleash SQL in your Java IDE

- Run SQL at design time without the need to compile the program
- Launch Visual Explain



The screenshot illustrates the workflow for running SQL and generating a visual explain plan in a Java IDE. It shows the following components:

- SQL Editor:** Contains a SQL query:
 

```
@Select(sql = "SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CUST_ADDR1,"
+ " CUST_ADDR2,"
+ " CUST_INFO,"
+ " FROM GOSALE...
+ " WHERE CUST..."
Cust getCust(int cust_code);
```
- Context Menu:** A menu is open over the SQL code, with options like "Undo Typing", "Revert File", "Save", "Open Declaration", and "Create Stored Procedure...".
- Specify Host Variable Values Dialog:** A dialog box for specifying variable values. It has a table with the following data:
 

Name	Type	Value
?	INTEGER	100
- Run SQL Menu Item:** A "Run SQL" menu item (Shift+F6) is highlighted in the context menu.
- Launch Visual Explain Menu Item:** A "Launch Visual Explain" menu item is highlighted in a secondary context menu.
- Data Output Window:** Shows the results of the SQL query in a table:
 

CUST_CODE	CUST_FRST...	CUST_LAS...	CUST_ADD...	CUST_P
100	John	Smith	77 Story R...	95141
- Visual Explain Plan:** A graphical execution plan showing the flow of data. The plan starts with a table scan (GOSALESC.T.CUST), followed by a join (CUST\_PK), a scan (XSCAN(6) 0.02), a fetch (FETCH(4) 7.58), and finally a return (RETURN(1) 7.58). An "Overview" window is also visible in the bottom right.



# Stored Procedure Support SQL and Java Debugger



The screenshot displays the SQL and Java Debugger interface. The top toolbar contains various icons, with a red box highlighting the 'Debug' button. The 'Servers' pane on the left shows the loaded stored procedure: GHUTCHIS.CHECK\_ORDERS(VARCHAR) - tuesday [DB2 Stored Procedure]. The 'Variables' pane on the right shows diagnostic information with a green box labeled 'variables' around it. The main editor shows the source code of the stored procedure, with a green box labeled 'source' around it. The console at the bottom shows the execution status and messages.

**Debug commands**

**variables**

```
P1: BEGIN
-- Declare cursor
DECLARE cursor1 CURSOR WITH RETURN FOR
SELECT PURCHASEORDER.POID, PURCHASEORDER.STATUS, PURCHASEORDER.ORDERDATE,
PURCHASEORDER.PORDER, PURCHASEORDER.COMMENTS, CUSTOMER.CID,
CUSTOMER.INFO, CUSTOMER.HISTORY
FROM CUSTOMER JOIN PURCHASEORDER ON CUSTOMER.CID = PURCHASEORDER.CUSTID
WHERE PURCHASEORDER.STATUS = V_STATUS;

-- Cursor left open for client application
OPEN cursor1;
```

**source**

Status	Action	Object Name
In progress	Debug	CHECK_ORDERS
Success	Run	CHECK_ORDERS
Success	Deploy	CHECK_ORDERS
Success	Run	s1.sql
Success	Run	s1.sql

GHUTCHIS.CHECK\_ORDERS(IN V\_STATUS VARCHAR(10))

Messages Parameters Results Profiling Data

GHUTCHIS.CHECK\_ORDERS - Run started.

Data returned in result sets is limited to the first 50 rows.

Data returned in result set columns is limited to the first 100 bytes or characters.

GHUTCHIS.CHECK\_ORDERS - Calling the stored procedure.

Unshipped

SAMPLE (SAMPLE: jdbc:db2://localhost:50000/SAMPLE:retrieveMessagesFromServerOnGetMessage=true;) Writable Smart Insert 3 : 73

*Integrated editor and debugger for SQL PL and Java (JDBC or SQLJ)*

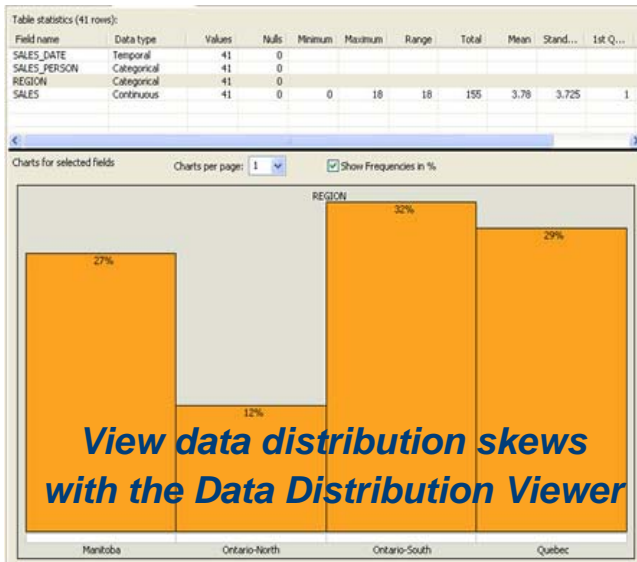
**SHARE** in Boston

# Data Management

## Edit Data



## Data Distribution Editor



EMPNO [CHAR(6)]	FIRSTNAME [VARCHAR(12)]	MIDINIT [CHAR(1)]	LASTNAME [VARCHAR(15)]	WORKDEPT [CHAR(3)]	PHONE [CHAR(4)]	HIREDATE [DATE]	JOB
000010	CHRISTINE	I	MAAS	AOO	3978	1/1/95	PRES
000020	MICHAEL	L	THOMPSON	BO1	3476	10/10/03	MANU
000030	SALLY	A	KUJAN	CO1	4738	4/5/05	MANU
000050	JOHN	B	GEYER	EO1	6789	8/17/79	MANU
000060	IRVING	F	STERN	O11	6423	9/14/03	MANU
000070	EVA	D	PULASKI	O21	7831	9/30/05	MANU
000090	EILEEN	W	HENDERSON	E11	5498	8/15/00	MANU
000100	THEODORE	Q	SPENCER	E21	0972	6/19/00	MANU
000110	VINCENZO	G	LUCCHESI	AOO	3490	5/16/88	SALE
000120	SEAN	A	O'CONNELL	AOO	2187	12/5/93	CLER
000130	CELORES	M	QUINTANA	CO1	4578	7/28/01	ANAL
000140	HEATHER	A	NICHOLS	CO1	1793	12/15/06	ANAL
000150	BRUCE	A	ADAMSON	O11	4510	2/12/02	DESD
000160	ELIZABETH	R	PSANKA	O11	3782	10/11/06	DESD
000170	MASATOSHI	J	YOSHIMURA	O11	2890	9/15/99	DESD
000180	MARILYN	S	SCOUTTEN	O11	1682	7/7/03	DESD
000190	JAMES	H	WALKER	O11	2986	7/26/04	DESD
000200	DAVID	B	BROWN	O11	4501	3/3/02	DESD
000210	WILLIAM	T	JONES	O11	0942	4/1/98	DESD
000220	JENNIFER	K	LUTZ	O11	0672	8/29/98	DESD
000230	JAMES	J	JEFFERSON	O21	2094	11/21/96	CLER
000240	SALVATORE	M	MARINO	O21	3780	12/5/04	CLER
000250	DANIEL	S	SMITH	O21	0961	10/30/99	CLER
000260	SYBIL	P	JOHNSON	O21	8953	9/11/05	CLER
000270	MARIA	L	PEREZ	O21	9001	9/30/06	CLER
000280	ETHEL	R	SCHNEIDER	E11	8997	3/24/97	OPEX
000290	JOHN	R	PARKER	E11	4502	5/30/06	OPEX
000300	PHILIP	X	SMITH	E11	2095	6/19/02	OPEX
000310	MAUDE	F	SETRUGHT	E11	3332	9/12/94	OPEX
000320	RAMLAL	V	MEHTA	E21	9990	7/7/95	FISLL
000330	WING	L	LEE	E21	2103	2/23/06	FISLL
000340	JACKSON	R	GOLDFOT	E21	5698	5/5/77	FISLL
000350	RYAN	T	HFBATHUR	AOO	1078	1/1/64	SU F

## Roles, Users, Privileges

The screenshot shows the Access Data Object Editor interface. On the left, the 'Database Explorer' pane shows a tree view of the database objects, with a red box highlighting the 'Roles' folder and a 'Manage Roles' label. The main window displays the 'Data Object Properties' dialog for a role, with a red box highlighting the 'Privileges By User' tab, which lists various permissions for different users. A second red box highlights the 'Privileges By Object' tab, which shows a table of permissions for a specific object (EMPLOYEE).

Object	Grantee	Grantee Type	Privilege	Grantor	WITH GRANT OPTION
EMPLOYEE	RBUGLJO	User	ALTER	SYSTEM	
	RBUGLJO	User	CONTROL	SYSTEM	
	RBUGLJO	User	DELETE	SYSTEM	
	RBUGLJO	User	INDEX	SYSTEM	
	RBUGLJO	User	INSERT	SYSTEM	
	RBUGLJO	User	REFERENCES	SYSTEM	
EMPLOYEE	RBUGLJO	User	SELECT	SYSTEM	
	RBUGLJO	User	UPDATE	SYSTEM	

# IBM Data Studio



United States [ change ]

Home Solutions ▾ Services ▾ Products ▾ Support & downloads ▾ My IBM ▾

**Software**

- Products
- Services
- Downloads
- Library
- News
- Training and certification
- Events
- Support

Communities:

- IBM Business Partners
- ISVs
- Developers

## IBM Data Studio (stand-alone)

**Downloads**

To properly configure your download, please review the information below. Select the appropriate offering. When you are done, press the "Continue" button at the bottom.

Offering	Platform	Format
<input type="radio"/> IBM Data Studio Version 2.2 Languages: English	Red Hat Linux SUSE Linux Enterprise Desktop(SLED) SUSE Linux Enterprise Server (SLES)	download
<input type="radio"/> IBM Data Studio Version 2.2 Languages: English	Windows Vista Business Windows Vista Enterprise Windows Vista Ultimate Windows XP Professional	download

No-charge

Download it today!



# IBM Optim Development Studio



**An integrated database development environment that speeds application design, development, and deployment while increasing data access performance and manageability.**



- Enhance developer productivity
  - Drag and drop creation of Web services
  - Provide a seamless SQL/Java experience
  - Generate a data access layer using Java objects, JSON, or, XML
  - Enhance problem isolation and impact analysis, even when using frameworks that generate the SQL
- Provide expert-equivalent performance
  - Facilitate use of JDBC and SQL data access best practices
  - Improve DB2 performance, predictability, and manageability by enabling transparent activation of static SQL (i.e. no change to the application) for Java and .NET
  - Facilitate DBA collaboration and optimization
- Enhance security
  - Eliminate SQL injection risk



“Development Studio enables us to bridge the gap between object-oriented design and relational database technology. As a result, we can speed the development of high quality applications and improve developer productivity by between 25 and 50 percent”

# pureQuery



- How Java and .NET applications can be managed like CICS/IMS applications
  - What problems do these application platforms present to the DBA?
    - ORM frameworks obscure the SQL generated
    - Dynamic SQL
  - Capturing SQL for review/tune/revisement
  - Converting to static SQL and the benefits of doing so

# Frequently Cited Concerns



**I have more and more Java workload coming into my data server driving up costs, but the budget is not keeping pace.**

**I don't even want to allow framework-generated SQL on my database. If I can't see it, I don't know how it will impact me.**

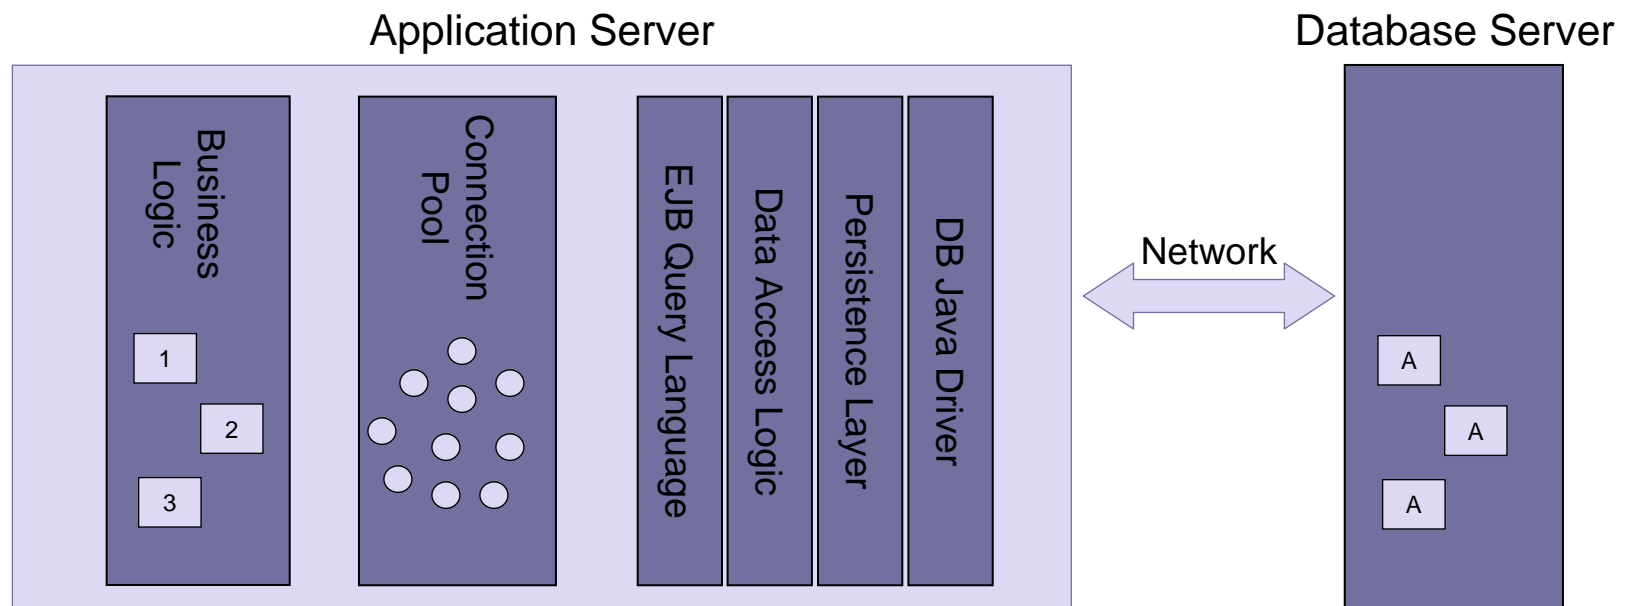


**Java and .NET performance problems are a real pain to resolve because I can't even tell what application issued the SQL.**



# Contemporary Application Stack Challenges

- Simplify development, but ...
  - Challenge problem resolution
  - Impact performance
  - Obscure impact analysis
  - Impede capacity planning



# Introducing pureQuery



A high-performance, data access platform to simplify developing, managing, securing, and optimizing data access for new and existing applications.



## *pureQuery Components:*

- Optim Development Studio
  - Integrated development environment with Java and SQL support
  - Improve problem isolation and impact analysis
- Simple and intuitive API
  - Enables SQL access to databases or in-memory Java objects
  - Facilitates best practices
- Optim pureQuery Runtime
  - Flexible static SQL deployment for DB2

SHARE in Boston



and



# pureQuery Improves Performance, Security, and Manageability for DB2 and Oracle... *Without Changing a Line of Code*

## Three steps

### 1. Capture the SQL

- Use pureQuery API, generate from WebSphere JPA, or capture while executing
- Use with custom-developed, framework-based, or packaged applications
- **Translation of literals to host variables (new in 2.2)**

### 2. Bind SQL to DB2 (Client Optimization)

- Use tooling in Data Studio Developer, WAS console or command line

### 3. Choose execution mode

- Dynamic or static
- Choose at deployment time instead of development time

## Static SQL value

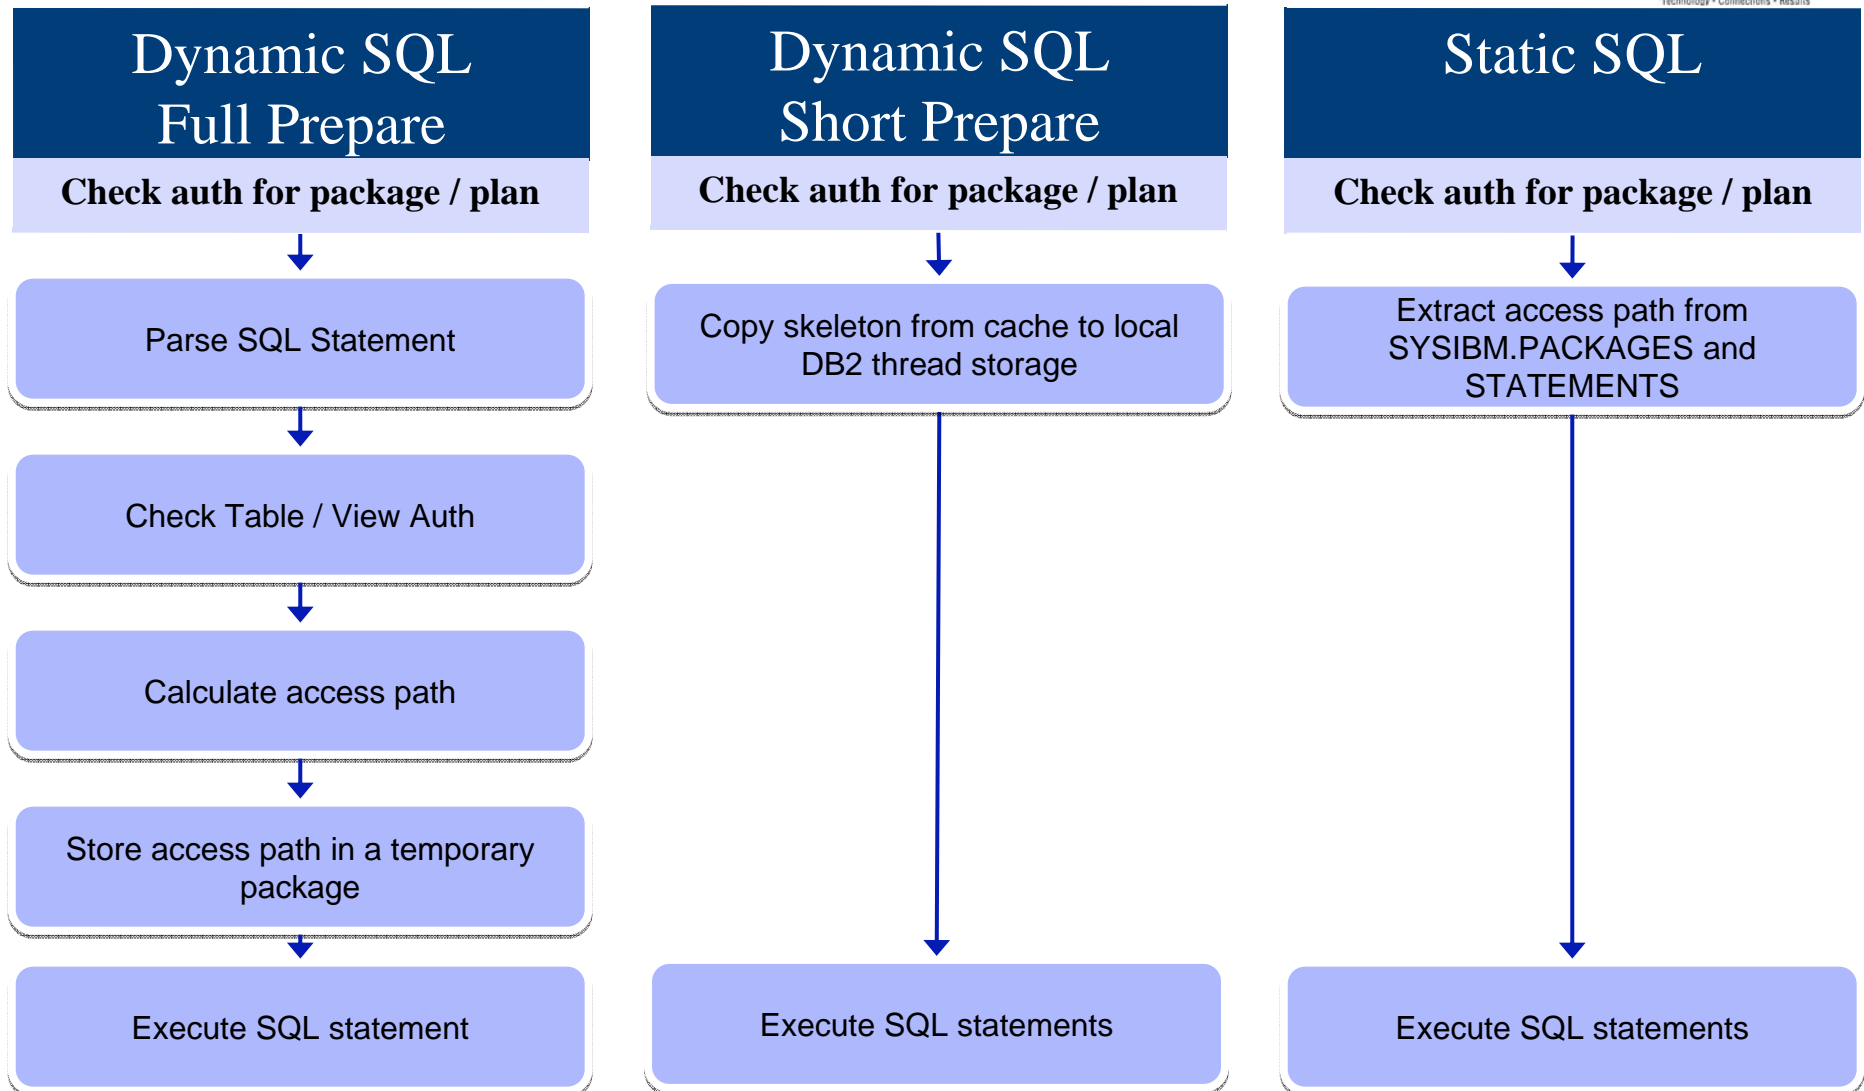
- Make response time predictable
  - Lock in the SQL access path pre-execution
- Limit user access and reduce SQL injection
  - Grant execute privileges on the query packages rather than access privileges on the table
- Accelerate problem resolution
  - Trace SQL execution to a specific package and the originating source
- Improve impact analysis and capacity planning
  - Visualize application SQL and correlation metadata
- Increase system capacity
  - Drive down DB cycles



**"The ability to use static SQL with pureQuery is huge. Recently, I worked with a client who could reduce CPU usage by 7 percent thanks to this one feature."**

**— David Beulke, Pragmatic Solutions Inc.**

# Dynamic vs. Static Execution



# Static Execution Advantages



Feature	Dynamic SQL (pureQuery, JDBC)	Static SQL (pureQuery, SQLJ)
Performance	Can approach static SQL performance with help from dynamic SQL cache. Cache misses are costly	All SQL parsing, catalog access, done at BIND time. Fully optimized during execution.
Access path reliability	Unpredictable – Any prepare can get a new access path as statistics or host variables change	Guaranteed – locked in at BIND time All SQL available ahead of time for analysis by EXPLAIN.
Authorization	Privileges handled at object level. All users or groups must have direct table privileges – Security exposure, and administrative burden	Privileges are package based. Only administrator needs table access. Users/Groups have execute authority. Prevent non-authorized SQL execution.
Monitoring, Problem determination	Database View is of the JDBC or CLI package – No easy distinction of where any SQL statement came from.	Package View of applications makes it simple to track back to the SQL statement location in the application
Capacity planning, Forecasting	Difficult to summarize performance data at program level.	Package Level Accounting gives program view of workload to aid accurate forecasting.
Tracking dependent objects	No record of which objects are referenced by a compiled SQL statement	Object dependencies registered in database catalog

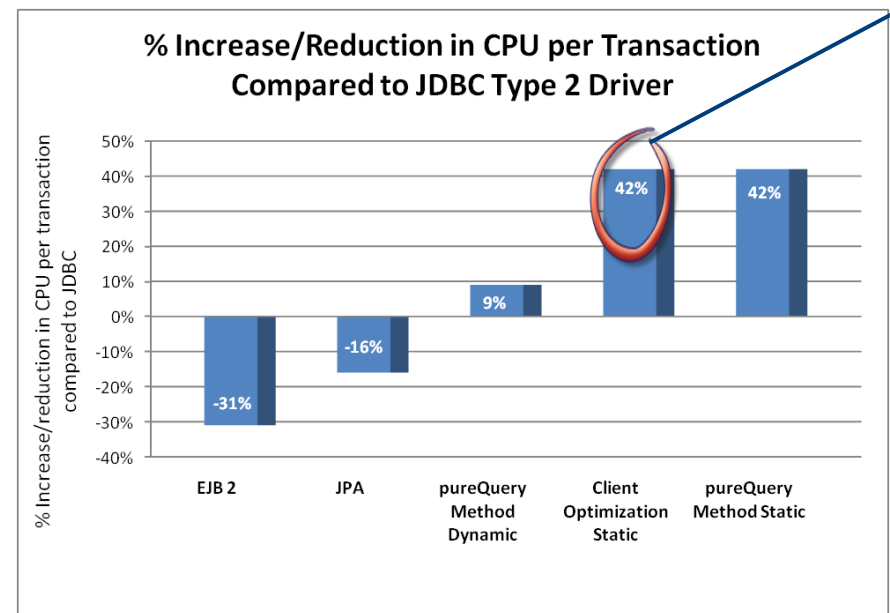
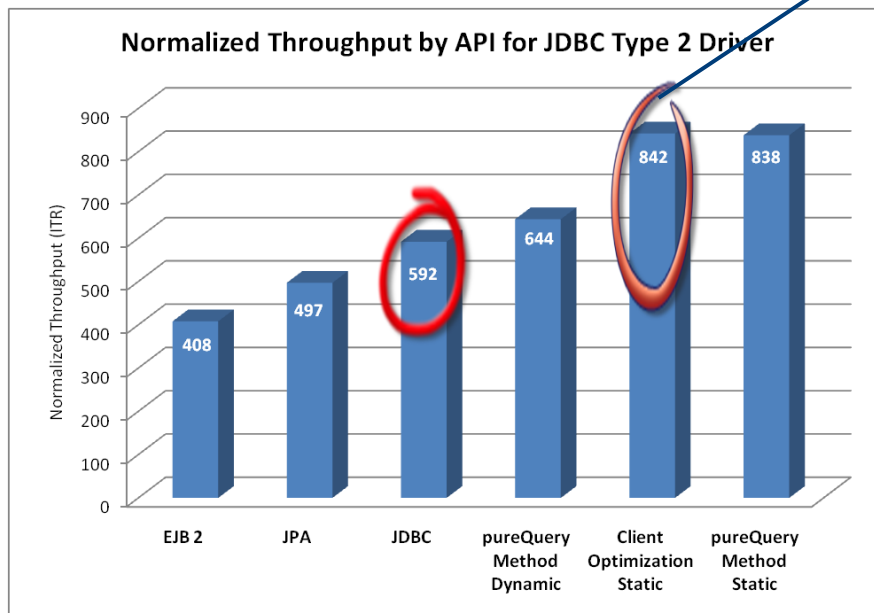


# Improving Throughput with pureQuery, a z/OS Example

- In-house testing shows **over 40% reduction** in CPU costs over dynamic JDBC using pureQuery and DB2 for z/OS (type-2)
  - Read [IBM Optim pureQuery Runtime for z/OS Performance](#)
  - IRWW – an OLTP workload, cache hit ratio between 70 and 85%

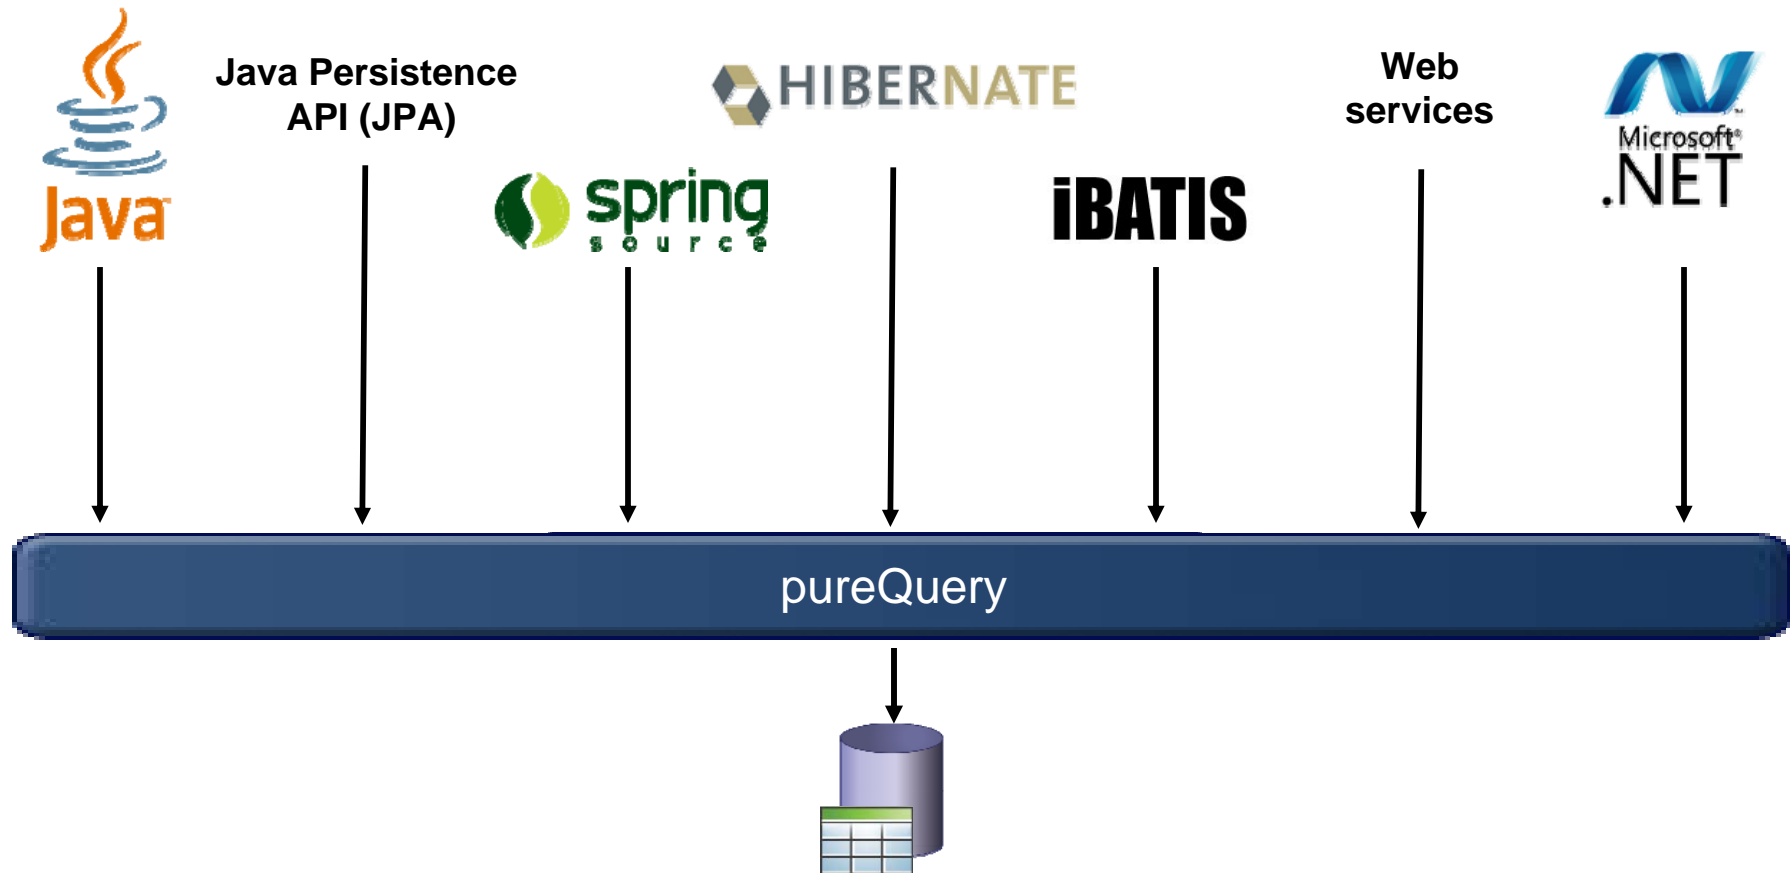
Preset application path never varies at runtime

Reduced CPU per transaction means more growth on existing resources



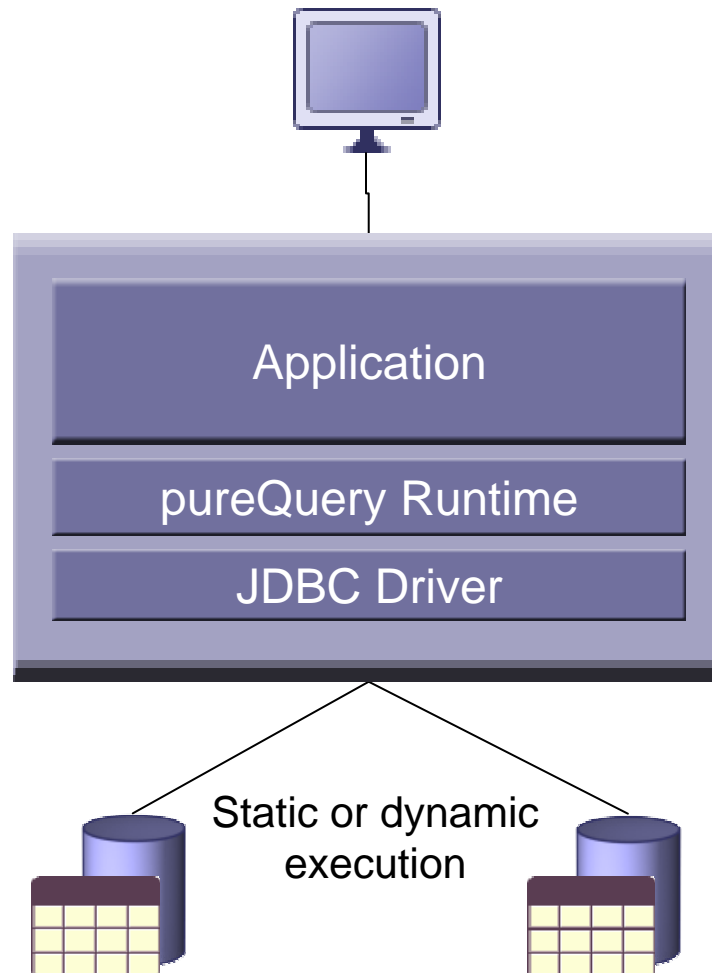
# Java Database Access and pureQuery

*Many on-ramps for new and existing applications*



**DB2 (z/OS, i & LUW), Informix, and Oracle now  
More coming**

# Deploying with pureQuery Runtime



## ***Application tier:***

- z/OS, Linux, UNIX, Windows

## ***Database tier:***

- DB2 for z/OS
- DB2 for i
- DB2 for Linux, UNIX, and Windows
- Informix Dynamic Server
- Oracle

# Unique Package Names Improves PD on DB2 for z/OS

- Most dynamic Java applications use packages SYSLNx00 making it hard to identify specific programs
- Unique package names link SQL to Java Beans, similar to CICS transaction names to programs.

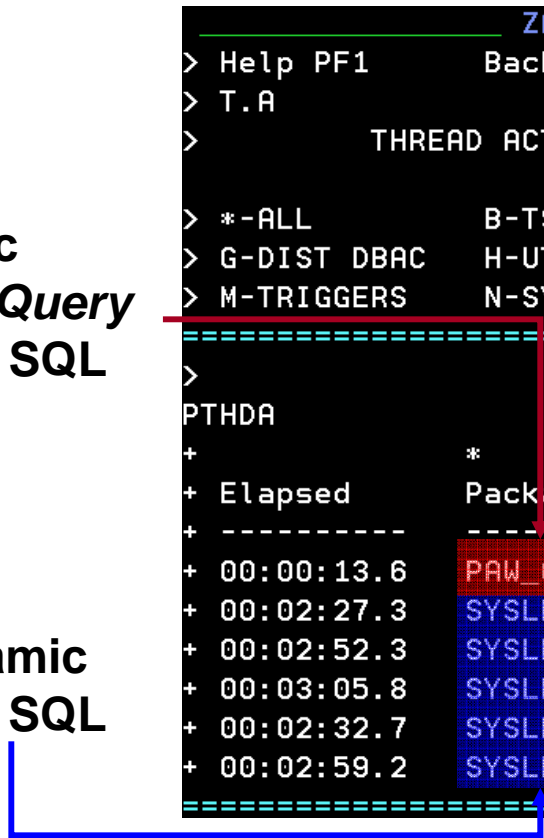
Static  
*pureQuery*  
Java SQL

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
=====

```

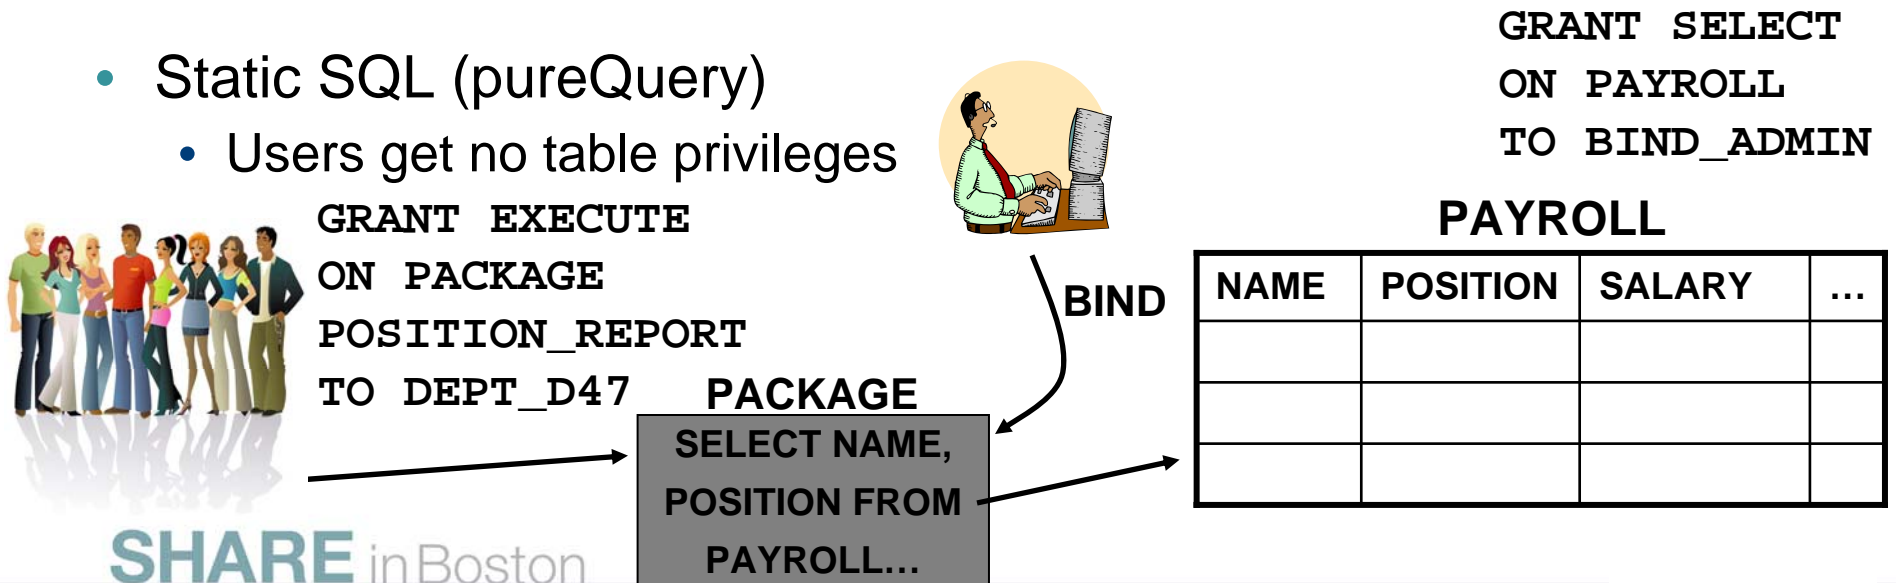


# Static SQL for security – Administering table privileges

- Dynamic SQL
  - Table privileges granted directly to users(groups)/secondary authids
    - Security exposure and administrative burden



- Static SQL (pureQuery)
  - Users get no table privileges





# Reduce Costs with zIIP and zAAP

- Using COBOL or SQL/PL stored procedures instead of executing SQL directly
  - Separation of application development from data access
  - DBAs retain better control over SQL including static execution
  - If not written in DB2 9 for z/OS's Native SQL/PL, **the stored procedure must use general purpose processors**
- pureQuery introduces alternatives
  - Create Java stored procedures to run on zAAP using the pureQuery runtime
    - DBAs retain control, Data Studio helps with development, pureQuery executes statically
  - Execute SQL directly from Java application or method to run on zIIP
    - Developers use Data Studio Developer to generate access layer with pureQuery, content assist helps with best practices and SQL validation, packages SQL for easy collaboration with DBA, pureQuery executes statically

# More Visibility and Control of Application SQL

- Capture SQL
- Share, review, and optimize SQL
- Revise and validate equivalency
- Bind for static execution or run dynamically
- Restrict SQL to eliminate SQL injection



IT PRO has been watching and charting the progress of what is one of the biggest and most high profile web security threats of this year - the SQL injection.

By Asavin Wattanajutra, 4 Aug 2008 at 11:55

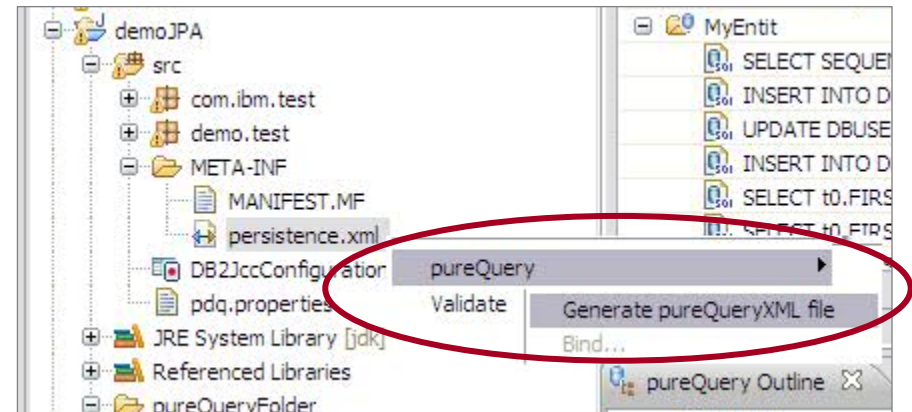


# Capture Application SQL: At Development or Later

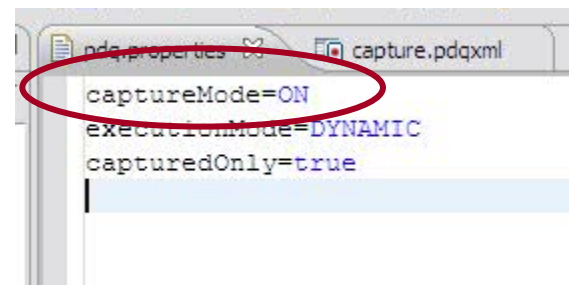


Three methods

1. Use pureQuery API
2. Use JPA and generate the pureQuery file
3. Set captureMode=ON and execute the program



**IBM Optim  
Development  
Studio**





# Visualize Application and SQL Metadata



- Review the captured SQL
- View metrics about execution frequency and duration
- Share captured SQL with DBA

**IBM Optim Development Studio**

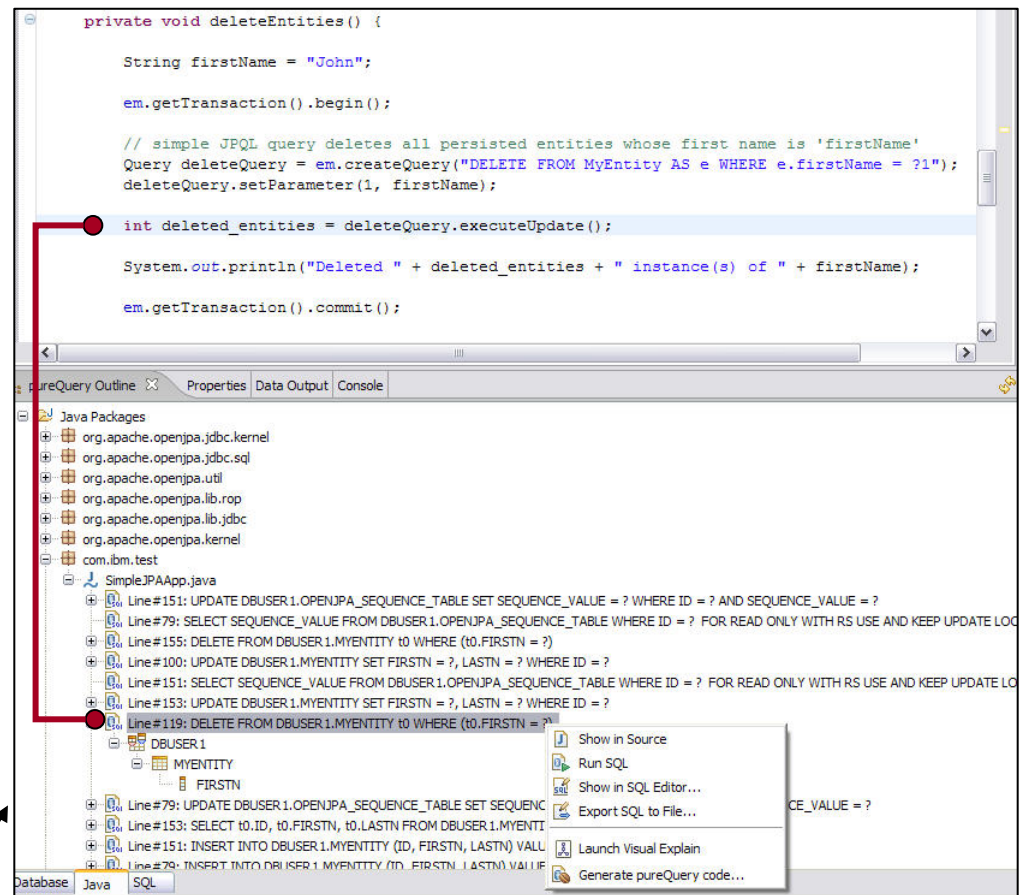
The screenshot displays the 'pureQuery Outline' window with a table of SQL execution metrics. The table includes columns for Execution Count, Max Time, Average, and Mir. The data is as follows:

SQL Statement	Execution Count	Max Time	Average	Mir
Line# 116: SELECT DEPTNO, DEPTNAME, MGRNO, LOCATION	49	12625627	5366204.755102041	41
Line# 123: getSql()	0	0	0.0	
Line# 108: SELECT DEPTNO, DEPTNAME, MGRNO, ADMRDEP	54	214289805	9491409.166666666	41
Line# 100: SELECT DEPTNO FROM DEPARTMENT	32	17681018	5702409.53125	41
Line# 124: SELECT DEPTNAME FROM DEPARTMENT	0	0	0.0	


# SQL Outline

## Speed up problem isolation for developers – even when using frameworks

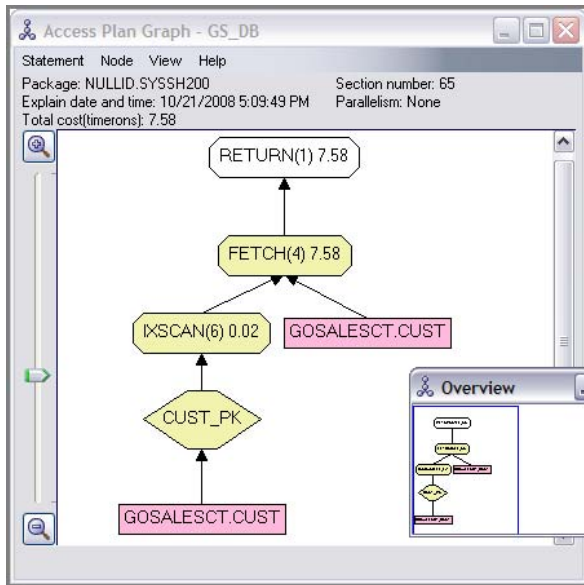
- Capture application-SQL-data object correlation (with or without the source code)
- Trace SQL statements to using code for faster problem isolation
- Enhance impact analysis identifying application code impacted due to database changes
- Answer “Where used” questions like “Where is this table used within the application?”
- Use with modern Java frameworks e.g. Hibernate, Spring, iBatis, OpenJPA



# Optimize SQL



- Launch Visual Explain



- Copy SQL to Optim Query Workload Tuner

Feature Details	Creator	Object Name	Columns	Estimated Disk Space
<input checked="" type="checkbox"/> LINEITEM <input checked="" type="checkbox"/> Index	DB2OE	LINEITEM_VIRT_IDX_1181...	I_RETURNFLAG(ASC) I_S...	717.08203125 M

```
CREATE INDEX 'DB2OE'.LINEITEM_VIRT_IDX_1181023618203 ON 'SYSADM'.LINEITEM (
  I_RETURNFLAG ASC, I_SUPPKEY ASC, I_RECEIPTDATE ASC, I_SHIPDATE ASC,
  I_SHIPMODE ASC, I_ORDERKEY ASC, I_PARTKEY ASC, I_LINENUMBER ASC,
  I_QUANTITY ASC, I_EXTENDEDPRICE ASC, I_DISCOUNT ASC, I_TAX ASC,
  I_LINESTATUS ASC, I_COMMITDATE ASC, I_SHIPINSTRUCT ASC, I_COMMENT ASC)
NOT PADDED FREEPAGE 0 PCTFREE 10;
```

# Query Tuning



# Introduction to query tuning

## Query performance causes application outage

- Customer example
  - Application service drops from 3 seconds to 5 minutes
  - Limited performance management expertise
  - Took 3 days to identify the problem query
  - 5 day application outage
  - Loss of confidence in IT team
  - Time and money in analysts
- Expert analysis revealed
  - Relevant statistics were not being collected
  - Wasted CPU resources due to default statistics collection
  - Index design was very poor
  - Rapid development practices drove untuned queries into production

**Poor data**

**Poor design**

**Wasted resources**

**Poor process**

# Introduction to query tuning

## Why is query tuning such a pain?

- Development
  - Lack of query tuning skills
  - Use of frameworks that generate SQL
  - Inadequate testing environments to drive the workload and data scale
  - Disconnect between application and data groups
- Production
  - Performance problems appear without warning
  - Aggregate information from multiple places
  - Complexity of analysis requires specialized skills
  - Need to consider entire workloads, not just single queries

# Introduction to query tuning

- Why do we have to perform query and system tuning? Isn't this why I have DB2?
  - DB2 is powerful, but it's not all-knowing
    - DB2 doesn't know what queries are going to be submitted
    - DB2 doesn't know object sizes or cardinality unless statistics are gathered
    - DB2 doesn't know about cross-table correlation
    - DB2 doesn't know how many rows will be fetched, unless told
  - The biggest reason for a non-optimal access path is because the desired statistics are not collected
- Defining the task "query tuning" -- Developer versus DBA tasks
  - Developer tasks
    - Follow shop standards
    - Avoid Stage 2 (and Stage 3) predicates
  - DBA tasks
    - Design indexing scheme
    - Collect and maintain needed statistics
    - Analyze access paths and influence them when needed

# Introduction to query tuning

- Techniques used
  - Rules-based analysis of SQL statement text
  - Invoking EXPLAIN to understand the access path and cost of the query blocks
  - Invoking EXPLAIN to understand what statistics are interesting to DB2
  - Simulation of index schemes
- Other issues
  - Sometimes the SQL is generated and cannot be reviewed





# IBM Tooling that is available to assist with this

- IBM Tooling that is available to assist with this
  - DB2 SQL Performance Analyzer (SQL PA)
  - Optim Query Workload Tuner



# SQL Performance Analyzer

- ISPF and batch
- Enhanced Explain
  - Compare access paths
  - What if
  - Retro explain
  - Migrate DB2 Statistics
- Provides advice on each SQL statement (helps users write better SQL)
  - Warnings, alerts, guidelines and recommendations
  - Performance notes
- Forecasts SQL performance (response/wait times, CPU, I/O counts)
- Forecasts the cost of the query, in terms of
  - Charge back (monetary, in national currency)
  - QUNITS (query service units)
- Acts as a governor for Static or Dynamic SQL
  - in QMF, as a preemptive exit
  - in DRDA and IMS or CICS, via a Stored Procedure call

Suitable for SQL Develop, Tuning and Control



# Query Workload Tuner

- Eclipse-based interface
- Gathers SQL from various sources (catalog, statement cache, captured file)
- Visualize queries and costs (query annotation and visual explain)
- View DB2 Optimizer's query transformations
- Analyze one query or a whole workload
  - Access path
  - Statistics
  - Index
- Validate changes

Suitable for SQL Tuning, Statistics Collection and Index Design by DBAs

# Streamlined Analysis



## Define or select workload

Select a workload and then click an item in the Open menu to view information about the workload.

New Workload Open Refresh Remove Tools

Wizard... Statement Cache... Catalog... QMF... QMF HPO... File... Category... Other Workloads... Import... Export...

Summary Status	Owner
ENABLED/STARTED	SYSADM
ENABLED/STARTED	SYSADM
ANALYZING	B3OSC07
ANALYZING	B3OSC12
ANALYZING	B3OSC12

## Execute advisors

Workload Statements

Immediately capture statements or multiple sources to this workload, launch workload advisors, use tools capture, consolidation, and analysis.

Capture Workload Tools Query Tools Schedule History

Filter name: DEF

Run All Advisors  
Run Workload Statistics Advisor  
Run Workload Index Advisor  
Run Workload Query Advisor

Generate Workload  
Workload Report  
Workload Environment Capture

Execution Count		Average Elapsed Time (sec)	Accumulated
5		21.151768	0.0696
5		23.351143	2.2789
5		18.033714	10.322
5	CACHE 90.168571	5.907151	0.0459
5	CACHE 29.535755	20.778751	5.5212

## Drill into advice

Advisor Recommendation Overview

These recommendations are part of a sample project. [Learn More](#)

Advisor	Priority	Description
Recommendations		
Statistics Advisor	HIGH	Repair statistics problems for this
Query Advisor	MEDIUM	Consider adding join predicates between
Query Advisor	MEDIUM	Consider adding the following predicate
Access Path Advisor	LOW	The DSN8910.EMP table is accessed by
Index Advisor	LOW	Index recommendations found.
Statistics Advisor	MAINTENANCE	Gather and recollect all of relevant stat

Context

Query 1 Initial Analysis

## Validate improvement

Workloads List

Select a workload and then click an item in the Open menu to view information about the workload. To create a workload, click

New Workload Open Refresh Remove Tools

Name	Summary Status	Owner	Execution Time
WorkloadWithTypicalStats	ANALYZING	B3OSC12	CPU time: 97.32 (second...
WorkloadTunedWithStatsAdvisor	ANALYZING	B3OSC12	CPU time: 53.19 (second...
WorkloadTunedWithIndexAdvisor	ANALYZING	B3OSC07	CPU time: 40.67 (second...
AbsoluteCPUTimeExceptionMonitor	ENABLED/STARTED	SYSADM	N/A
NormalMonitor	ENABLED/STARTED	SYSADM	N/A



# Visualize Queries and Costs to Speed Analysis

Formatted Query	Annotation	Additional Information
<pre> SELECT A.EMPNO       , A.FIRSTNME       , A.LASTNAME       , A.JOB       , A.SALARY       , A.BONUS       , A.COMM       , B.LOCATION       , C.PROJNAME FROM DSN8910.DEPT AS B       , DSN8910.EMP AS A       , DSN8910.EPROJ AS C WHERE ( A.EMPNO IN ( SELECT DSN8910.DEPT.MGRNO                      FROM DSN8910.DEPT                      WHERE DSN8910.DEPT.MGRNO IS NOT NULL                    )       AND A.WORKDEPT = B.DEPTNO       AND B.DEPTNO = C.DEPTNO       ) ORDER BY A.EMPNO ASC       , A.FIRSTNME ASC       , A.LASTNAME ASC           </pre>	<p>Easily see tables, sections, join predicates, etc.</p>	<p>Examine table statistics and additional information</p>
	CARDF=14 QUALIFIED_ROWS=	
	CARDF=42 QUALIFIED_ROWS=	
	CARDF=(missing) QUALIFIED_ROWS=	
	COLCARDF=42 MAX_FREQ=(r	
	CARDF=14 QUALIFIED_ROWS=	
	COLCARDF=9 MAX_FREQ=42	DSN8910.DEPT.MGRNO contain(s) skewed data
	COLCARDF=8/14 MAX_FREQ=:	DSN8910.EMP.WORKDEPT contain(s) skewed data
	COLCARDF=14/(missing) MAX_FREQ=	

- Accelerate analysis, reduce downtime
  - Spot human errors
  - Identify where filtering should occur



# View Optimizer Transformations



Original Transformed

Annotation to display: All

Expand All Collapse All Customize Save Print

Formatted Query

```
SELECT A.EMPNO
      , A.FIRSTNAME
      , A.LASTNAME
      , A.JOB
      , A.SALARY
      , A.BONUS
      , A.COMM
      , B.LOCATION
      , C.PROJNAME
FROM DSN8910.DEPT AS B
      , DSN8910.EMP AS A
      , DSN8910.EPROJ AS C
WHERE ( A.EMPNO IN ( SELECT DSN8910.DEPT.MGRNO
                    FROM DSN8910.DEPT
                    WHERE DSN8910.DEPT.MGRNO IS NOT NULL
                  )
      AND A.WORKDEPT = B.DEPTNO
      AND B.DEPTNO = C.DEPTNO
)
ORDER BY A.EMPNO ASC
      , A.FIRSTNAME ASC
      , A.LASTNAME ASC
```



Original Transformed

Annotation to display: All

Expand All Collapse All Customize Save Print Clear Highlights

Formatted Query

```
SELECT A.EMPNO
FROM DSN8910.DEPT AS B
      , DSN8910.EMP AS A
      , DSN8910.EPROJ AS C
WHERE ( A.WORKDEPT = B.DEPTNO
      AND A.WORKDEPT = C.DEPTNO
      AND B.DEPTNO = C.DEPTNO
      AND A.EMPNO = SYSADM."DSNWFQB(02)".MGRNO
      DB2 creates a virtual table, SYSADM.DSNWFQB(02), to process the following correlated subquery
          SELECT DSN8910.DEPT.MGRNO
          FROM DSN8910.DEPT
          WHERE ( DSN8910.DEPT.MGRNO IS NOT NULL
                AND A.EMPNO = DSN8910.DEPT.MGRNO
              )
      End of the subquery for the SYSADM.DSNWFQB(02) virtual table.
)
ORDER BY A.EMPNO ASC
      , A.FIRSTNAME ASC
      , A.LASTNAME ASC
```

- Accelerate analysis, reduce downtime
- Spot transformations which will occur



# Execute Advisors

Workload Statements

Immediately capture statements or multiple sources to this workload, launch workload advisors, use tools capture, consolidation, and analysis.

Capture Workload Tools Query Tools Schedule History

Filter name: DEF

< Previous View Cus

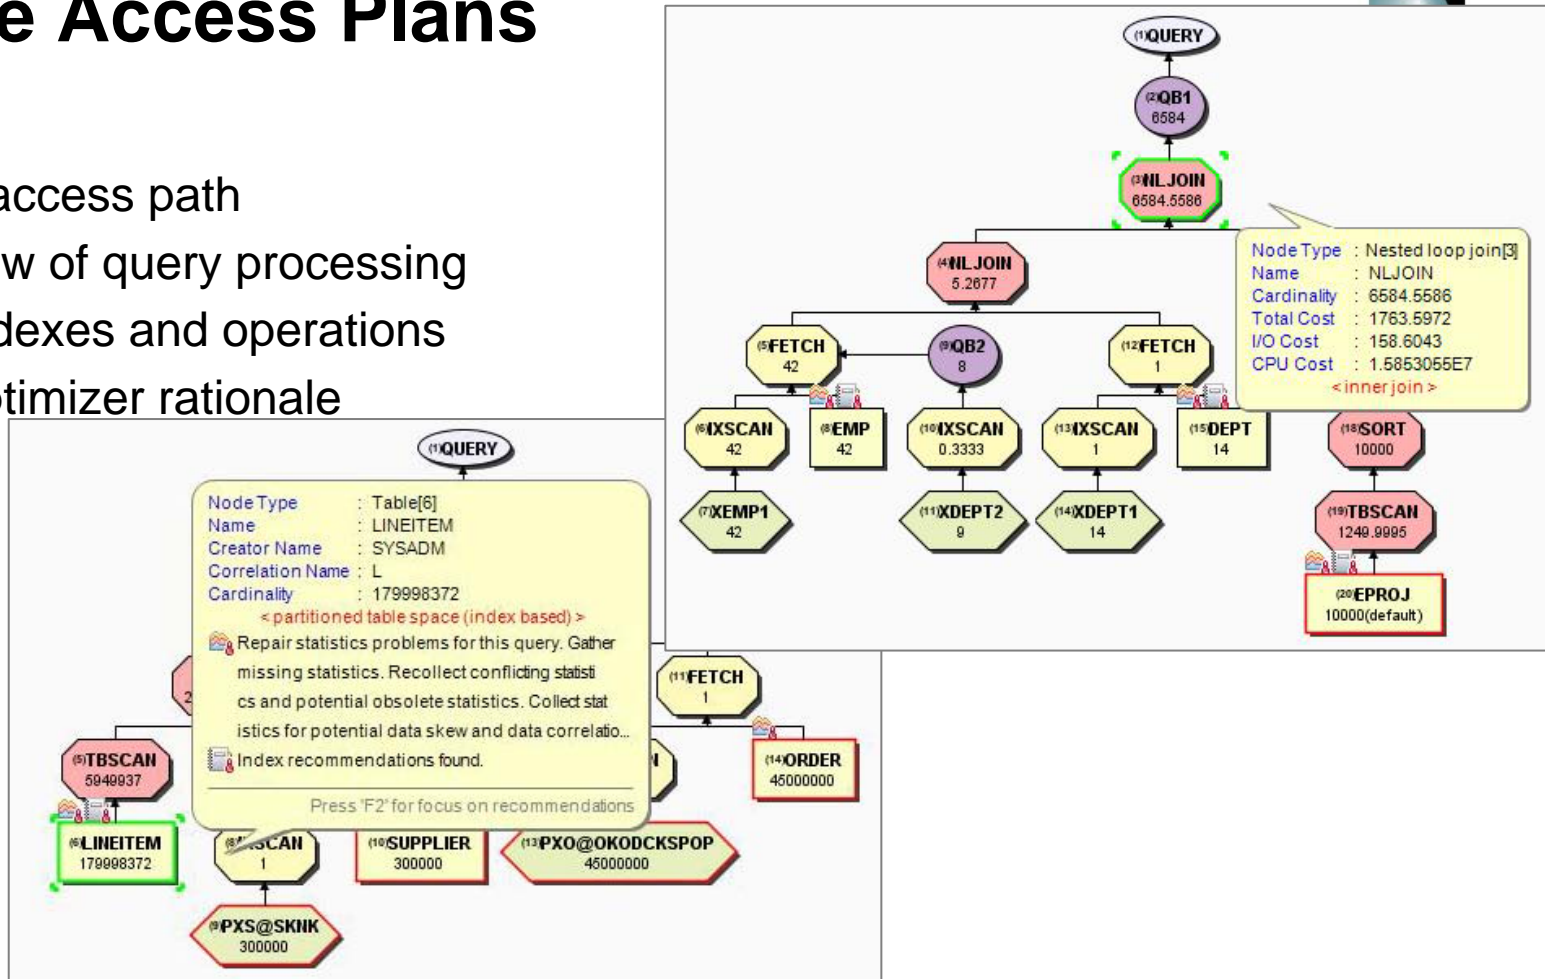
- Run All Advisors
- Run Workload Statistics Advisor
- Run Workload Index Advisor
- Run Workload Query Advisor
- Generate Workload
- Workload Report
- Workload Environment Capture

Execution Count	Cache	Average Elapsed Time (sec)	Accum
5		21.151768	0.0696
5		23.351143	2.2789
5	CACHE 90.168571	18.033714	10.322
5	CACHE 29.535755	5.907151	0.0459
5	CACHE 103.893761	20.778751	5.5212
5	CACHE 498.326721	99.665344	0.5974
5	CACHE 67.203255	13.440651	2.4417
9	CACHE 83.786201	9.309578	0.3156
9	CACHE 170.230835	18.914537	0.1205
11	CACHE 193.929184	17.629927	14.428
11	CACHE 248.512375	22.592033	5.0555
11	CACHE 208.377579	18.943417	0.1524
11	CACHE 208.310974	18.937361	0.0849
11	CACHE 37.513599	3.410327	2.4663
11	CACHE 205.77478	18.706799	0.0596
13	CACHE 244.390244	18.79925	1.2846
13	CACHE 353.33551	27.179655	6.4435
16	CACHE 84.005241	5.250328	0.0872
16	CACHE 15.534015	0.970876	0.3689

- **Statistics**
  - Get recommendations on the best statistics to capture to influence access path selection
- **Query**
  - Get recommendations regarding how to rewrite the query for better efficiency
- **Index**
  - Get recommendations on additional indexes that can reduce database scans

# Analyze Access Plans

- Visualize access path
  - See flow of query processing
  - See indexes and operations
  - See optimizer rationale



- Assess access path stability to reduce risk of performance regression
  - 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 Statistics Quality and Collection



- Provides advice on
  - Missing statistics
  - Conflicting statistics
  - Out-of-date statistics
- Results
  - Accurate estimated costs
  - Better query performance
  - Less CPU consumption
  - Improved maintenance window throughput

*“Half of access path PMRs could be resolved by statistics advisor before calling IBM support.”*  
– IBM Support

SHARE in Boston

The screenshot shows the 'Statistics Recommendation Detail' window. It is divided into several sections:

- Recommendations:** Contains 'RUNSTATS Control Statements' and a list of table and column statistics recommendations, such as 'RUNSTATS INDEX(SYSADM.PXS@SKNK FREQVAL NUMCOLS 1 COUNT 15)'. A callout bubble points to this section with the text: **Generates RUNSTATS control statements**.
- Table, index, column, and column group details:** Shows details for two objects: 'O\_ORDERKEY' and 'O\_ORDERPRIORITY'. For 'O\_ORDERKEY', the 'Uniform statistics status' is 'conflicting' and the 'Frequency statistics status' is 'conflicting'. For 'O\_ORDERPRIORITY', the 'Uniform statistics status' is 'OK' and the 'Frequency statistics status' is 'missing'. A callout bubble points to the 'conflicting' status of 'O\_ORDERKEY' with the text: **Indicates conflicting and missing statistics**.
- Conflicts detail:** Provides an explanation for the conflict on 'O\_ORDERKEY', stating: 'One of the frequency records (-1.0) of the L\_ORDERKEY column group is out of range [0,1] Tolerance: 0.0010'. A callout bubble points to this section with the text: **Conflicting statistics explanation**.

# Improve Query Design

Guard against errors and oversights:  
Further constrain query, increase index utilization, and reduce data reads

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	0

View statements that meet the following:

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

Restore Defaults Save as Defaults

Statements Index Advisor Query Advisor Summary Advisors

View analysis summary

Filter recommendations by severity

Query Advisor checks for

- Missing join predicate for referential constraint
- Predicates that can be rewritten as indexable
- Stage 2 predicates that can be rewritten as stage 1 predicates

Query Recommendation Detail

SQL Text

```
SELECT A.EMPNO
, A.FIRSTNAME
, A.LASTNAME
, A.JOB
, A.SALARY
, A.BONUS
, A.COMM
, B.LOCATION
, C.PROJNAME
FROM DSN8910.DEPT AS B
, DSN8910.EMP AS A
, DSN8910.EPROJ AS C
WHERE ( A.EMPNO IN ( SELECT DSN8910.DEPT.MGRNO
FROM DSN8910.DEPT
WHERE DSN8910.DEPT.MGRNO IS NOT NULL
)
```

Highlights relevant components of the query

Selected Recommendation:

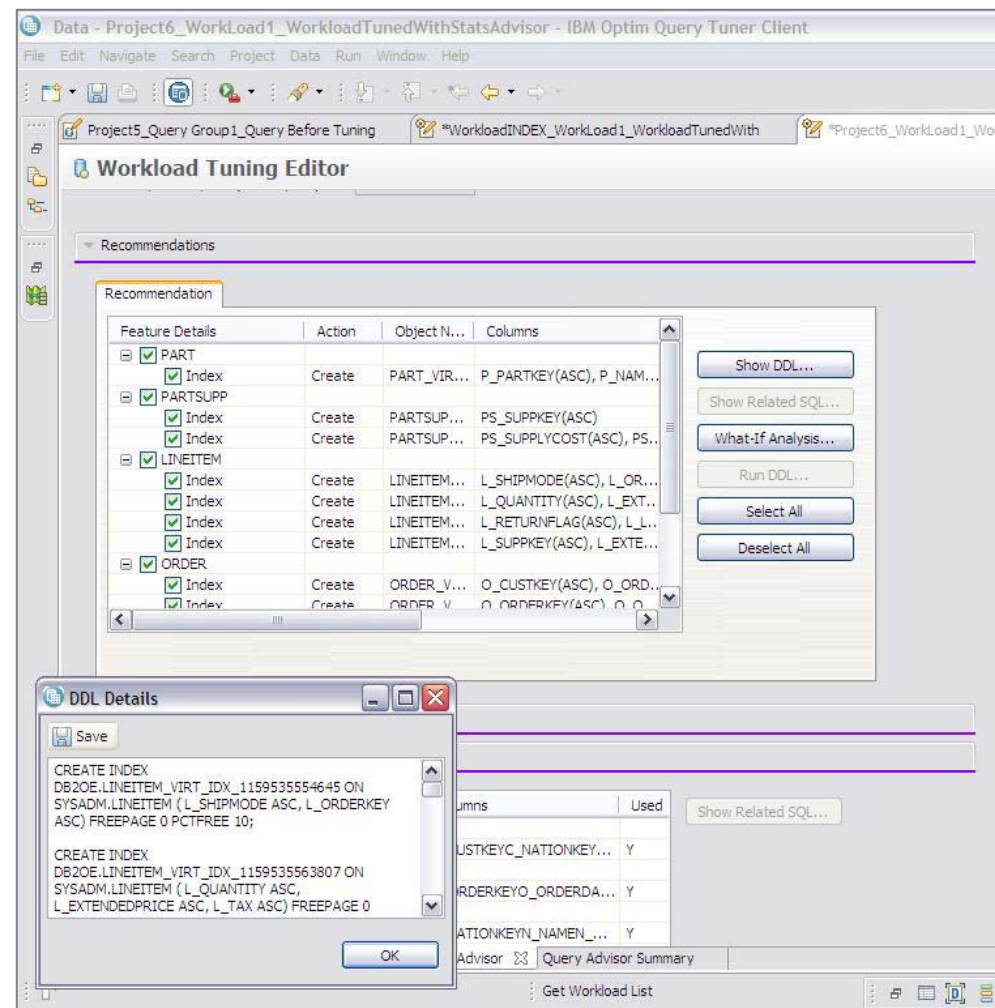
Description
Consider adding join predicates between columns EMPNO in table DSN8910.EMP and columns MGRNO in table DSN8910.DEPT which use the referential constraints between table DSN8910.EMP and table DSN8910.DEPT to avoid a potentially costly Cartesian join. Check the explanation for this warning for more details about possible impact and examples.
If a referential constraint is defined between corresponding join predicates that map exact RI1 is defined between table T1 and T2, in w that an SQL statement joins the two tables as
<pre>SELECT T1.C2, T2.C2 FROM T1, T2 WHERE T1.C2 = :charHV3 AND T2.C2 = :charHV4</pre> <p>As written, the SQL statement does not include between the two tables. The result is likely to rewriting the SQL statement as follows:</p> <pre>SELECT T1.C2, T2.C2 FROM T1, T2 WHERE T1.C2 = :charHV3</pre>

Recommendation and rationale



# Indexing Advice to Improve Query Efficiency

- Improve query efficiency
  - Indexing foreign keys 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
  - Enables what-if analysis
  - Provides DDL to create indexes
  - Run immediately or save
  
- Test before deployment
  - Utilize virtual index capabilities built into the DB2 engine



The screenshot shows the IBM Optim Query Tuner Client interface. The main window is titled "Workload Tuning Editor" and displays a table of recommendations for creating indexes. A "DDL Details" dialog box is open, showing the SQL DDL for two indexes.

Feature Details	Action	Object N...	Columns
<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> PART</li> <li><input checked="" type="checkbox"/> Index</li> </ul>	Create	PART_VIR...	P_PARTKEY(ASC), P_NAM...
<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> PARTSUPP</li> <li><input checked="" type="checkbox"/> Index</li> <li><input checked="" type="checkbox"/> Index</li> </ul>	Create	PARTSUP...	PS_SUPPKEY(ASC)
<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> LINEITEM</li> <li><input checked="" type="checkbox"/> Index</li> <li><input checked="" type="checkbox"/> Index</li> <li><input checked="" type="checkbox"/> Index</li> <li><input checked="" type="checkbox"/> Index</li> </ul>	Create	LINEITEM...	L_SHIPMODE(ASC), L_OR...
<ul style="list-style-type: none"> <li><input checked="" type="checkbox"/> ORDER</li> <li><input checked="" type="checkbox"/> Index</li> <li><input checked="" type="checkbox"/> Index</li> </ul>	Create	ORDER_V...	O_CUSTKEY(ASC), O_ORD...

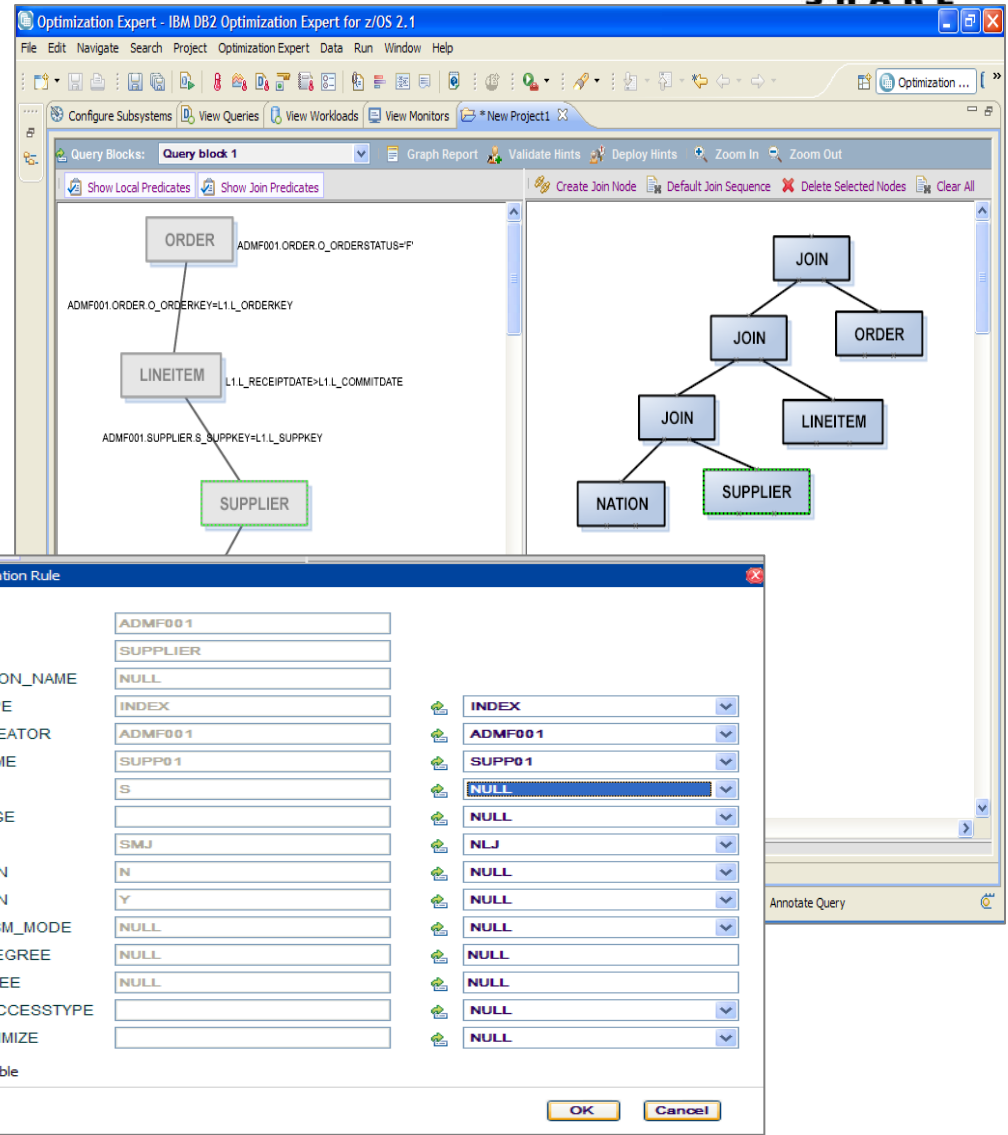
```

CREATE INDEX
DB20E.LINEITEM_VIRT_IDX_1159535554645 ON
SYSADM.LINEITEM (L_SHIPMODE ASC, L_ORDERKEY
ASC) FREEPAGE 0 PCTFREE 10;

CREATE INDEX
DB20E.LINEITEM_VIRT_IDX_1159535563807 ON
SYSADM.LINEITEM (L_QUANTITY ASC,
L_EXTENDEDPRICE ASC, L_TAX ASC) FREEPAGE 0
  
```

# Visual Plan Hints for Experienced DBAs

- Why hints?
  - Version to version migration
  - Assumptions made by optimizer when SQL contains parameter markers or host variables
  
- Reduce errors
  - Hint editor
  - Hint validation
  - Hint deployment



The screenshot displays the IBM DB2 Optimization Expert interface. The main window shows a visual query plan for 'Query block 1'. The plan consists of several nodes: 'ORDER' (ADMFO01 ORDER\_O\_ORDERSTATUS='F'), 'LINEITEM' (ADMFO01.ORDER\_O\_ORDERKEY=L1.L\_ORDERKEY), and 'SUPPLIER' (ADMFO01.SUPPLIER\_S\_SUPPKEY=L1.L\_SUPPKEY). The 'SUPPLIER' node is highlighted with a green dashed border. To the right, a hierarchical tree view shows the join structure: 'JOIN' (root) connects to 'JOIN' and 'ORDER'; the lower 'JOIN' connects to 'JOIN' and 'LINEITEM'; the lowest 'JOIN' connects to 'NATION' and 'SUPPLIER'. The 'SUPPLIER' node in this tree is also highlighted with a green dashed border.

In the foreground, the 'Hint Customization Rule' dialog is open. It contains the following fields:

CREATOR	ADMFO01
TNAME	SUPPLIER
CORRELATION_NAME	NULL
ACCESSTYPE	INDEX
ACCESSCREATOR	ADMFO01
ACCESSNAME	SUPP01
PREFETCH	S
PAGE_RANGE	
METHOD	SMJ
SORTN_JOIN	N
SORTC_JOIN	Y
PARALLELISM_MODE	NULL
ACCESS_DEGREE	NULL
JOIN_DEGREE	NULL
PRIMARY_ACCESTYPE	
WHEN_OPTIMIZE	

At the bottom of the dialog, there is a checkbox for 'Leading table' and 'OK' and 'Cancel' buttons.



# Query tuning tools for z/OS – Unique benefits

## Optim Query Workload Tuner

- **Better support for modern dev environment and dynamic SQL**
- **Eclipse-based, GUI and stored procedure-driven**
- **Integrates with OM/PE and QM GUI interfaces, Data Studio, Optim**
- **Provides virtual index capability for single queries and across workloads**
- **Recommends rewriting queries only when improvements result**
- **Has limited support to create statistics in Query Environment Capture and Workload Environment Capture service support**
- **Strategic investment spans databases**

## DB2 SQL Performance Analyzer

- **Better support for classic z/OS env, DBRMs, batch analysis, and z/OS libraries**
- **ISPF-based, Command and JCL- driven**
- **Integrates with OM/PE VTAM and QM 3270 interfaces**
- **Creates real indexes for “WHAT-IF” scenarios for single queries**
- **Offers a best practices Query Advisor with approx. 150 rules**
- **Has full support for cloning statistics**
- **Continued z/OS-specific investment**





# Optimization Service Center (the fate of)

- Current OSC remains supported through DB2 9
- OSC functionality will be split among the following offerings:
  - Data Studio (no charge)
  - Query Tuner (single query tuning capabilities already overlap with Q[W]T today)
  - Base function in DB2 for z/OS customers
    - SQL environment capture
    - ~~Profile monitoring~~

# Contents of Eclipse-based Query Tuning offerings

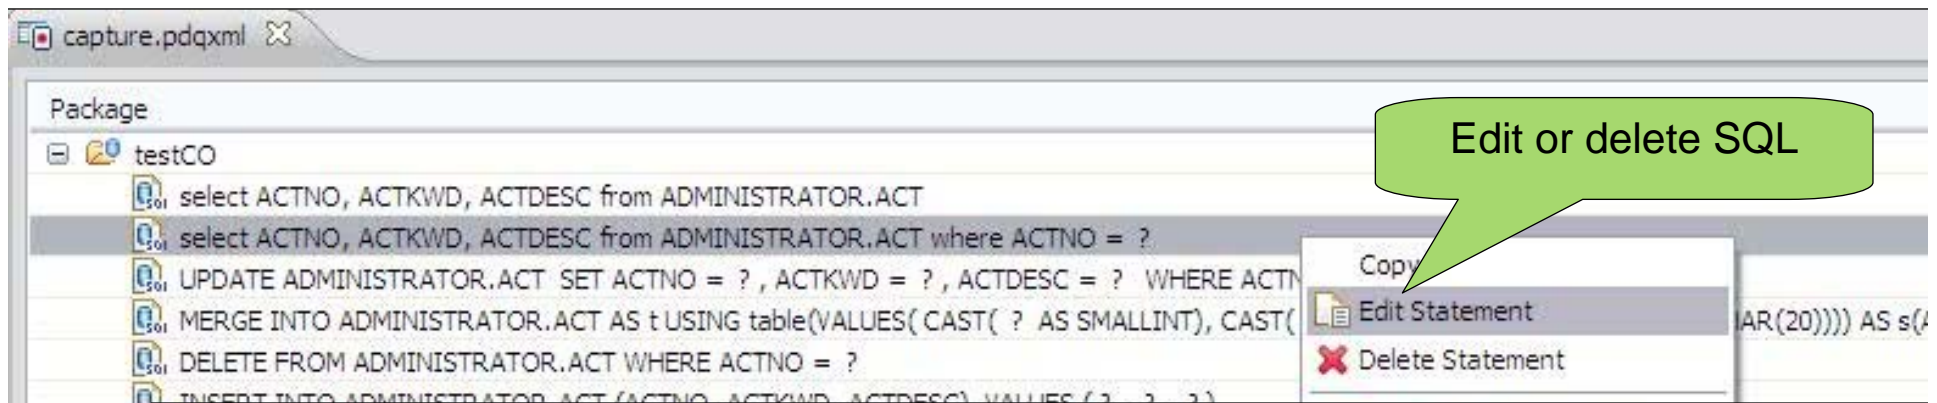


	Data Studio	Optim Query Tuner for z/OS	Optim Query Workload Tuner for z/OS
Queries from all sources	✓	✓	✓
Reports	✓	✓	✓
Query Formatter	✓	✓	✓
Access Plan Graph	✓	✓	✓
Query Statistics Advisor	✓	✓	✓
Query Annotation		✓	✓
Visual Plan Hint		✓	✓
Query Index Advisor		✓	✓
Query Advisor		✓	✓
Access Path Advisor		✓	✓
Workload Statistics Advisor			✓
Workload Index Advisor			✓
Workload Query Advisor			✓

# End of query tuning



# Revise SQL Without Modifying the Application



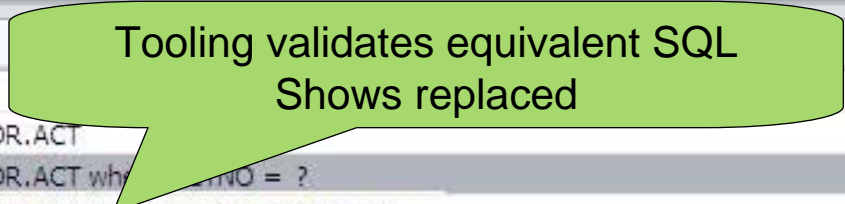
Package

- testCO
  - select ACTNO, ACTKWD, ACTDESC from ADMINISTRATOR.ACT
  - select ACTNO, ACTKWD, ACTDESC from ADMINISTRATOR.ACT where ACTNO = ?
  - UPDATE ADMINISTRATOR.ACT SET ACTNO = ? , ACTKWD = ? , ACTDESC = ? WHERE ACTNO = ?
  - MERGE INTO ADMINISTRATOR.ACT AS t USING table(VALUEs(CAST( ? AS SMALLINT), CAST( ? AS VARCHAR(20)))) AS s(t)
  - DELETE FROM ADMINISTRATOR.ACT WHERE ACTNO = ?
  - INSERT INTO ADMINISTRATOR.ACT (ACTNO, ACTKWD, ACTDESC) VALUES ( ? , ? , ? )

Context Menu:

- Copy
- Edit Statement
- Delete Statement

Callout: Edit or delete SQL



Package

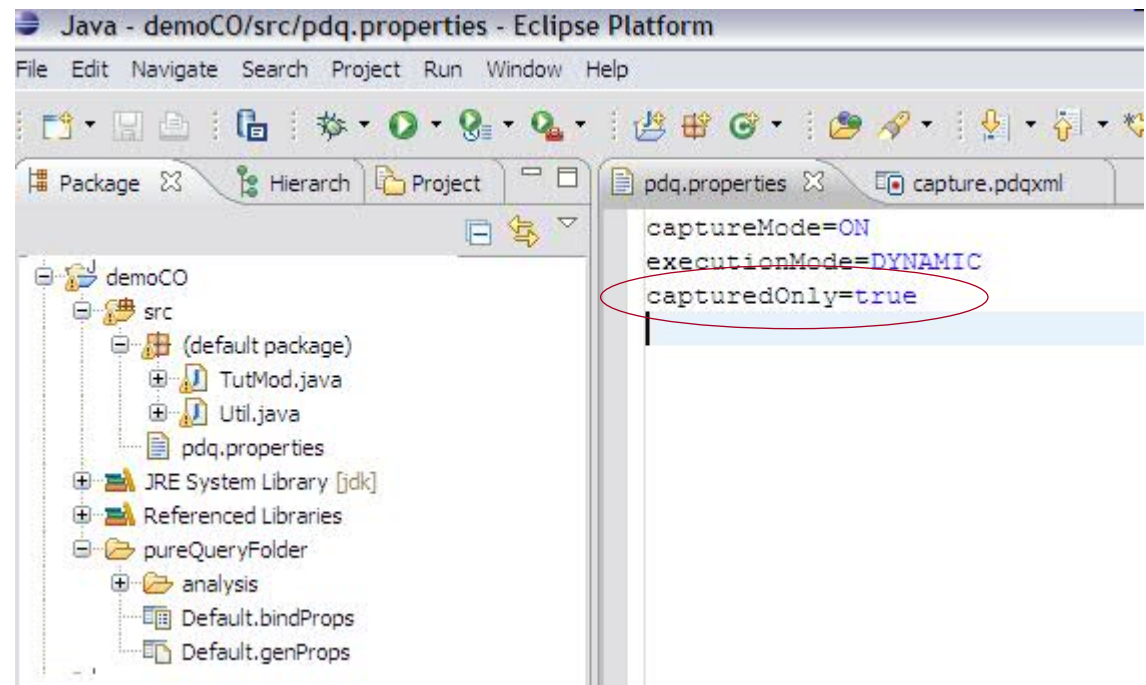
- testCO
  - select ACTNO, ACTKWD, ACTDESC from ADMINISTRATOR.ACT
  - select ACTNO, ACTKWD, ACTDESC from ADMINISTRATOR.ACT where ACTNO = ?
  - select ACTNO, ACTKWD, ACTDESC from ADMINISTRATOR.ACT where ACTNO = ?
  - UPDATE ADMINISTRATOR.ACT SET ACTNO = ? , ACTKWD = ? , ACTDESC = ? WHERE ACTNO = ?
  - MERGE INTO ADMINISTRATOR.ACT AS t USING table(VALUEs(CAST( ? AS SMALLINT), CAST( ? AS CHAR(6)), CAST( ? AS VARCHAR(20)))) AS s(t)
  - DELETE FROM ADMINISTRATOR.ACT WHERE ACTNO = ?
  - INSERT INTO ADMINISTRATOR.ACT (ACTNO, ACTKWD, ACTDESC) VALUES ( ? , ? , ? )

Callout: Tooling validates equivalent SQL Shows replaced

# Eliminate SQL Injection



- Restrict SQL execution to only those statements captured
- Set capturedOnly=true in pdq.properties
- pureQuery Runtime looks for it in the classpath



```
captureMode=ON
executionMode=DYNAMIC
capturedOnly=true
```

# Optimize for WebSphere and DB2 with pureQuery

## Capture metadata from existing applications

- Capture from JPA without executing
- Derive performance, costs, security and manageability value

## Jump start application design

- Generate SQL and Code from Database Objects
- Setup basic DAO Pattern

## Enhance development productivity

- Code generation, content assist
- Database aware, Java SQL Editor

## Reduce HW and SW costs

- Up to 42% lower CPU/Trans
- Move workload to zIIP and zAAP

## Replace SQL without changing the source

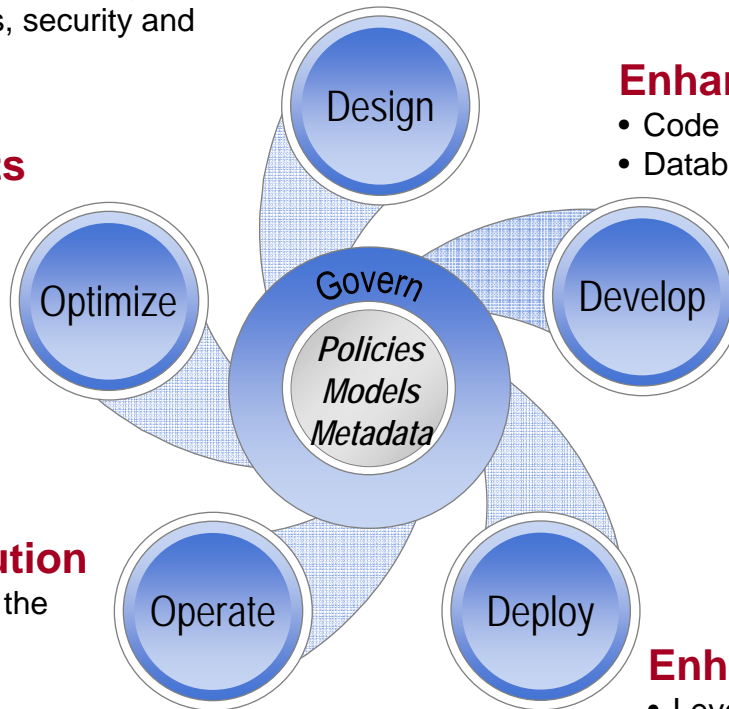
- Editor validates equivalency

## Speed up problem resolution

- Trace SQL back to line of code in the application

## Prevent SQL Injection

- Lock down SQL for dynamic or static execution



## Simplify impact analysis

- Categorize by Java, SQL, Database, Packages, track back to line of code

## Focus tuning efforts

- Find and sort by query elapsed time from Java

## Enhance performance

- Leverage best practices, automatically for JPA
- Use static execution, automatically for JPA
- Lock in access plans for consistent performance

## Reduce security exposure

- Grant access to queries, not tables



# Optim Development Studio and pureQuery Runtime



Country/region [select]

IBM

Home Solutions Services Products Support & downloads My IBM

developerWorks > Information Management > Downloads >

## Trial: IBM Optim Development Studio and Optim pureQuery Runtime

Learn Try Buy Support

**Download Optim Development Studio**, which provides an integrated database development environment for Oracle, DB2, and Informix. Optim Development Studio improves development productivity up to 50% for developing and testing SQL and XQuery queries, stored procedures, Web services, and Java data access layers.

### Trial download

The evaluation period for this trial is 30 days. All of the product's features are enabled.

A trial version of Optim pureQuery Runtime V2.2 is available within the Optim Development Studio V2.2 trial package for use on the same Windows or Linux machine. To try the pureQuery Runtime on another platform, contact your IBM sales representative or IBM Business Partner.

➔ [System requirements](#)

The estimated download time using Download Director over a 1.5Mbps connection is 1 hour and 24 minutes; 43 minutes using Installation Manager.

Choose one of these options to download the trial. A Web install downloads and installs the product from a Web based repository. A local Install downloads and installs the product to your local hard drive.

Operating system	Version	Size	Download method
<a href="#">Web install using IBM Installation Manager (recommended)</a> Windows, Linux	V2.2	480MB	
<a href="#">Web install if you already have IBM Installation Manager installed</a> Windows	V2.2	372MB	

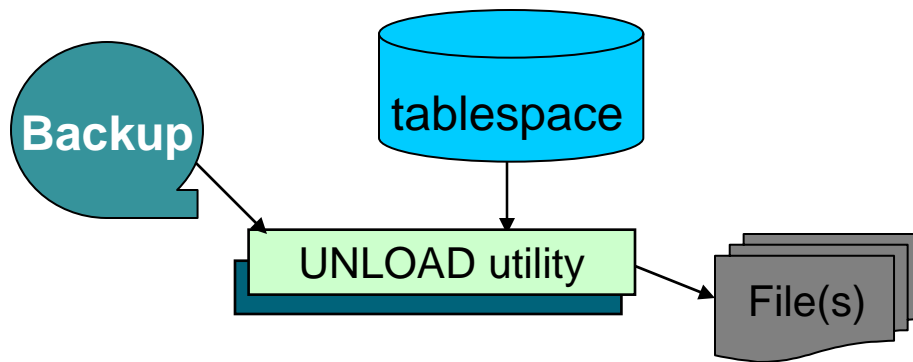
Download it today!



# Data movement options


## UNLOAD Utility

- Designed for loading back into DB2 for z/OS



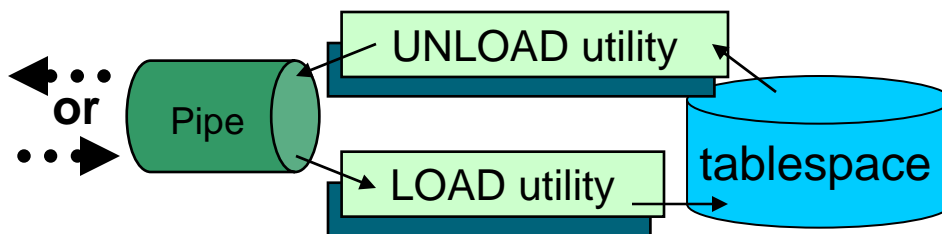
## Crossloader (LOAD from INCURSOR DD)

- Ultra fast loading of data when the source is in a relational table – even remote



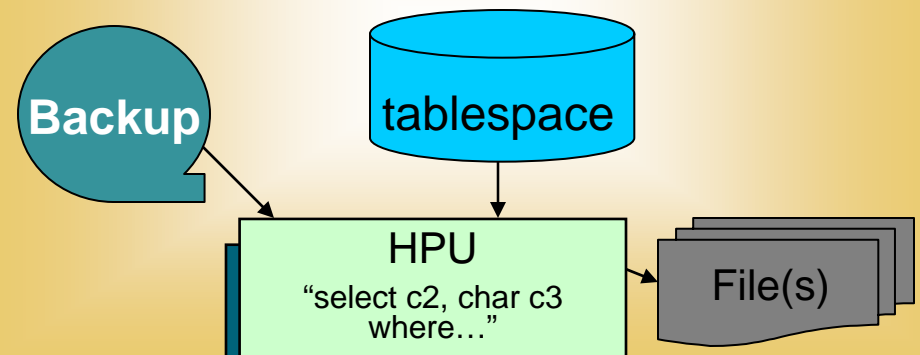
## LOAD/UNLOAD via pipes

- Perfect for remote loading/unloading



## High Performance Unload (z/OS and LUW)

- Ideal data moving to another platform
- SQL interface enables developer usage
- Low CPU consumption



# Recovery Tooling

- What value is there to extending the DB2 backup and recovery utilities?
  - Safety net for a complex task
  - Ability to view and make surgical changes using the DB2 recovery log
  - Allows you to use DB2 9 function on a DB2 Version 8 system

# Recovery management

- DB2 Change Accumulation Tool
- DB2 Log Analysis Tool
- DB2 Object Restore Tool
- DB2 Recovery Expert
  
- DB2 Cloning Tool

# DB2 Recovery Expert for z/OS



- Wizard driven
  - Object selection (Database, Table space, Table,..... Multiple objects supported)
  - Point in time selection (Time, Quiesce point, Backup point, Log point)
  - Recovery Paths / alternate recovery resources
- Analyzes all possible recovery paths for a version and attaches a relative cost to each
- Recommends the least cost recovery path while allowing other paths to be chosen
- Recovery paths include
  - Traditional restore and log apply (forward recovery)
  - Generating undo operations (backward recovery)
- Recommend set of objects to recover
  - Prompts the user to show related objects (next slide)
  - Assists in including these objects in the recovery process
- Manage recovery via policies: Recovery Policy – Insure that I can recover an application (set of DBMS objects) within twenty minutes to any point in time in the last week. {Supersedes the need for a backup policy}

ACTION	ROW STATUS	EMPNO	FIRSTNAME	MIDINIT	LASTNAME	WORKDEPT	PHONENO	HIREDATE
INSERT	POST-CHAN...	001100	Bryan	F	Smith	Z99	3474	2002-09-0
	PRE-CHANGE	-	-	-	-	-	-	-
INSERT	POST-CHAN...	001200	Dan		Wardman	Z99	4574	2002-09-0
	PRE-CHANGE	-	-	-	-	-	-	-
UPDATE	POST-CHAN...	000010	CHRISTINE	I	HAAS	A01	3978	1965-01-0
	PRE-CHANGE	000010	CHRISTINE	I	HAAS	A00	3978	1965-01-0
UPDATE	POST-CHAN...	000110	VINCENZO	O	LUCCHESI	A01	3490	1958-05-1

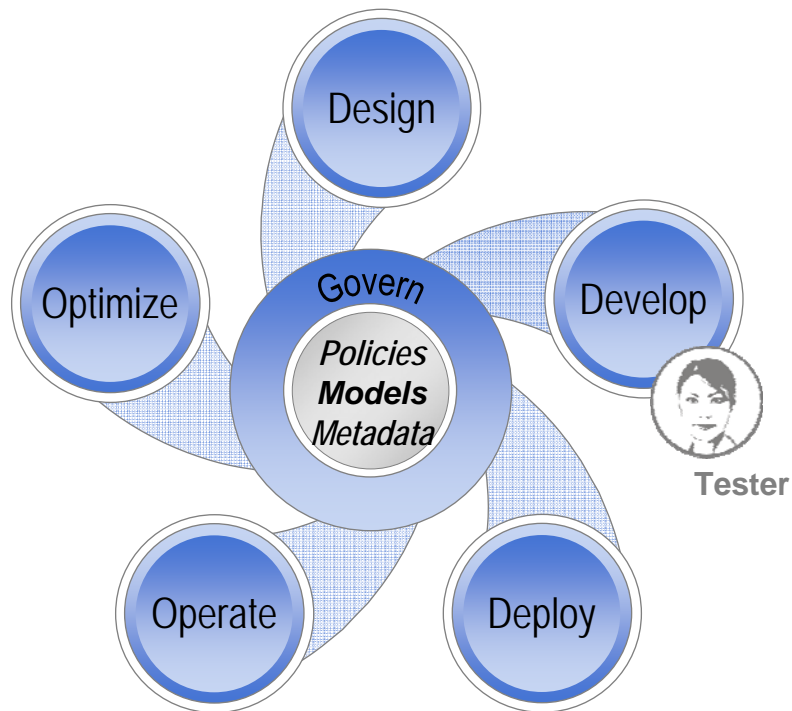


# Data Governance

- IBM Optim Data Growth
- IBM Optim Data Privacy
- IBM Optim Test Data Management
- DB2 Audit Management Expert
- IBM Database Encryption Expert
- Data Encryption for DB2 and IMS

# Create Production-like, Privatized Databases

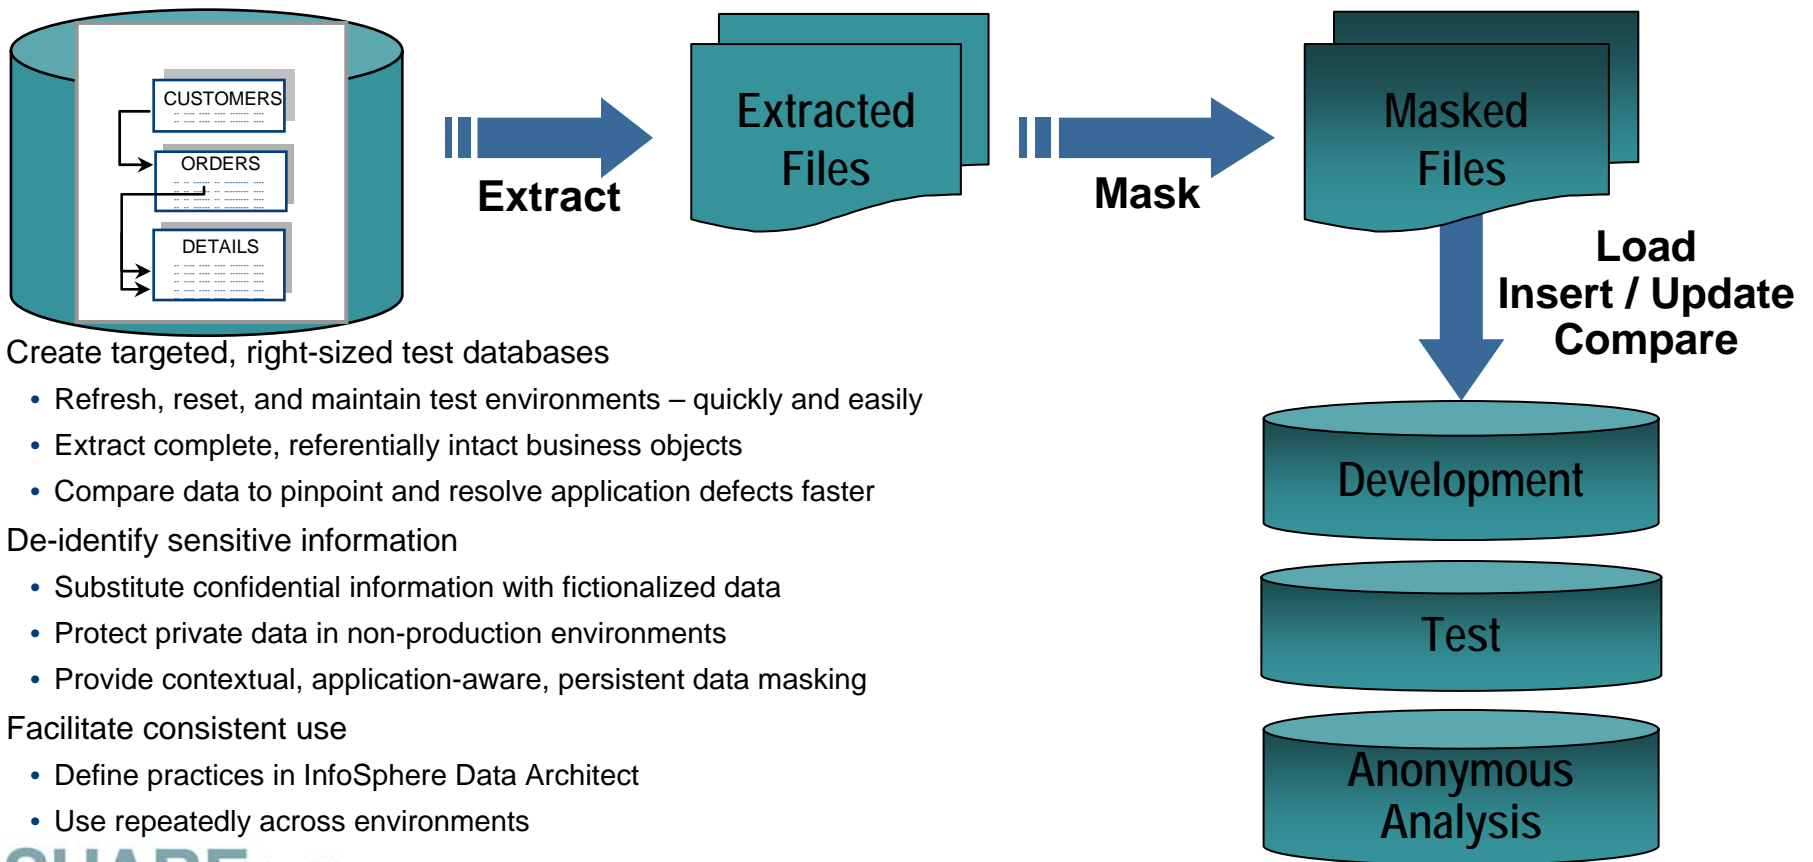
## *Optim Test Data Manager and Data Privacy Solutions*



- Increase software quality
  - Use test data that accurately reflects production systems
- Accelerate release schedules
  - Comparison to expected results and fast refresh
- Comply with privacy regulations
  - De-identify personal information with realistic substitutions
- Protect confidential information
  - Mask sensitive data
- Facilitate consistent use

# Optim Test Data Management and Data Privacy Solutions

**Increase software quality and comply with regulatory requirements based on test data that accurately reflects production systems while protecting client privacy and sensitive information.**



- Create targeted, right-sized test databases
  - Refresh, reset, and maintain test environments – quickly and easily
  - Extract complete, referentially intact business objects
  - Compare data to pinpoint and resolve application defects faster
- De-identify sensitive information
  - Substitute confidential information with fictionalized data
  - Protect private data in non-production environments
  - Provide contextual, application-aware, persistent data masking
- Facilitate consistent use
  - Define practices in InfoSphere Data Architect
  - Use repeatedly across environments



# Optimize Systems for Growth

## *Optim Data Growth Solution*



- Improve application performance
- Accelerate application upgrades
- Reduce the risk of application consolidation and retirement
- Comply with retention requirements
- Find information for legal discovery

Application  
Manager



**Optimize**  
*Archive,  
upgrade, retire*

**Design**

**Govern**

*Models  
Policies  
Metadata*

**Develop**

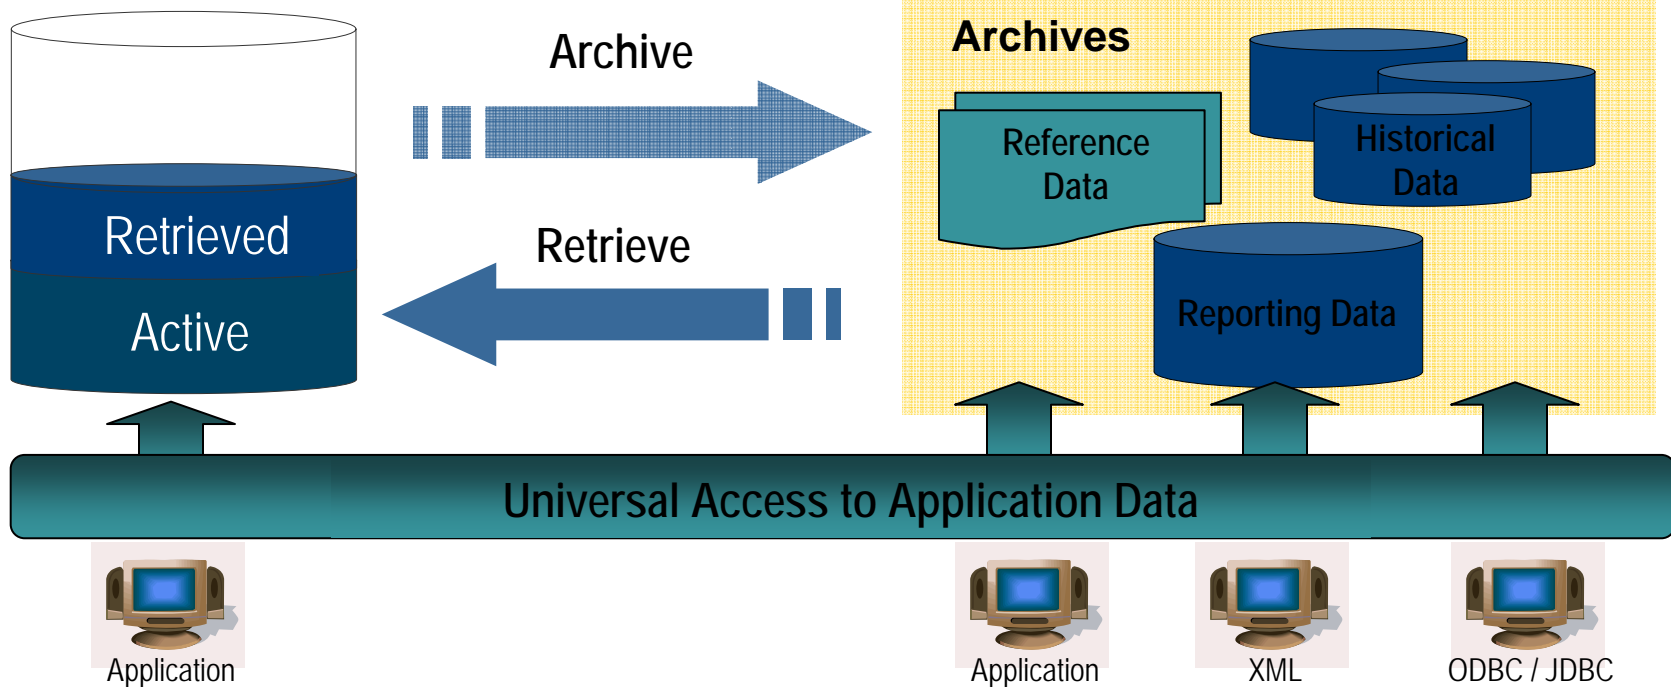
**Operate**

**Deploy**

# Optim Data Growth Solution



Facilitates application growth, application upgrades, application retirement, and retention management.



- Complete business object provides historical reference snapshot of business activity
- Storage device independence enables storage cost optimization
- Immutable file format enables data retention compliance

# What's cooking back at the lab?



- Performance Management

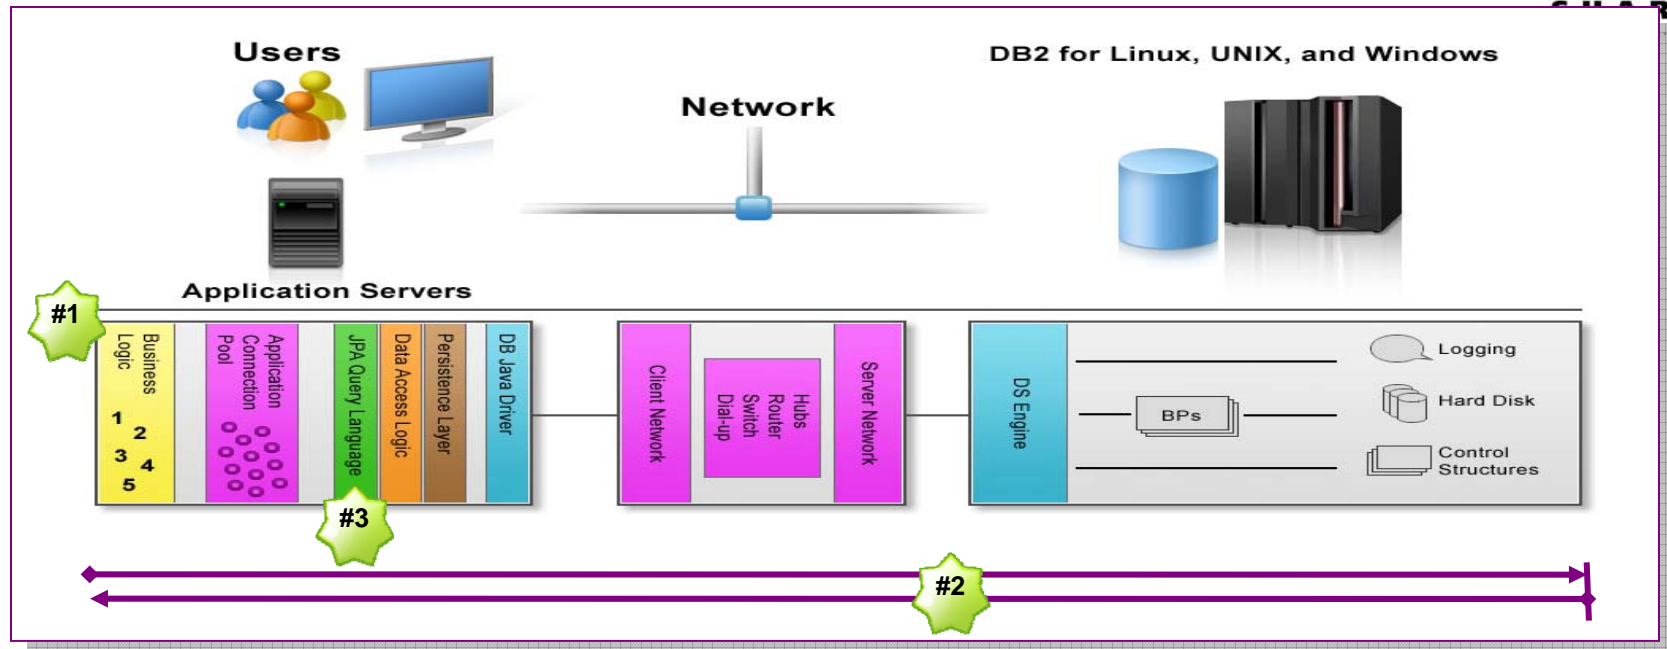
# The (not so new) differences between modern distributed apps and legacy Cobol and PL/I apps



Legacy apps	Distributed Java apps
<p>Mostly static SQL. Each program module was mapped to one DBRM under usually one DB2 PLAN.</p> <p>-&gt; Each app module automatically identified by auth ID, DBRM, and DB2 PLAN, and optionally CICS or IMS region</p>	<p>Primarily dynamic SQL. No DBRMs or DB2 PLANS.</p> <p>-&gt; Program had to identify itself via DB2 client info fields. Otherwise it would use default ids such as java.exe.</p>
<p>Program was running on mainframe only, no distributed access. Very small number of apps accessing database.</p> <p>-&gt; DBA knew each app personally ☺</p>	<p>Distributed apps can access database from anywhere. A ,new app' is developed in short period of time.</p> <p>-&gt; Increasing number of different apps accessing database. Not easy to keep an overview which app is running where.</p>
<p>DB2 pre-processor was well integrated into development languages</p> <p>-&gt; each SQL statement was identified by source line number</p>	<p>No pre-processor used. Program simply calls JDBC methods, or CLI functions to execute SQL statements</p> <p>-&gt; No information about location of SQL statement</p>
<p>Developer was coding SQL directly</p> <p>-&gt; easy to find and change SQL statement if necessary</p>	<p>Program is generating the statement text on-the-fly by concatenating strings, or frameworks and persistence layers like iBatis, Spring, EJB, JPA are generating the SQL statements out of programming language constructs</p> <p>-&gt; can be extremely difficult to identify origin of SQL statement. Changing SQL statement text can be impossible if you can't control the generation.</p>



# So, what is it what we are looking for?



Something ...

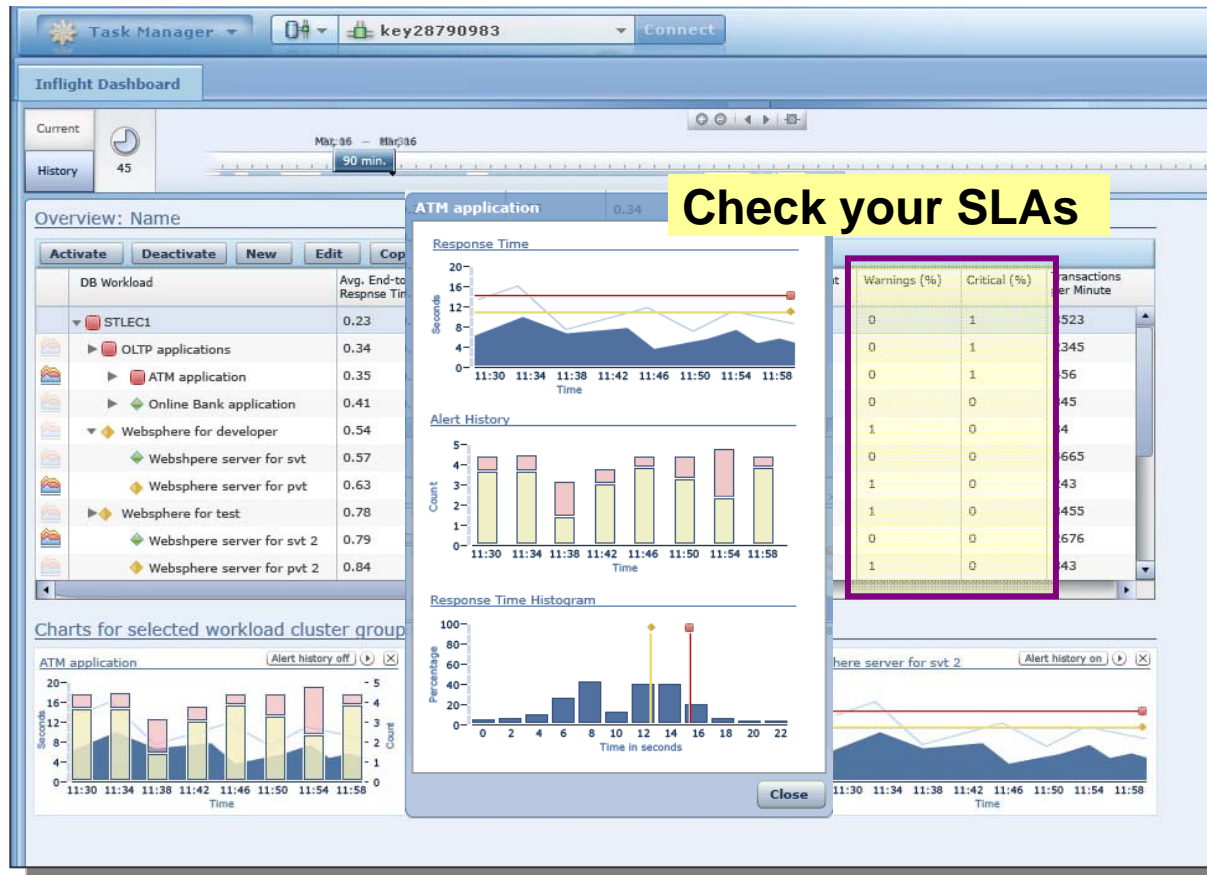
- #1 ...which tells me which application it is, and ultimately which business function is performing
- #2 ...where I can measure what my application/user is really experiencing (from application to database and back) and which tells me which components are involved and where my application is spending its time.
- #3 ...which tells me where the SQL statement is coded and let's me change it if necessary.



# Performance Expert Extended Insight Feature

- IBM offers a feature for its database monitor DB2 Performance Expert for LUW called ,Extended Insight‘ to collect exactly this information automatically ...
  - Supports standard applications (WebSphere, InfoSphere Warehouse, SAP, DataStage, Cognos, eXtremeScale) and homegrown Java applications setting the DB2 client information fields
  - Support for static and dynamic SQL
  - Works with DB2 for LUW V9.1 or higher. DB2 for LUW V9.7 will give you additional insight into database side – where your application is spending its time
  - Easy deployment: only a few files have to be copied into location of DB2 driver on application side; no configuration on client side necessary.
  - Collects statement and transaction level information with low overhead on client side (<= 5% overhead)
  - Allows to check for Service Level Agreement (SLA) violations

# Database Workload Overview - Extended Insight Feature



## End-to-end database monitoring

- ... the most advanced way to monitor your applications and solution accessing data in a database.
- introduced in DB2 Performance Expert V3.2 in 12/2008 as priced feature with focus on WebSphere applications
- now enhanced to support out-of-the-box: WebSphere, Cognos, DataStage, SQL Warehouse, SAP
- will also be used to monitor other applications using JCC or CLI database drivers (.NET support will follow)
- intuitive to use because looking at the applications/solutions behind it and not into DB2 resources or DB connections
- will auto-detect the application behind the DB connection



# Database Workload Details - Extended Insight Feature



Recent aggregation : 2  
History 60 sec  
12:01 14:01 05/14/2009 13:58  
120

End-to-End Dashboard Overview TEST Config Connection

Overview : Test

End to End response time  
Partition : Selected AreaSeries : E2E\_DS

Partition / Sql Statements / Clients  
Top 10 by End-to-end respons...

Statement text	Execution	Avg. data ...	End-to-en...
Select username, account, address from user where...	10	25	45
Select username, account, address from user where...			
Select student, course, address from student			
Select firstname, lastname, address from user			

Detail Area

Statement information  
select \* from db where.....

Package Name perf.widonet.test  
Package Schema schema1  
Package Version version2.1

Java Class	pack	Method	Source Line No	Build Version	Build Time
event		clone	12	version1.0	04-06-2009
event		clone	12	version1.0	04-06-2009
event		clone	12	version1.0	04-06-2009

Transfer Volume

Average bytes sent	12456 KB
Average bytes received	45781 KB
Average rows returned	12148

Statement Performance

No. of executions :	execution number 1
Average end-to-end elapsed time	45
Average client time	25
Average driver time	
Average network time	
Average data server time	

Time Distributions

25% 25% 25% 25%

First negative SQL code :

- showing you the real response time of your application and where it is spending its time
- including database side – new with DB2 V9.7 FP1

- seamlessly integrated with Optim Query Tuner; identify problematic SQL and tune it

# Tell me where my SQL is and let me control it

## Or why JPA and other frameworks/persistency layers are then a nightmare for a DBA

- SQL statement are not coded directly and generated on-the-fly during execution

<pre> @Entity @Table(name="ORDER_TABLE") public class Order {     private int id;     private String address;     private Customer customer;      @Id     @Column(name="ORDER_ID")     public int getId() {         return id;     }      public void setId(int id) {         this.id = id;     } </pre>	<pre> @Column(name="SHIPPING_ADDRESS") public String getAddress() {     return address; }  public void setAddress(String address) {     this.address = address; }  @ManyToOne() @JoinColumn(name="CUSTOMER_ID") public Customer getCustomer() {     return customer; }  public void setCustomer(Customer customer) {     this.customer = customer; } </pre>
--	---

Java annotations define OR mapping, so that Entity Manager can generate SQL

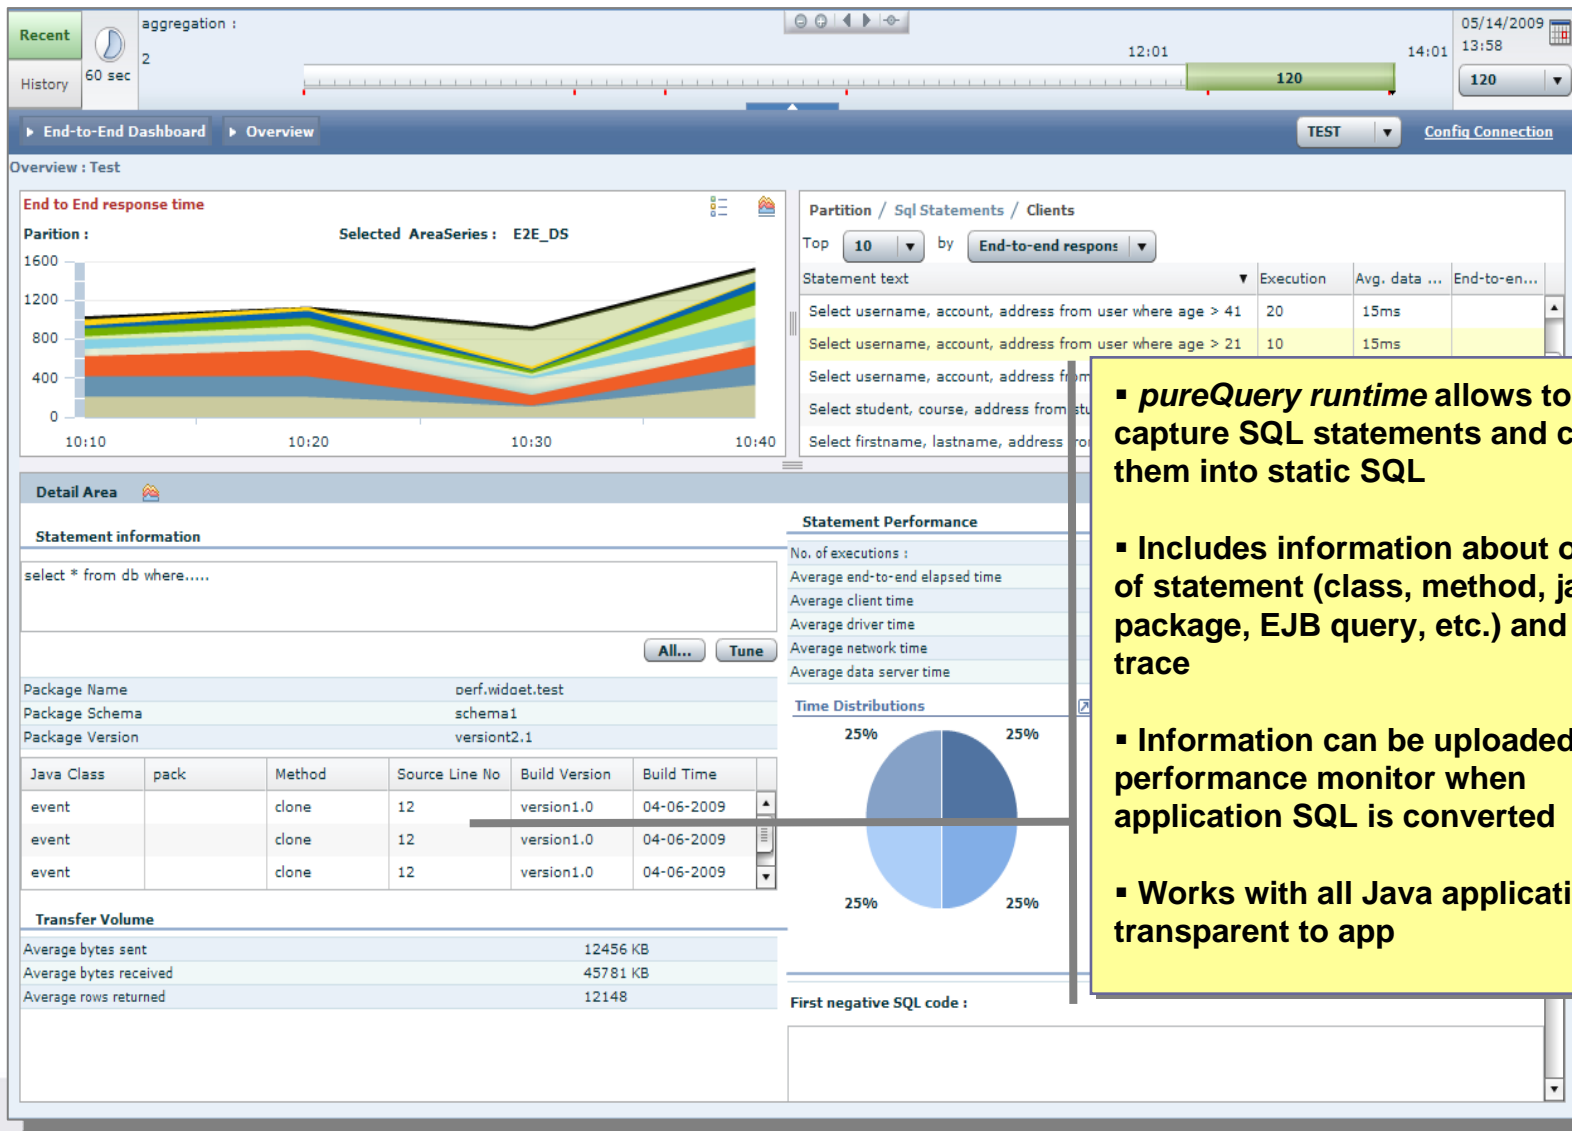
- The developer can code additional queries – in SQL but also in its own EJB Query Language, which is once again ‘translated’ into SQL statements on the fly

<pre> orm.xml: &lt;entity class="Order"&gt; ... &lt;named-query name="SoldOut"&gt;   &lt;![CDATA[ SELECT p FROM Order p               WHERE p.status == "Sold Out" ]]&gt; &lt;/named-query&gt; &lt;/entity&gt; </pre>	<pre> Query query = em.createNamedQuery("SoldOut"); List results = query.getResultList(); </pre>
---	--

**How can you fix a problematic SQL statement if you can't even find it?**

# Tell me where my SQL is and let me control it

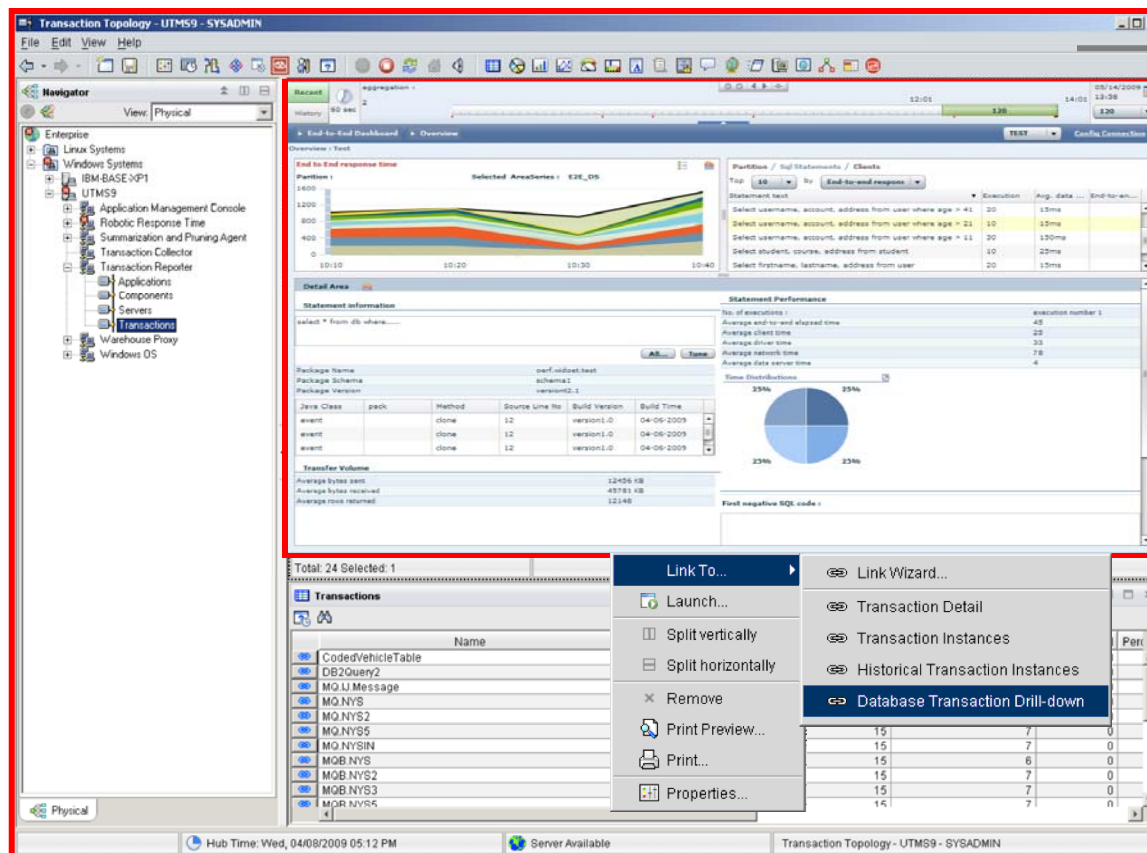
How pureQuery runtime can help



- pureQuery runtime allows to capture SQL statements and convert them into static SQL
- Includes information about origin of statement (class, method, java package, EJB query, etc.) and stack trace
- Information can be uploaded into performance monitor when application SQL is converted
- Works with all Java applications; transparent to app

# Database Performance Monitoring + Tivoli Enterprise portal

- Can run stand alone, or inside Tivoli Enterprise Portal
- Integrated into ITCAM for Transactions; high application response times can be directly analyzed with database performance monitor



The screenshot displays the Tivoli Enterprise Portal interface for database performance monitoring. The main window is titled 'Transaction Topology - UTM59 - SYSADMIN'. The interface is divided into several sections:

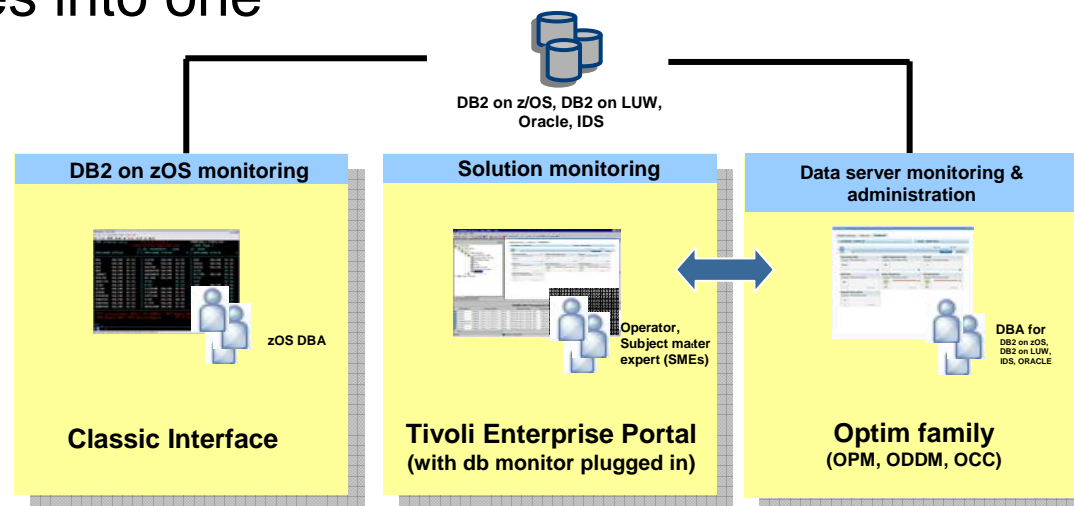
- Navigator:** A tree view on the left showing the system hierarchy, including 'Enterprise', 'Linux Systems', 'Windows Systems', 'IBM-BASE>VP1', 'UTM59', and various components like 'Application Management Console', 'Robotic Response Time', 'Summarization and Pruning Agent', 'Transaction Collector', 'Transaction Reporter', 'Applications', 'Components', 'Servers', 'Warehouse Proxy', and 'Windows OS'.
- Overview:** A central area showing a line chart for 'End-to-end response time' with a selected area series 'E2E\_OS'. Below the chart is a table for 'Statement Performance' with columns for 'Statement text', 'Execution', 'Avg. data', and 'End-to-end...'. The table lists several SQL queries and their execution times.
- Detail Area:** A section for 'Statement Performance' showing 'Statement information' and 'Time Distributions' with a pie chart.
- Transactions:** A table at the bottom listing various transactions such as 'CodedVehicleTable', 'DB2Query2', 'MQ.U.Message', 'MQ.NYS', 'MQ.NYS2', 'MQ.NYS5', 'MQ.NYSIN', 'MOB.NYS', 'MOB.NYS2', 'MOB.NYS3', and 'MOR.NYS4'. A context menu is open over this table, showing options like 'Link To...', 'Launch...', 'Split vertically', 'Split horizontally', 'Remove', 'Print Preview...', 'Print...', and 'Properties...'. The 'Database Transaction Drill-down' option is highlighted.

Tivoli Enterprise Portal

transaction topology

# How to prepare yourself for the future ...

- New functionality will not replace existing Optimization Expert, Query Monitor, or OMEGAMON XE products, but be an add-on to them ...
- It will most likely have a pre-requisite for *OMEGAMON XE Performance Expert for DB2 on z/OS (OMPE)*
- Long-term: It might consolidate OMPEs graphical user interfaces into one




# What's cooking back at the lab ?

- Configuration Management
  - Speeding up problem determination by analyzing what has changed
  - Simplifying implementation of and checking adherence to best practices
  - Centralizing control of client configuration
- Maintenance Management
  - Managing large, complex enterprise environments by exception
  - Shifting from time-based maintenance via utilities to policy-based maintenance
  - Visualization of space management



# The problem ... Our solution!






# Manage by Exception to Lower TCO

Name	Monitoring Status		Data Server Status		Alert		System		Database			
	Critical	Warning	CPU Usage	Disk Space	Memory Usage	Locking	SQL Performance	Connections	Transactions	Logging	Maintenance	
Production	3	8	⬇️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Web	1	1	⬇️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Retail	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Accounts	2	3	⬇️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Marketing	0	4	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Test	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
Development	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️

## Disk Efficiency Alert

 *In excess of 40% of allocated space for database objects can be reclaimed.  
Recommended courses of action:*



- ▷ Turn off alert - "I don't care or I will handle this all on my own"
- ▷ Change alert threshold - "Customize to my environment or workload"
- ▷ Reclaim space now - "I need to fix it now, but need some options to tailor."
- ▷ Reclaim space later - "Schedule during maintenance window, possibly repeatedly."
- ▷ Automate space reclamation - "Just do it for me, I have more important things to do."

# Enable Fully-manual to Fully-automated Management


Name	Monitoring Status		Data Server Status		Alert		System			Database			
	Critical	Warning	Critical	Warning	CPU Usage	Disk Space	Memory Usage	Locking	SQL Performance	Connections	Transactions	Logging	Maintenance
▼ Production	3	8	3	8	⬇️	⬆️	⬆️	⚠️	⬆️	⬆️	⬆️	⬆️	⬆️
▶ Web	1	1	1	1	⬇️	⬆️	⬆️	⚠️	⬆️	⬆️	⬆️	⬆️	⬆️
▶ Retail	0	0	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
▢ Accounts	2	3	2	3	⬇️	⬆️	⬆️	⚠️	⬆️	⬆️	⬆️	⬆️	⬆️
▢ Marketing	0	4	0	4	⬆️	⬆️	⬆️	⚠️	⬆️	⬆️	⬆️	⬆️	⬆️
▶ Test	0	0	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️
▶ Development	0	0	0	0	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️	⬆️

## Best Practice Violation Alert: **Log and Data Files Collocated on Single Disk**

⬇️ *Log files are on same disk as table space container files which negatively impacts performance and may potentially prevent full recovery upon disk failure.*

Recommended courses of action:

- ▶ Turn off alert - "I will figure out how to handle this all on my own"
- ▶ Manually reconfigure now - "I need to fix it myself now, but a little guidance would be nice."
- ▶ Reconfigure later - "Script the reconfiguration and schedule during maintenance window."
- ▶ Automate this best practice - "Just do it for me, I have more important things to do."



# Manage by Exception to Lower TCO



Heat Chart Alerts Dashboard

Next Refresh 0:01

Name	Monitoring Status		Data Server Status		Alert		System			Database			
	Critical	Warning	Critical	Warning	CPU Usage	Disk Space	Memory Usage	Configuration	SQL Performance	Connections	Transactions	Logging	Maintenance
▼ Production	3	8	3	8	1	1	0	0	0	0	0	0	0
▶ Web	1	1	0	0	0	0	0	0	0	0	0	0	0
▶ Retail	0	0	2	3	0	0	0	0	0	0	0	0	0
▢ Accounts	0	4	0	0	0	0	0	0	0	0	0	0	0
▢ Marketing	0	0	0	0	0	0	0	0	0	0	0	0	0
▶ Test	0	0	0	0	0	0	0	0	0	0	0	0	0
▶ Development	0	0	0	0	0	0	0	0	0	0	0	0	0

## Configuration Change Alerts:

- Table space container 'tblsp25003' is no longer on disk
- Index 'idx00453' has been dropped from table 'mkt.tblA'
- Column 'A' of table 'mkt.tblA' changed data type from INT to FLOAT
- Configuration parameter 'Automaintenance' changed from ON to OFF



# IBM Integrated Data Management

- Getting started
- Downloads
  - IBM Data Studio
  - IBM Optim
- Community
  - developerWorks
    - forums
    - zone and space
    - articles
    - tutorials
- Documentation
- Support



A screenshot of the IBM Integrated Data Management website. The page has a dark blue header with the IBM logo and navigation links. The main content area features a central diagram with "Governance" in the center, surrounded by "Design", "Develop", "Deploy", "Operate", and "Optimize". To the right of the diagram is the text "Develop Code, debug, test, tune, package". Below this is a section titled "What is Integrated Data Management" followed by "What we offer" which lists various solutions like "Accelerate Solution Delivery", "Facilitate Integrated Database Administration", etc. There are also sections for "All Integrated Data Management products" and "Browse by application". The right sidebar contains several promotional boxes, including "E-mail us", "IBM Information Demand", "Get IBM Data Studio", "New! Optim Java Acceleration", "New! Software bundle DBAs", "Optim trial software!", and "Optim Downloads".





# Data Management Communities for DB2

- IDUG – the worldwide community of DB2 users
  - Membership is FREE – join today! [www.idug.org](http://www.idug.org)
- Data Management Community – share and interact with peers around the world
  - [www.ibm.com/software/data/management/community.html](http://www.ibm.com/software/data/management/community.html)
- Information Champions – recognizes individuals who have made the most outstanding contributions to the Information Management community
  - [www.ibm.com/software/data/champion](http://www.ibm.com/software/data/champion)

# Disclaimer



© Copyright IBM Corporation [current year]. All rights reserved.

**U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.**

**THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM’S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS AND/OR SOFTWARE.**

IBM, the IBM logo, ibm.com, DB2, Optim, Tivoli, Rocket Software, and Data Studio are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml)

Other company, product, or service names may be trademarks or service marks of others.