IBM Information Management

Introduction to Integrated Data Management

Extending Tooling Solution to Manage Data Over its Lifetime

Hong Sang Tie
Senior Development Manager,
Optim Query Tuner
IBM Silicon Valley Laboratory

DB2 Technical Forum
Taipei, Taiwan
Oct. 5th-6th, 2009

Agenda

- Integrated Data Management
  - Business Challenges Today
  - Solution Overview

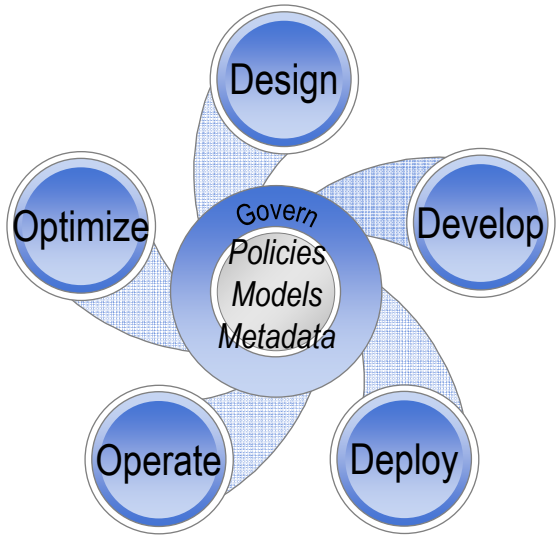

What do Businesses Have?

A Collection of Disparate, Single-Purpose Products

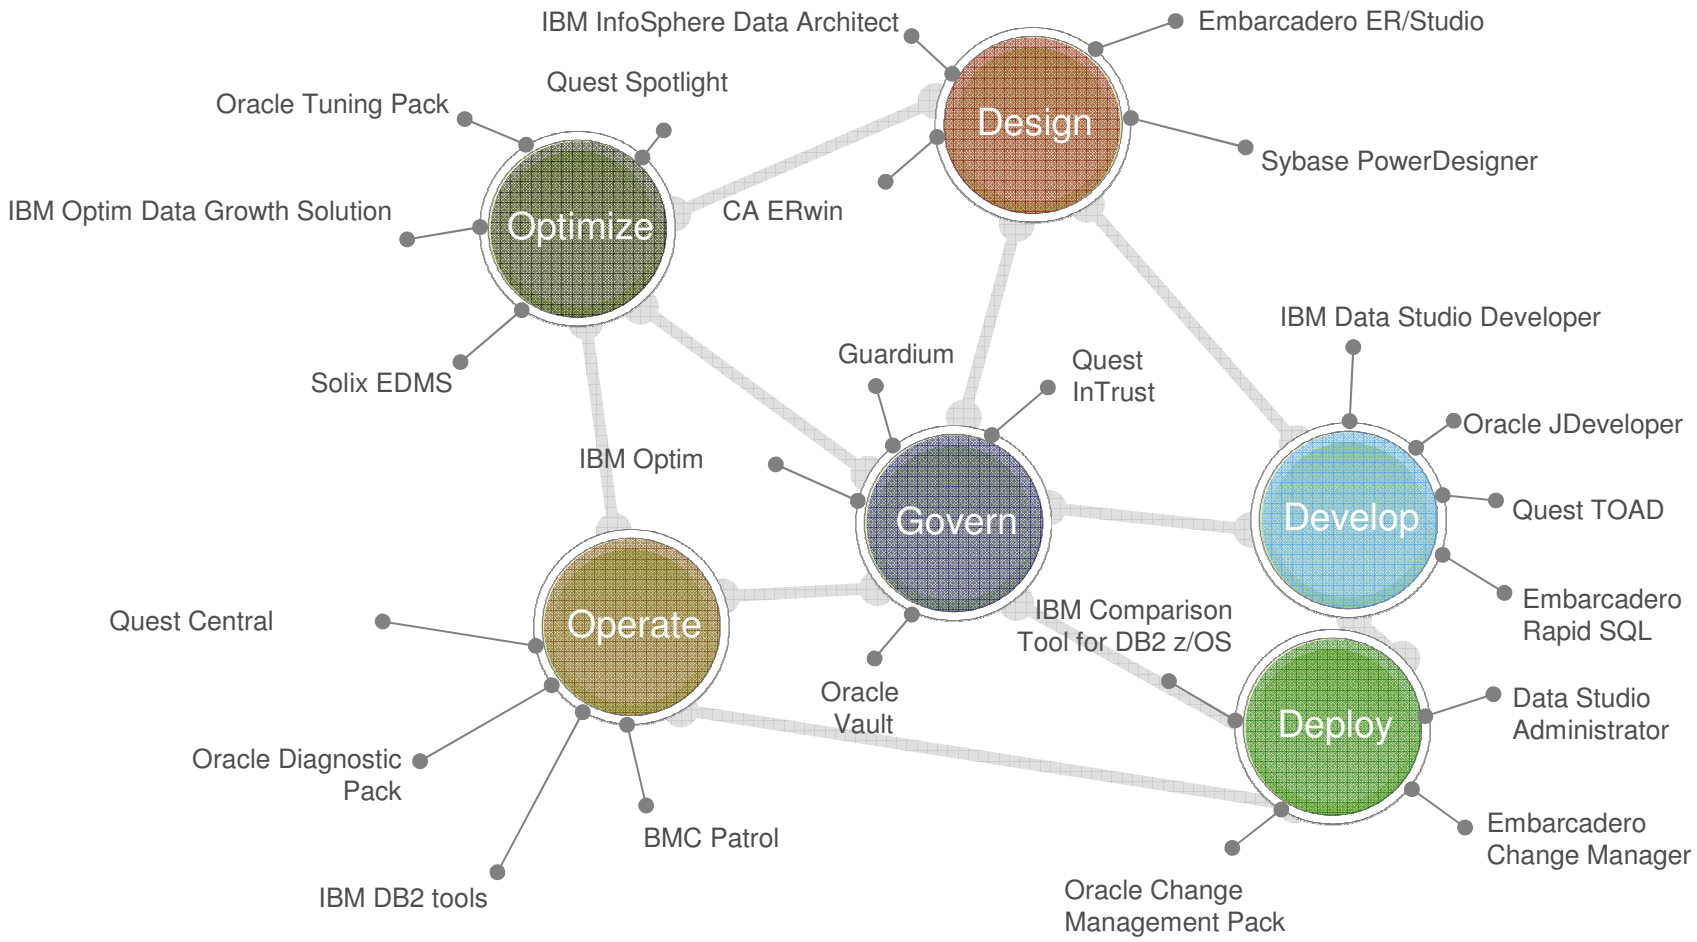

The gaps create risk …

- Loss of customers
  - Average customer churn rate up 2.5% after a breach
- Loss of revenue
  - $197 USD per customer record leaked
  - Average cost was ~ $6.3 million / breach in this study
  - Average cost for financial services organizations was 17% higher than average
- Fines, penalties or inability to conduct business based on non-compliance
  - PCI
  - Sarbanes-Oxley (SOX)
  - HIPAA
  - Data Breach Disclosure Laws
  - Gramm-Leach-Bliley Act
  - Basel II

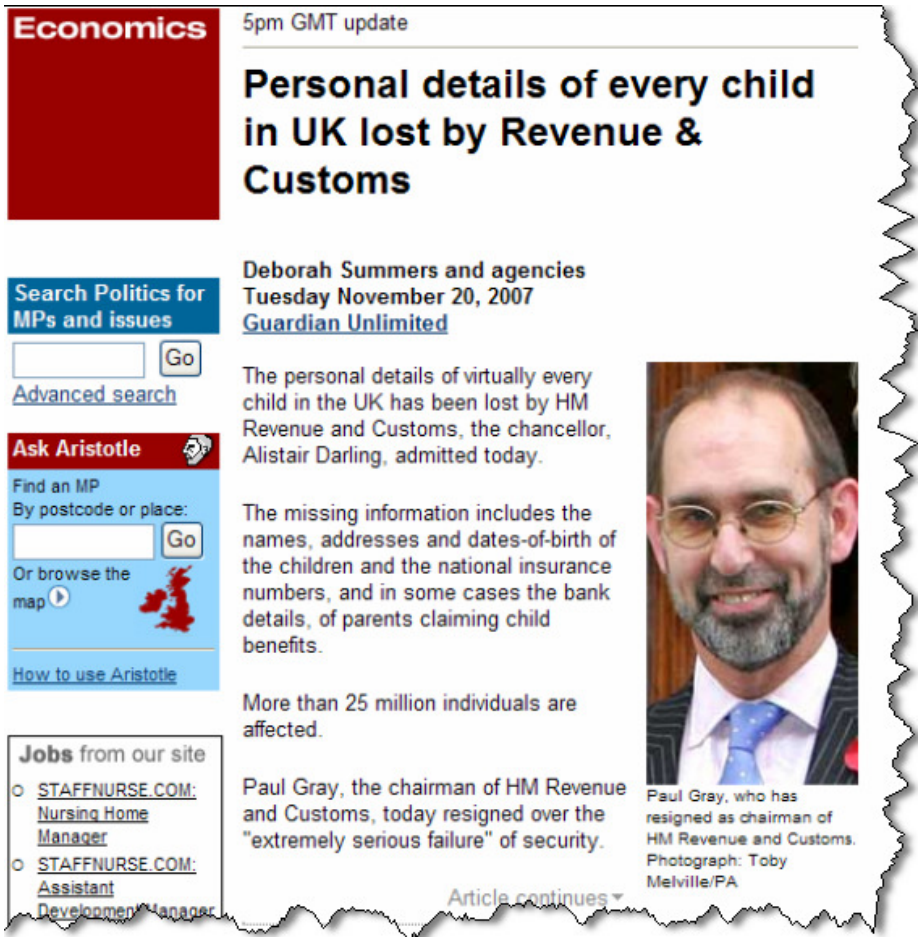

Economics

5pm GMT update

Personal details of every child in UK lost by Revenue & Customs

Search Politics for MPs and issues

Go

Advanced search

Ask Aristotle

Find an MP
By postcode or place:

Go

Or browse the map

How to use Aristotle

Jobs from our site

- STAFFNURSE.COM: Nursing Home Manager
- STAFFNURSE.COM: Assistant Development Manager

Deborah Summers and agencies
Tuesday November 20, 2007
Guardian Unlimited

The personal details of virtually every child in the UK has been lost by HM Revenue and Customs, the chancellor, Alistair Darling, admitted today.

The missing information includes the names, addresses and dates-of-birth of the children and the national insurance numbers, and in some cases the bank details, of parents claiming child benefits.

More than 25 million individuals are affected.

Paul Gray, the chairman of HM Revenue and Customs, today resigned over the "extremely serious failure" of security.

Article continues

Paul Gray, who has resigned as chairman of HM Revenue and Customs. Photograph: Toby Melville/PA

Source: "2007 Annual Study: Cost of a Data Breach" , The Ponemon Institute

Worldwide IT Spending on Servers, Power and Cooling, and Management/Administration

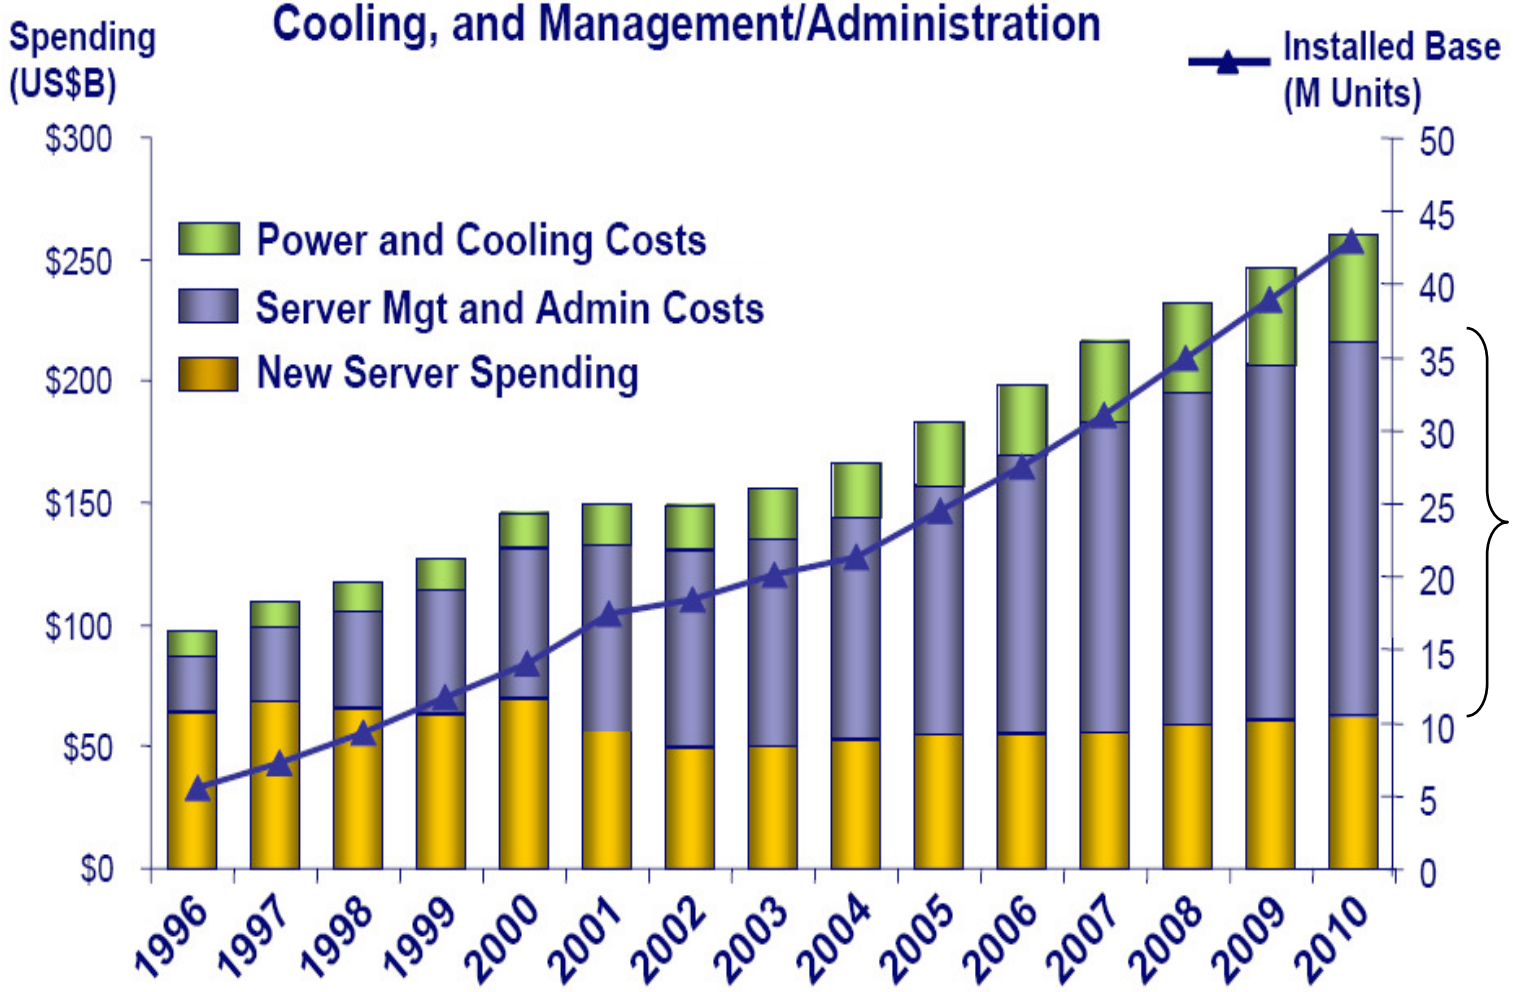

Driven by the increasing numbers of physical systems, system management has become the main component of IT costs and is growing rapidly

Many Servers, Much Capacity, Low Utilization = $140B unutilized server assets

Source: IDC, 2006

What do Businesses Need?

An integrated environment to span today's flexible roles

- **Manage data throughout its lifecycle**
  - From design to sunset
- **Manage data across complex IT environments**
  - Multiple interrelated databases, applications and platforms
- **Facilitate cross-functional collaboration**
  - Within IT
  - Among Line of Business, Compliance functions
  - Across disparate skill sets
- **Optimize business value**
  - Respond quickly to emerging opportunities
  - Improve quality of service
  - Reduce cost of ownership
  - Mitigate risk

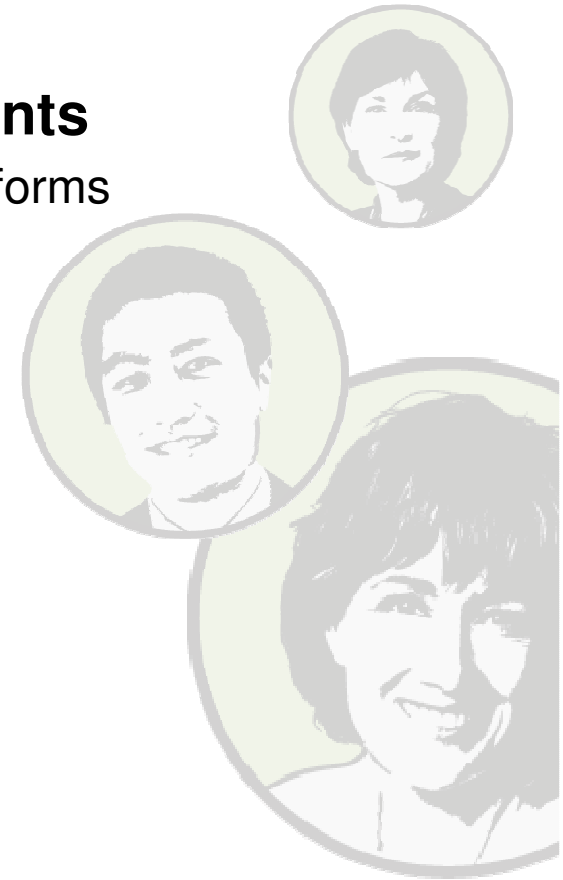

Introducing Integrated Data Management

An integrated, modular environment to design, develop, deploy, operate, optimize and govern enterprise data throughout its lifecycle on the System z platform

Enabling organizations to more efficiently and effectively

- **Respond to emergent, data-intensive business opportunities**
- **Meet service level agreements for data-driven applications**
- **Comply with data privacy and data retention regulations**
- **Grow the business while driving down total cost of ownership**

Integrated Data Management

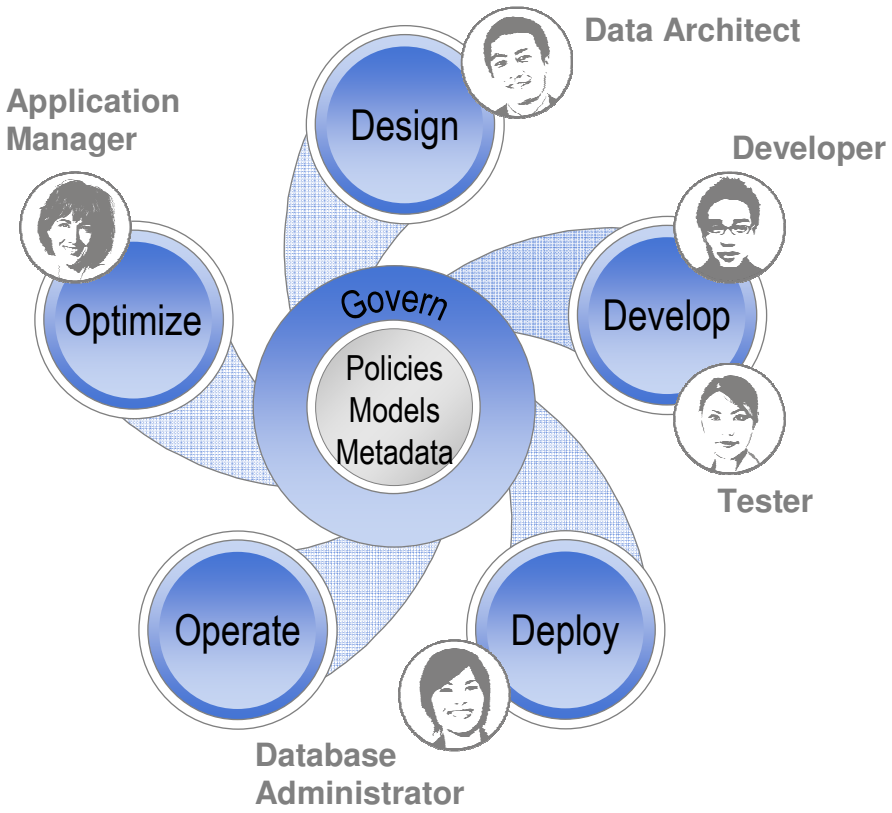

- **Deliver increasing value across the lifecycle**, *from requirements to retirement*
- **Facilitate collaboration and efficiency across roles,** *via shared artifacts, automation and consistent interfaces*
- **Increase ability to meet service level agreements,** *improving problem isolation, performance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

Integrated Data Management Solutions

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

InfoSphere Data Architect**

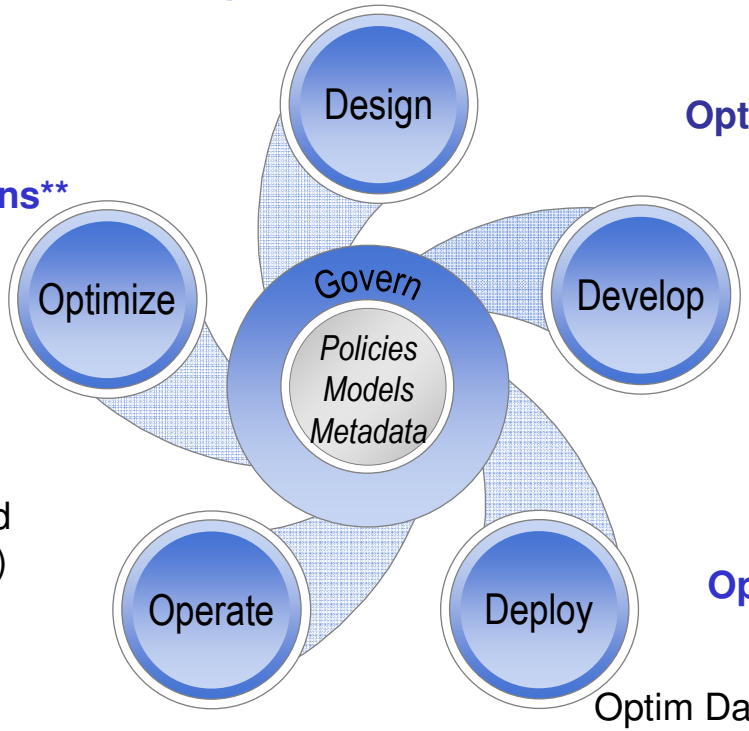

Optim Development Studio**
Optim Query Tuner**

Optim Data Growth Solutions**

Optim Test Data Management**

Optim Query Workload Tuner (z)

Optim Data Privacy Solutions**

DB2 Performance Expert and Extended Insight (DB2 LUW)
Omegamon XE (z)

Optim pureQuery Runtime**

DB2 z/OS Tools

Optim Database Administrator (LUW)
DB2 Admin Tool/Object Compare (z)

DB2 z/OS Utilities Suite

DB2 Audit Management Expert**
Database Encryption Expert

Heterogeneous Offerings
***** z and LUW***

IBM Information Management

InfoSphere Data Architect

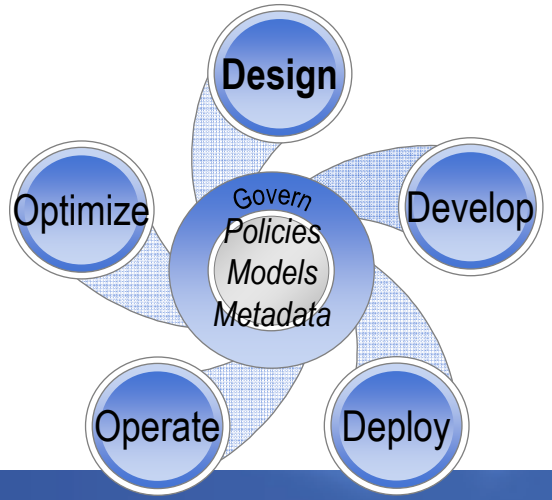

InfoSphere Data Architect

A collaborative, data design solution to discover, model, relate, and standardize diverse data assets.

Key Features

- Create logical and physical data models
- Discover, explore, and visualize the structure of data sources
- Discover or identify relationships between disparate data sources
- Compare and synchronize the structure of two data sources
- Analyze and enforce compliance to enterprise standards
- Support across heterogeneous databases
- Integration with the Rational Software Delivery Platform, Optim, IBM Information Server, and IBM Industry Models

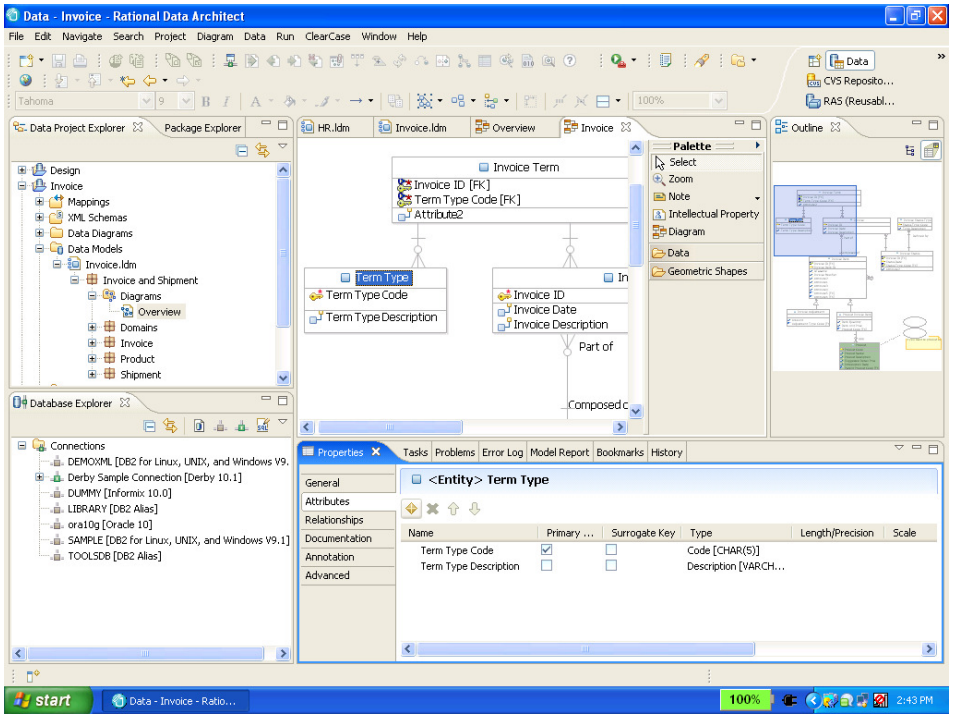

Automate Data Design via Model-driven Transformation

- **Built-in transformation**
- **Compare and sync facilitates merge**

Rational Software Architect

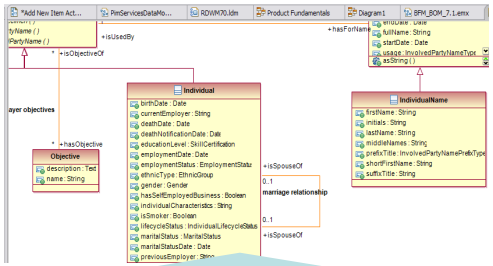

UML

SOLUTION ARCHITECT

WebSphere Business Modeler

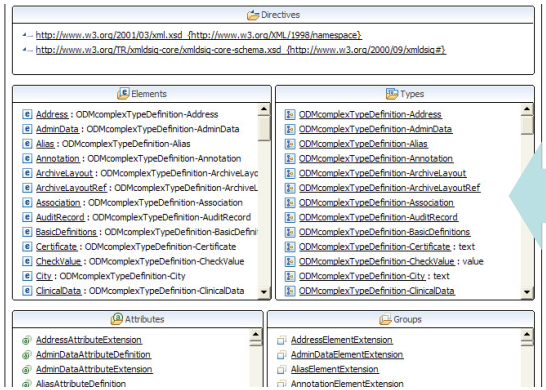

XSD

InfoSphere Data Architect

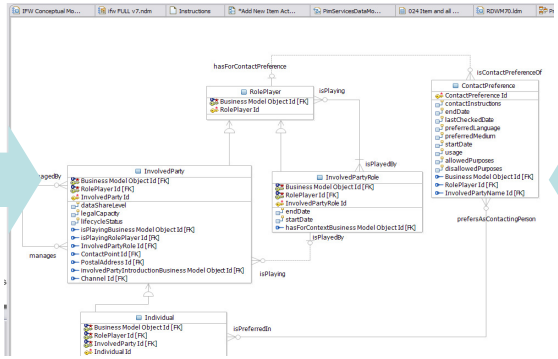

PDM

Optim DataBase Administrator

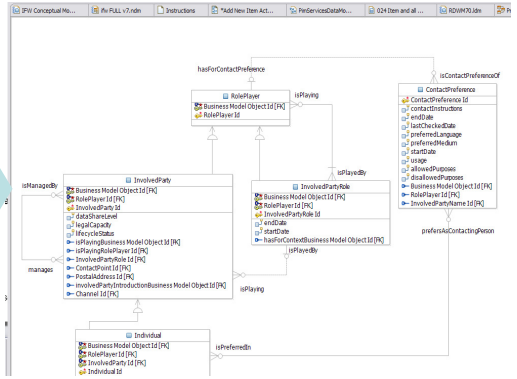

INTEGRATION DEVELOPER

DATA ARCHITECT

DATABASE ADMINISTRATOR

IBM Information Management

Optim Test Data Management Solution

Optim data Privacy Solutions

Optim Test Data Management Solution

Streamline building test databases, improve application quality, cut IT costs and accelerate solution delivery

- **Accelerate time to market**
  - Create “right sized” test databases
  - Extract referentially intact subsets
  - Compare baseline data against test results to pinpoint and resolve application defects faster
  - Edit test data to create error and boundary conditions
  - Easily refresh, reset and maintain test environments
- **Cut storage costs**
  - Reduce storage requirements by using smaller subsets for testing
- **Enable compliance**
  - De-identify or mask data

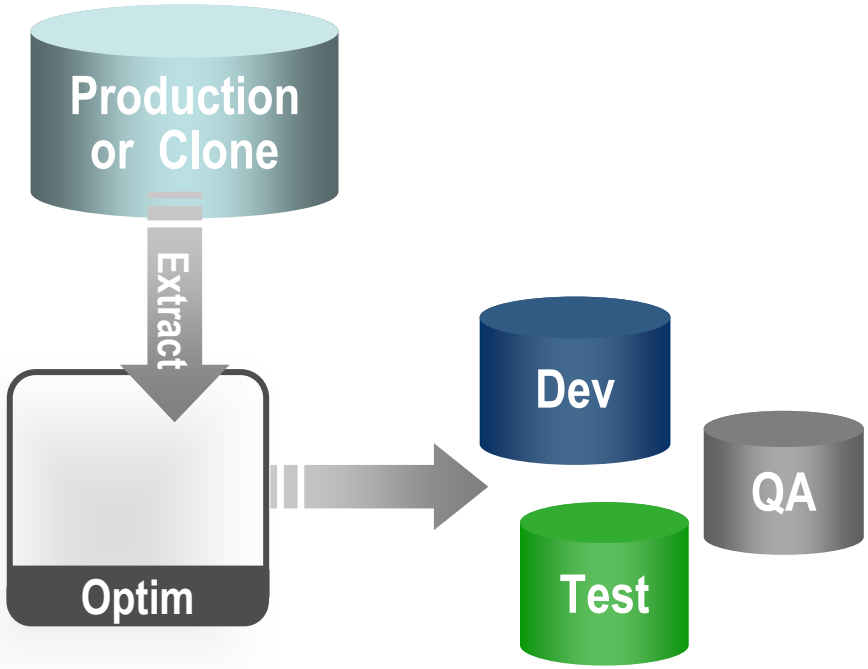

Enterprise Challenge: Data Privacy

Optim Data Privacy Solution

- *Application-aware masking capabilities ensure data is realistic but fictional*
- *Prepackaged data masking routines make it easy to de-identify elements*
  - *E.g. credit card numbers & email addresses*

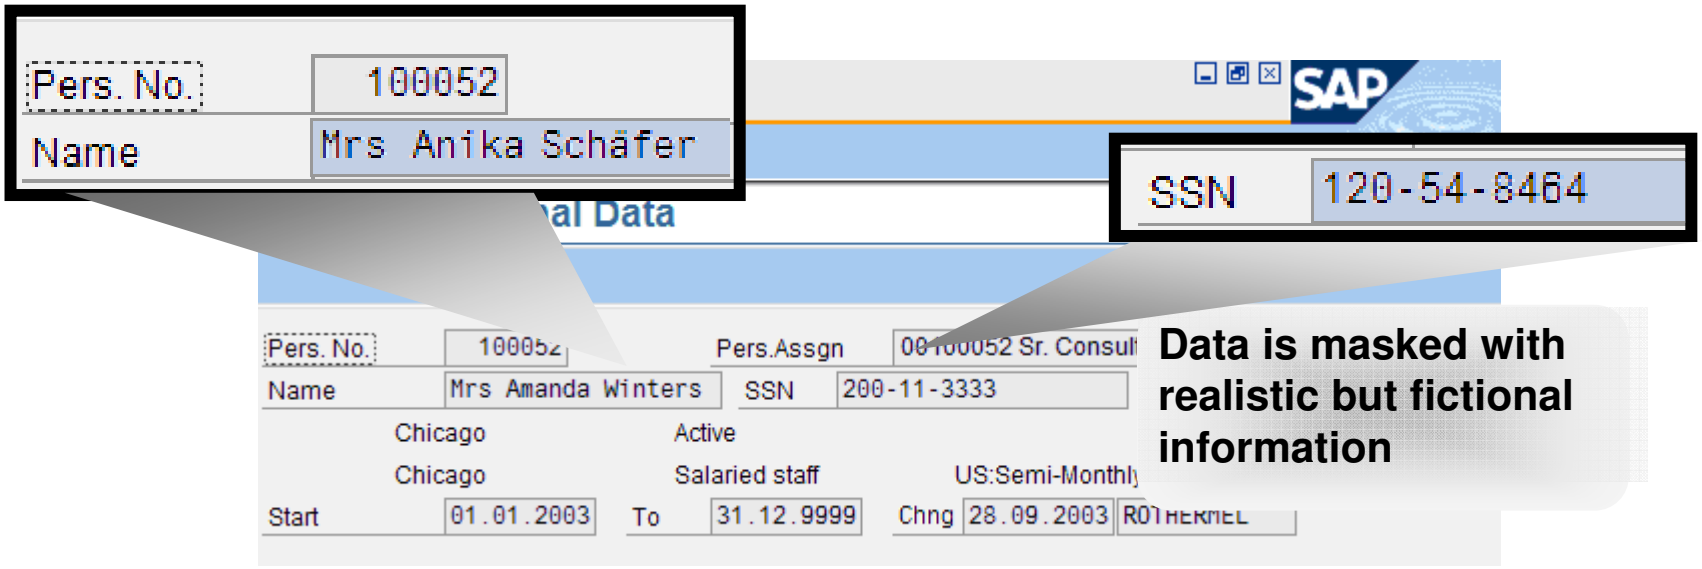

A comprehensive set of data masking techniques to transform or de-identify data, including:

- **String literal values**
- **Character substrings**
- **Random or sequential numbers**
- **Arithmetic expressions**
- **Concatenated expressions**
- **Date aging**
- **Lookup values**
- **Intelligence**

IBM Information Management

Optim Query Tuner
Optim Query Workload Tuner

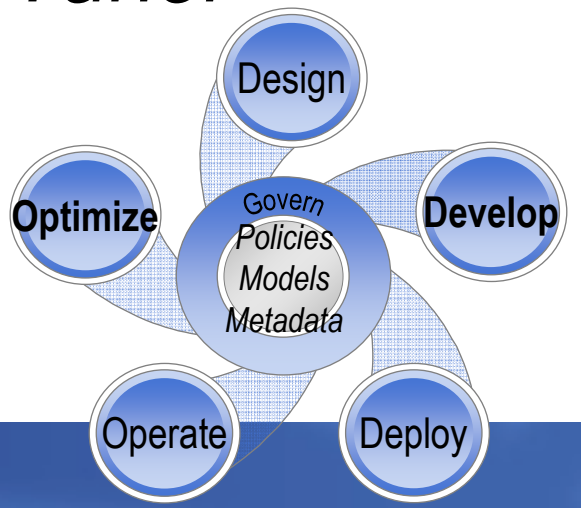

Optim Query Tuner

Maximize performance and reduce specialized skill requirements

- ***Optim Query Tuner (a.k.a. Optimization Expert)***
- ***Empowers Developers and DBAs to proactively tune queries for improved SQL performance***
- ***Reduces the complexity and manual effort required to perform query tuning***
- ***Decreases the reliance on specialized Query Optimization and SQL tuning skills***
- ***Provides a rich set of SQL tuning tools and design advisors that can be used over Data Lifetime***
- ***Eclipse Based with integration and shell sharing with:***
  - *InfoSphere data Architect*
  - *Optim Development Studio (with pureQuery)*
  - *Optim Database Administrator*
- ***Support both DB2 for z/OS and LUW***

Optim Query Tuner Key Features at a Glance

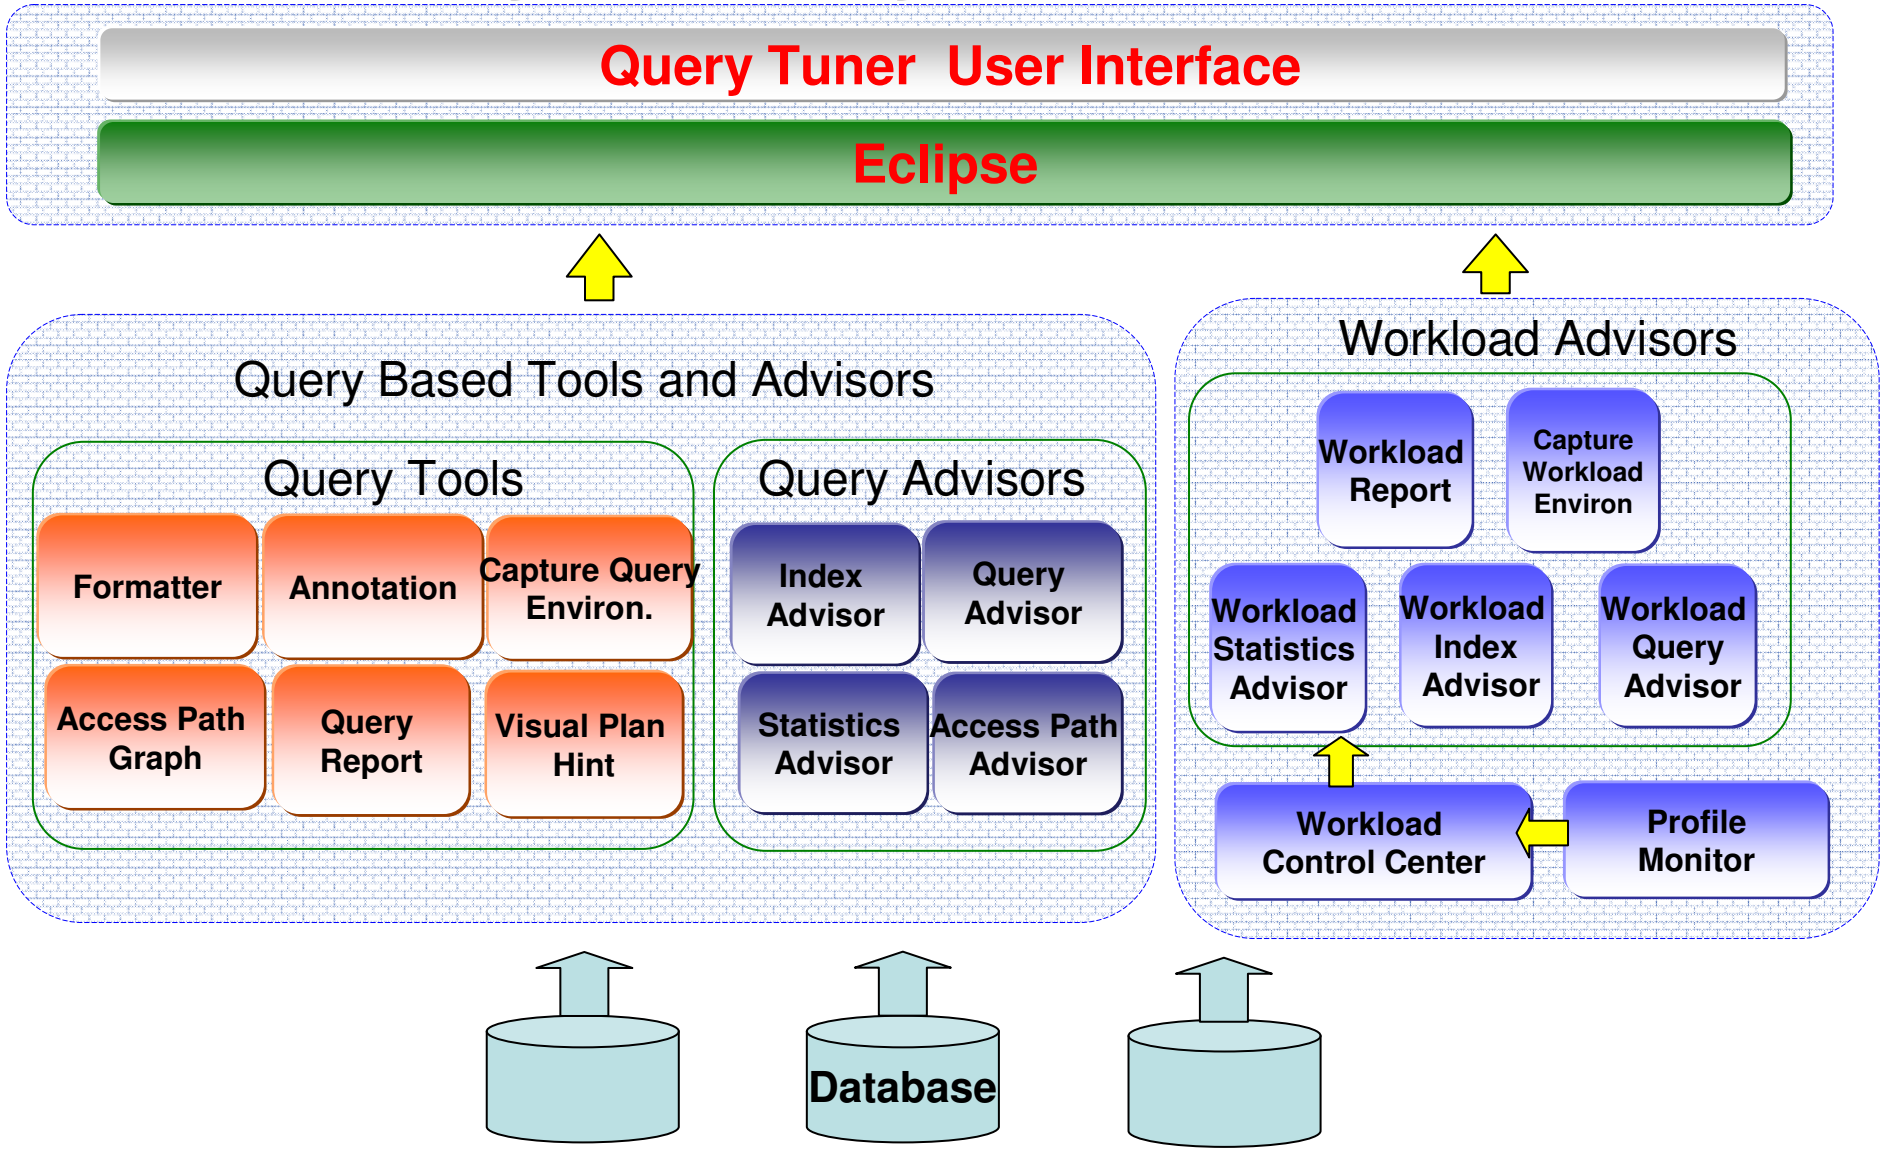

IBM Optim Query Tuner key functions

Functions	Query Tuner for DB2 for z/OS and LUW (single Query)	Query Workload Tuner for DB2 for z/OS
Query Formatter	Yes	Yes
Query Annotation	Yes	Yes
Access Plan Graph	Yes	Yes
Visual Plan Hint**	Yes	Yes
Query Advisor	Yes	Yes
Access Path Advisor	Yes	Yes
Statistics Advisor	Yes	Yes
Index Advisor	Yes	Yes
Query Reports	Yes	Yes
Query Environment Capture**	Yes	Yes
Workload Query Advisor		Yes
Workload Statistics Advisor		Yes
Workload Index Advisor		Yes
Workload Query Reports		Yes
Workload Environment Capture		Yes
Profile Based Monitor *		Yes

- *DB2 for z/OS V9.1 NFM
- ** functions for DB2 for z/OS

IBM Optim Query Tuner

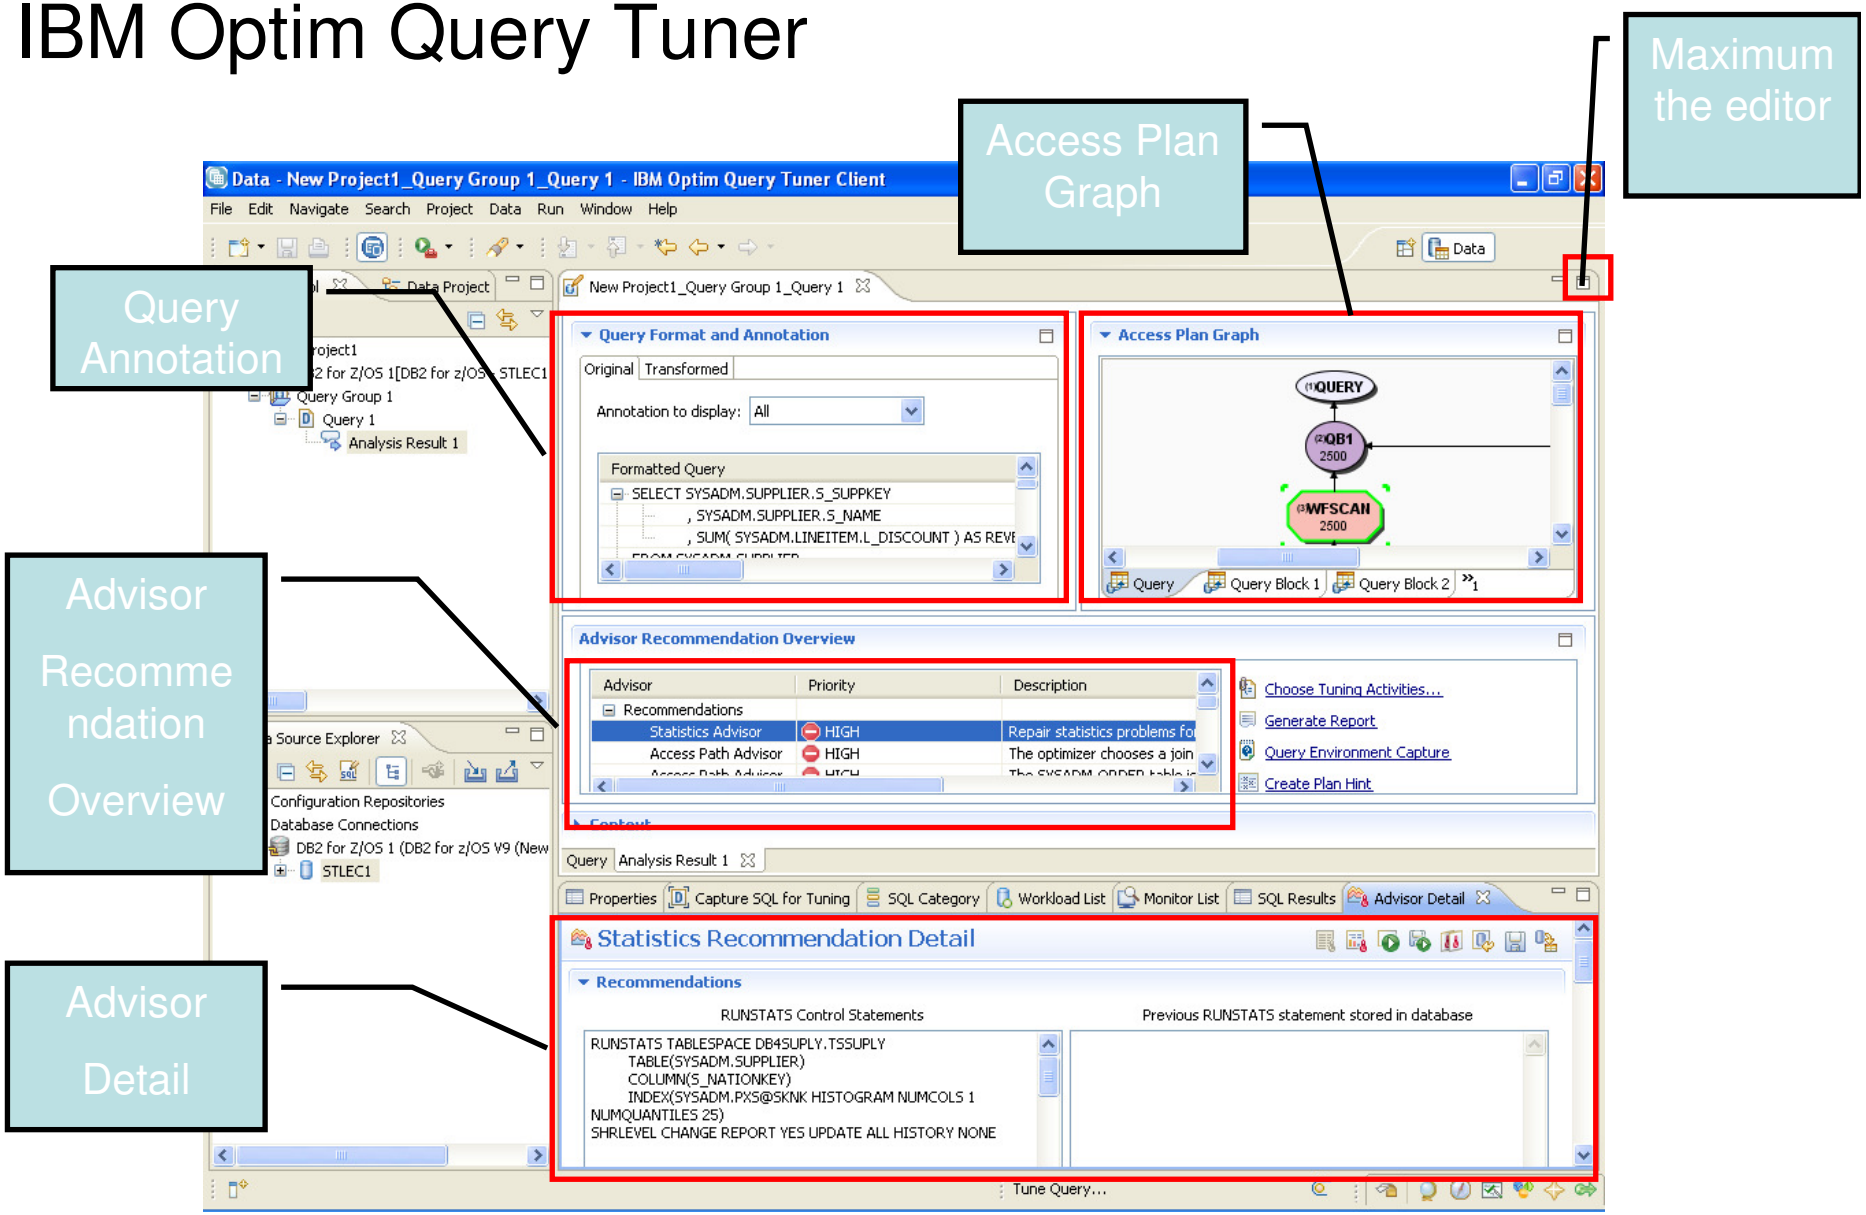

Accelerate SQL Performance with OMPE and Optim Development Studio

- **Omegamon XE (z)**

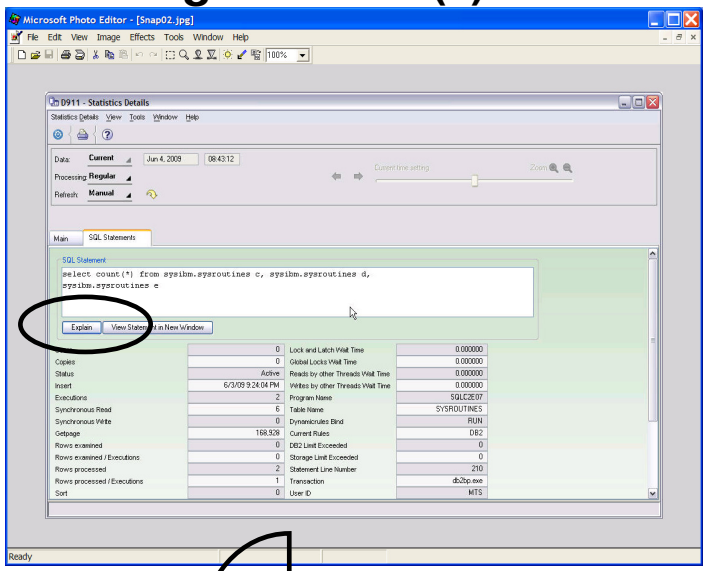

- **Optim Development Studio**

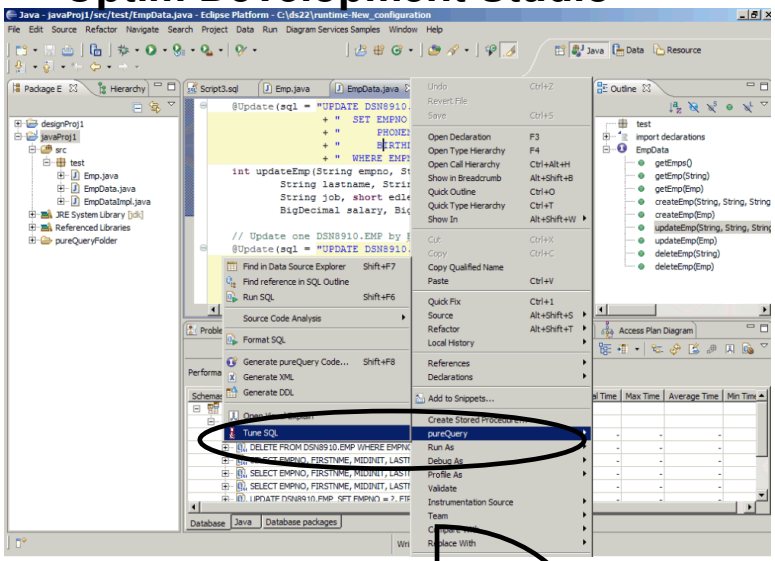

- **Optim Query Tuner**

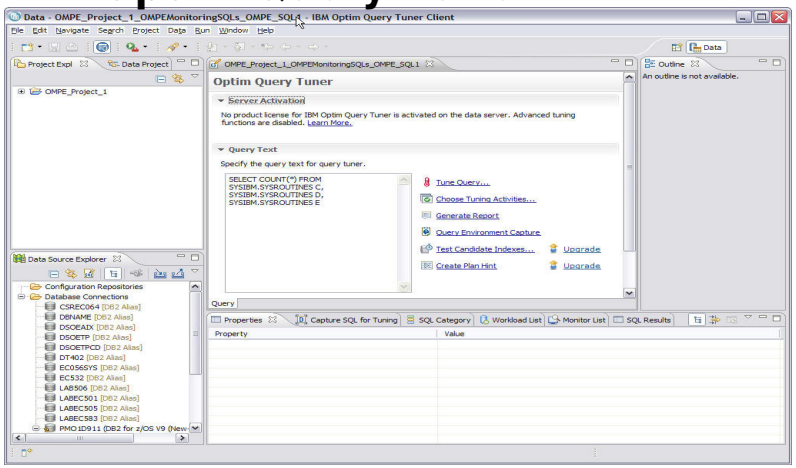

- OPME launch
- OQT to tune
- SQL performance

- ODS Launch OQT
- For application tuning

Use IBM Optim Query Tuner/Workload Query Tuner Solutions

- **Identify query/workload candidates**
  - DB2 Catalog,
  - Dynamic statement Cache,
  - Text, File, package, QMF and more
- **Prevent problems before they impact the business**
  - Get early warning from Statistics advice, Access path advice, Query advice, index advice of emergent problems
  - Isolate problems quickly
  - Identify and optimize high cost queries proactively from workload Query Tuner tooling and advices
- **Improves quality of service/application development**
  - Use expert advice for performance optimization
  - Accelerate responses to performances issues
  - Enable fixes in applications before run on production systems
- **Reduce costs**
  - Increase capacity of existing systems
  - Free up DBA time for value creation activities
  - Optimize SQL in development while costs and impact are low

Why Query Formatting/Annotation?

A peek at an unformatted Query

```
SELECT B.BRANCH_NO, B.BRANCH_NAME ,B.BRANCH_ACCT_NO,
B.BRANCH_CITY, B.BRANCH_ST, A.BRANCH_ADDR_TYPE_CD ,
S.TRANS_SETL_AMT FROM TRANS_SETLMNT S ,BRANCH C ,
BRANCH_ADDR A
WHERE S.TRANS_NO = ? AND S.TRANS_PROC_DT < '9999-12-31‘
AND  YEAR(S.TRANS_TARGET_DT) = ‘2002’ S.TRANS_TYPE  IN
(‘A1', ‘A2', ‘A3', ‘Z9')
AND S.TRANS_CD IN  ('EOD', 'IMD', ‘UGT') AND
S.TRANS_SETL_DT = ? AND S.BRANCH_NO = C.BRANCH_NO AND
B.BRANCH_EFF_DT <= ? AND B.BRANCH_INACTIVE_DT >  ? AND
A.BRANCH_NO = C.BRANCH_NO AND A.BRANCH_EFF_DT    <= ?
 AND A.BRANCH_INACTIVE_DT    >  ?  AND
A.BRANCH_ADDR_TYPE_CD   = ' '
```

Query Annotation - Annotate a statement

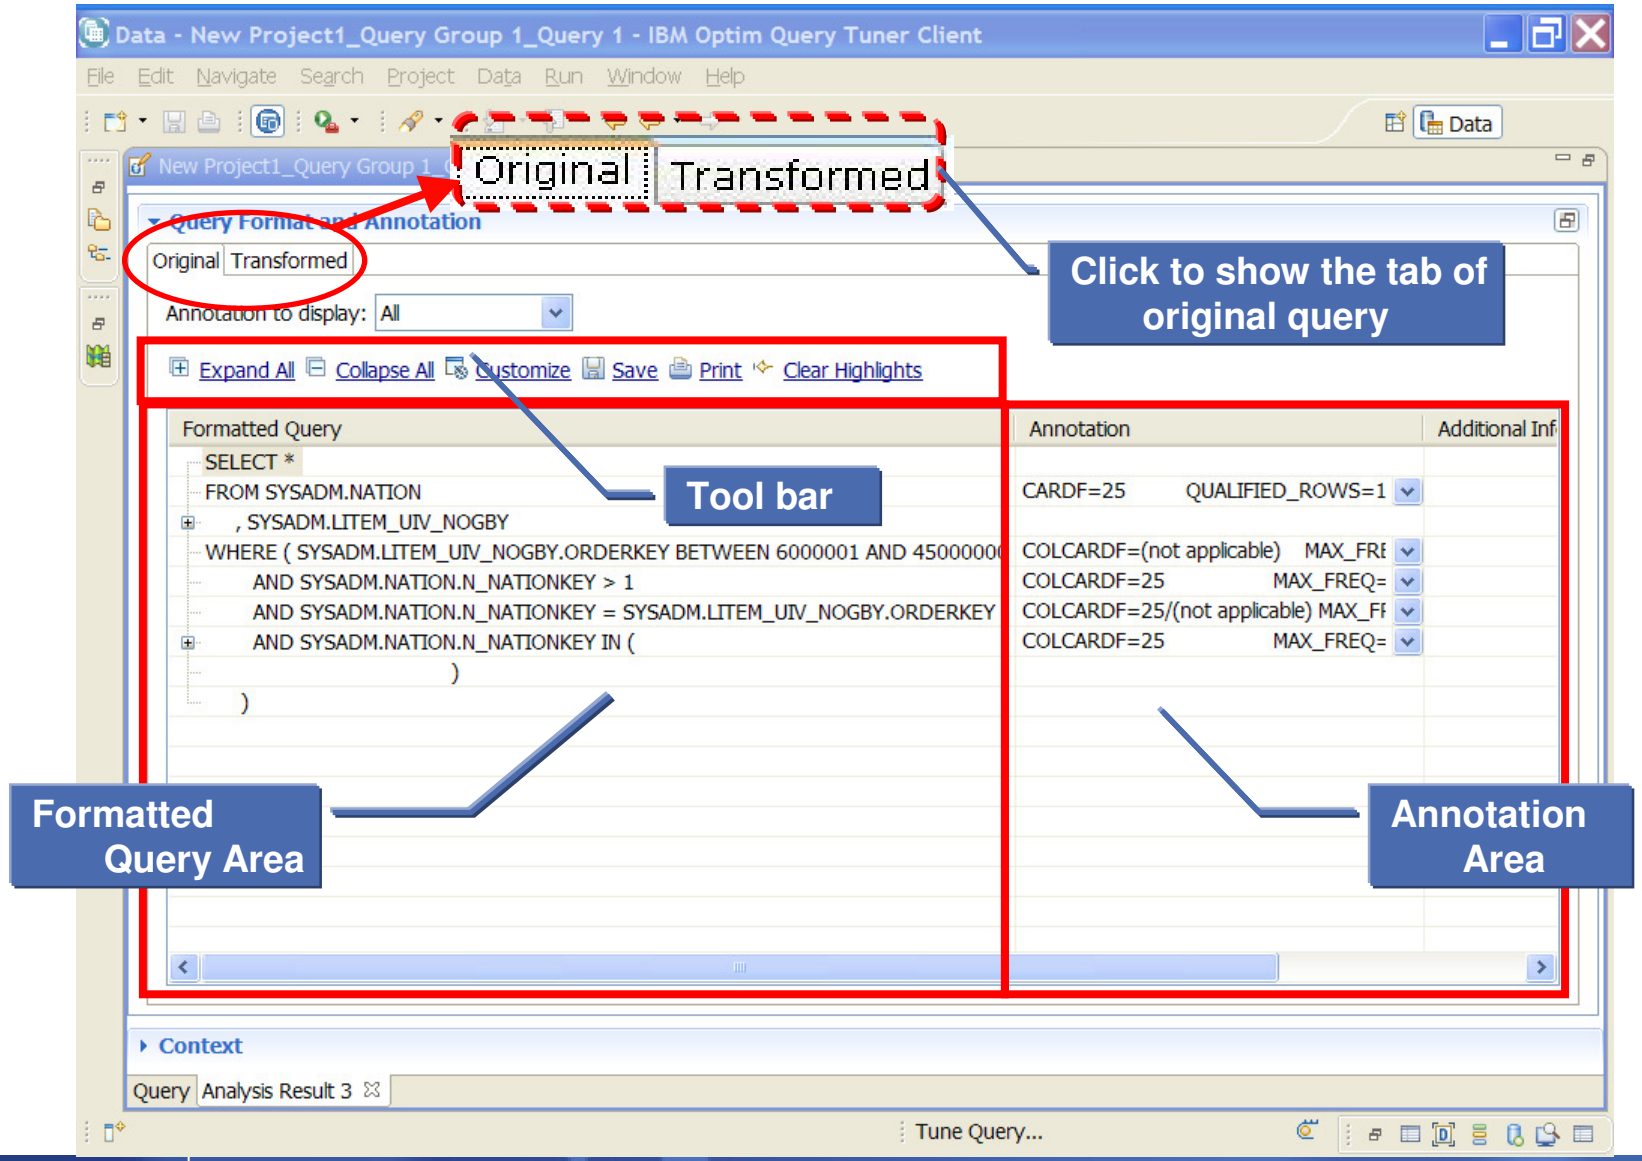

Query Annotation - Transformed Query Annotation

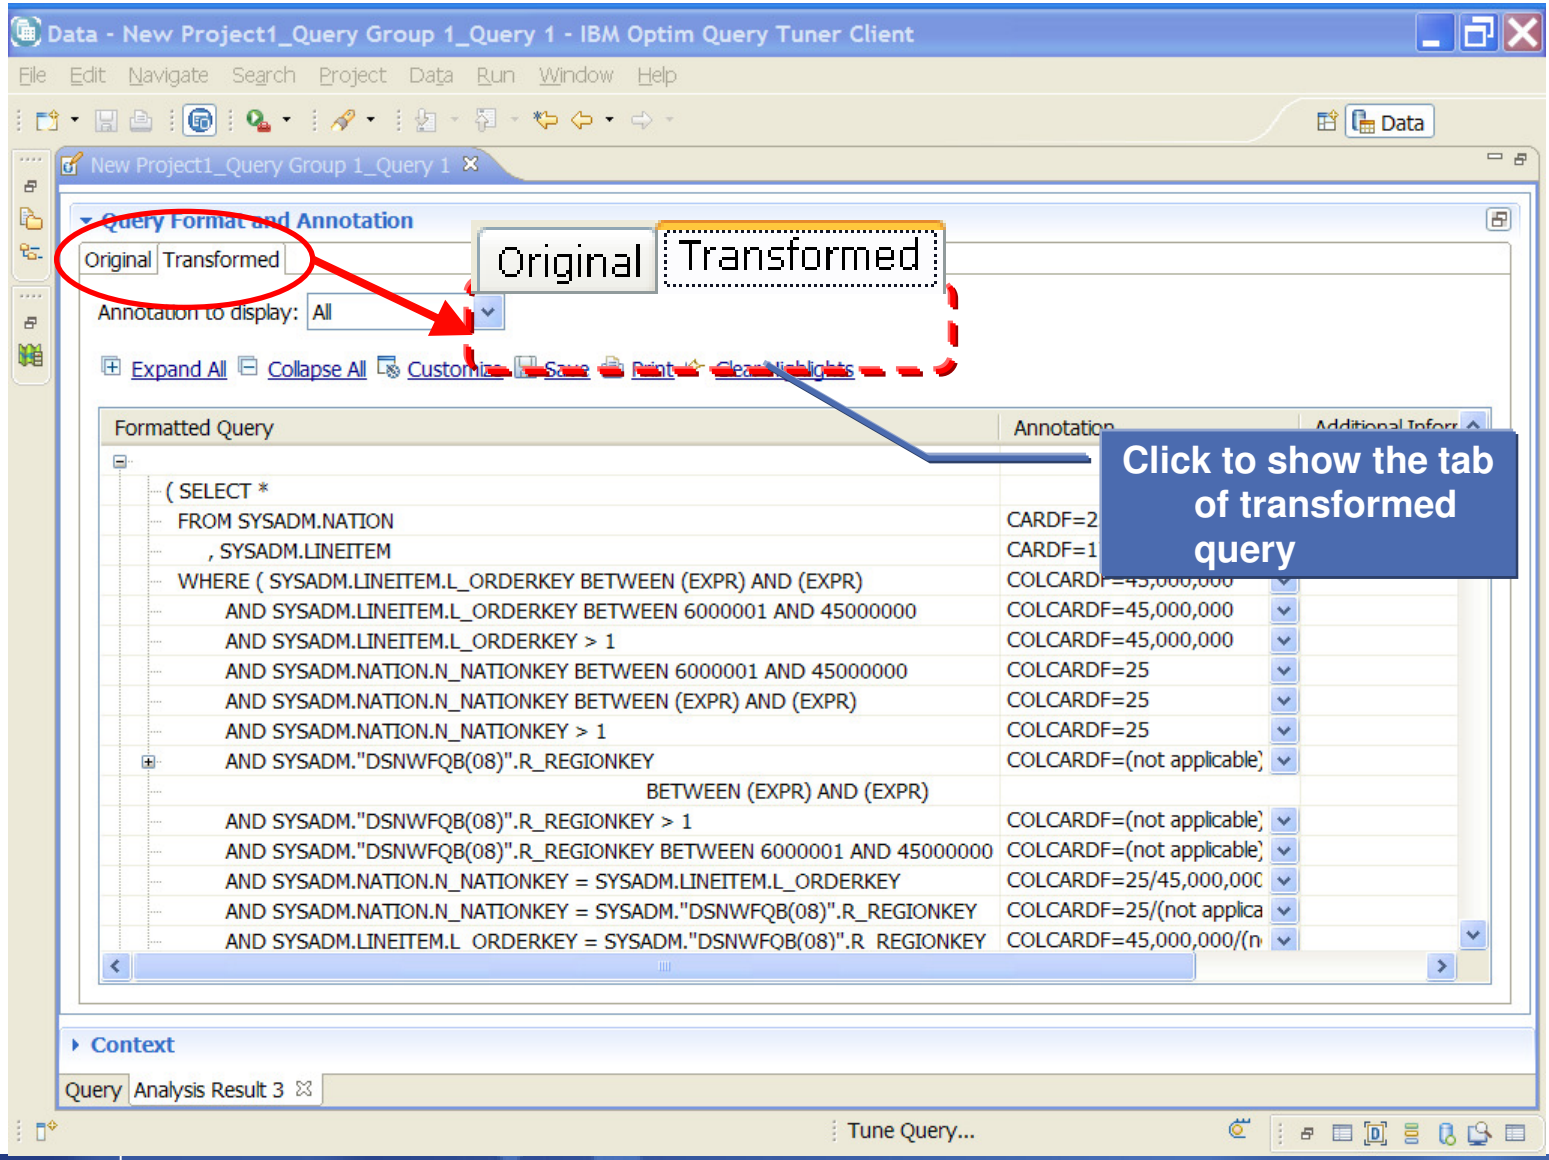

Stats Advisor – Recommendation Summary

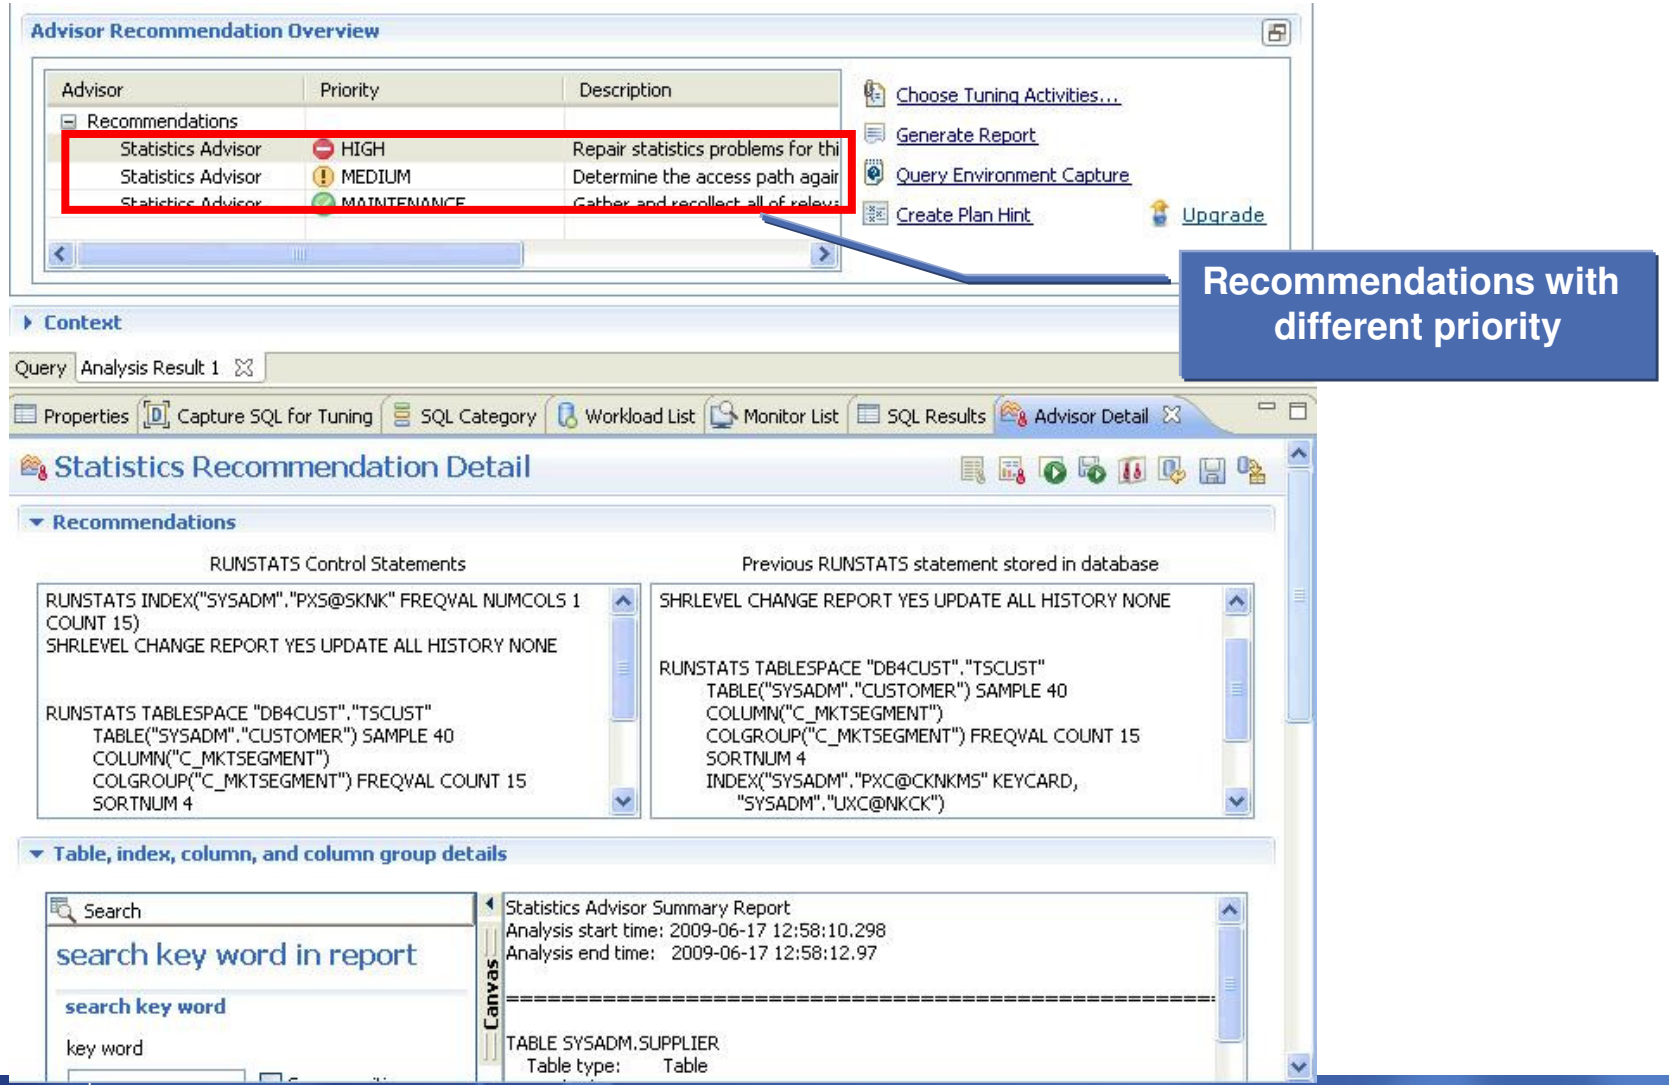

Access Path Advisor – Access Path Warning

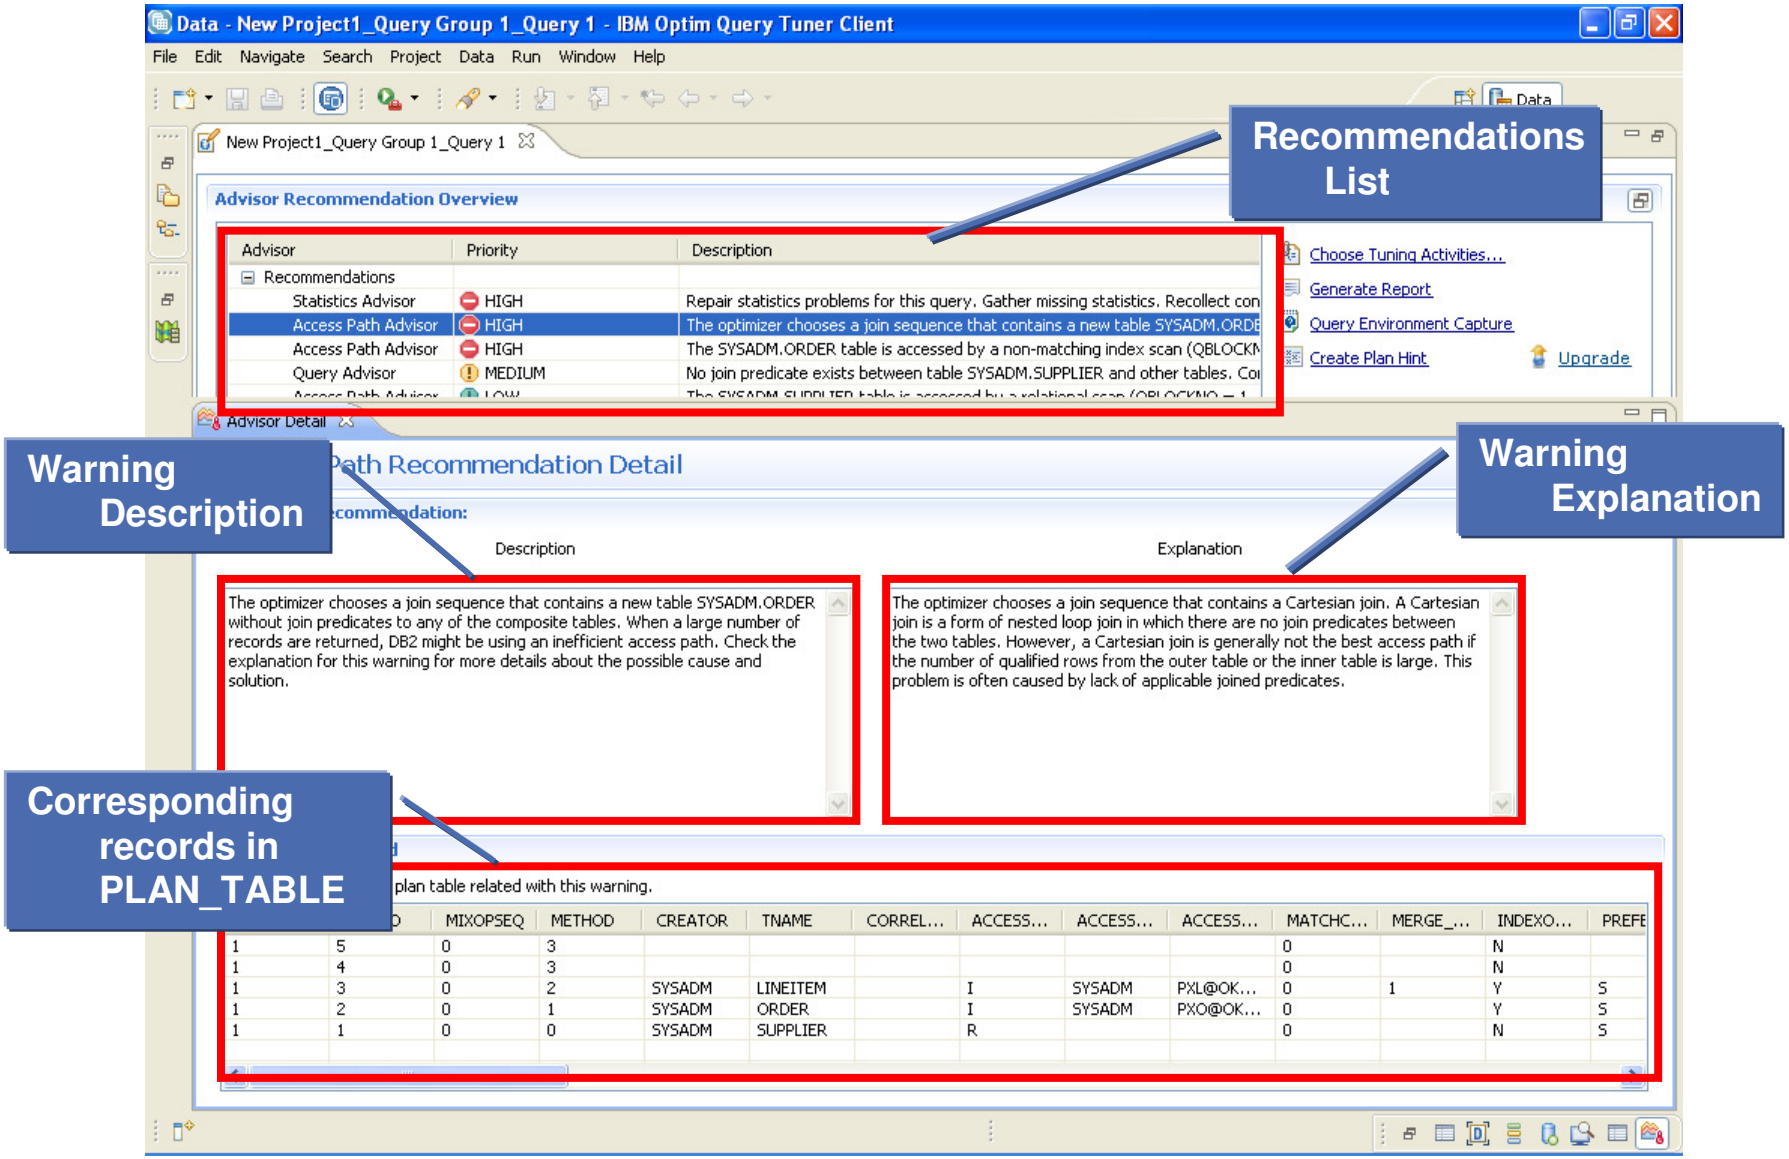

Query Advisor - Recommendations

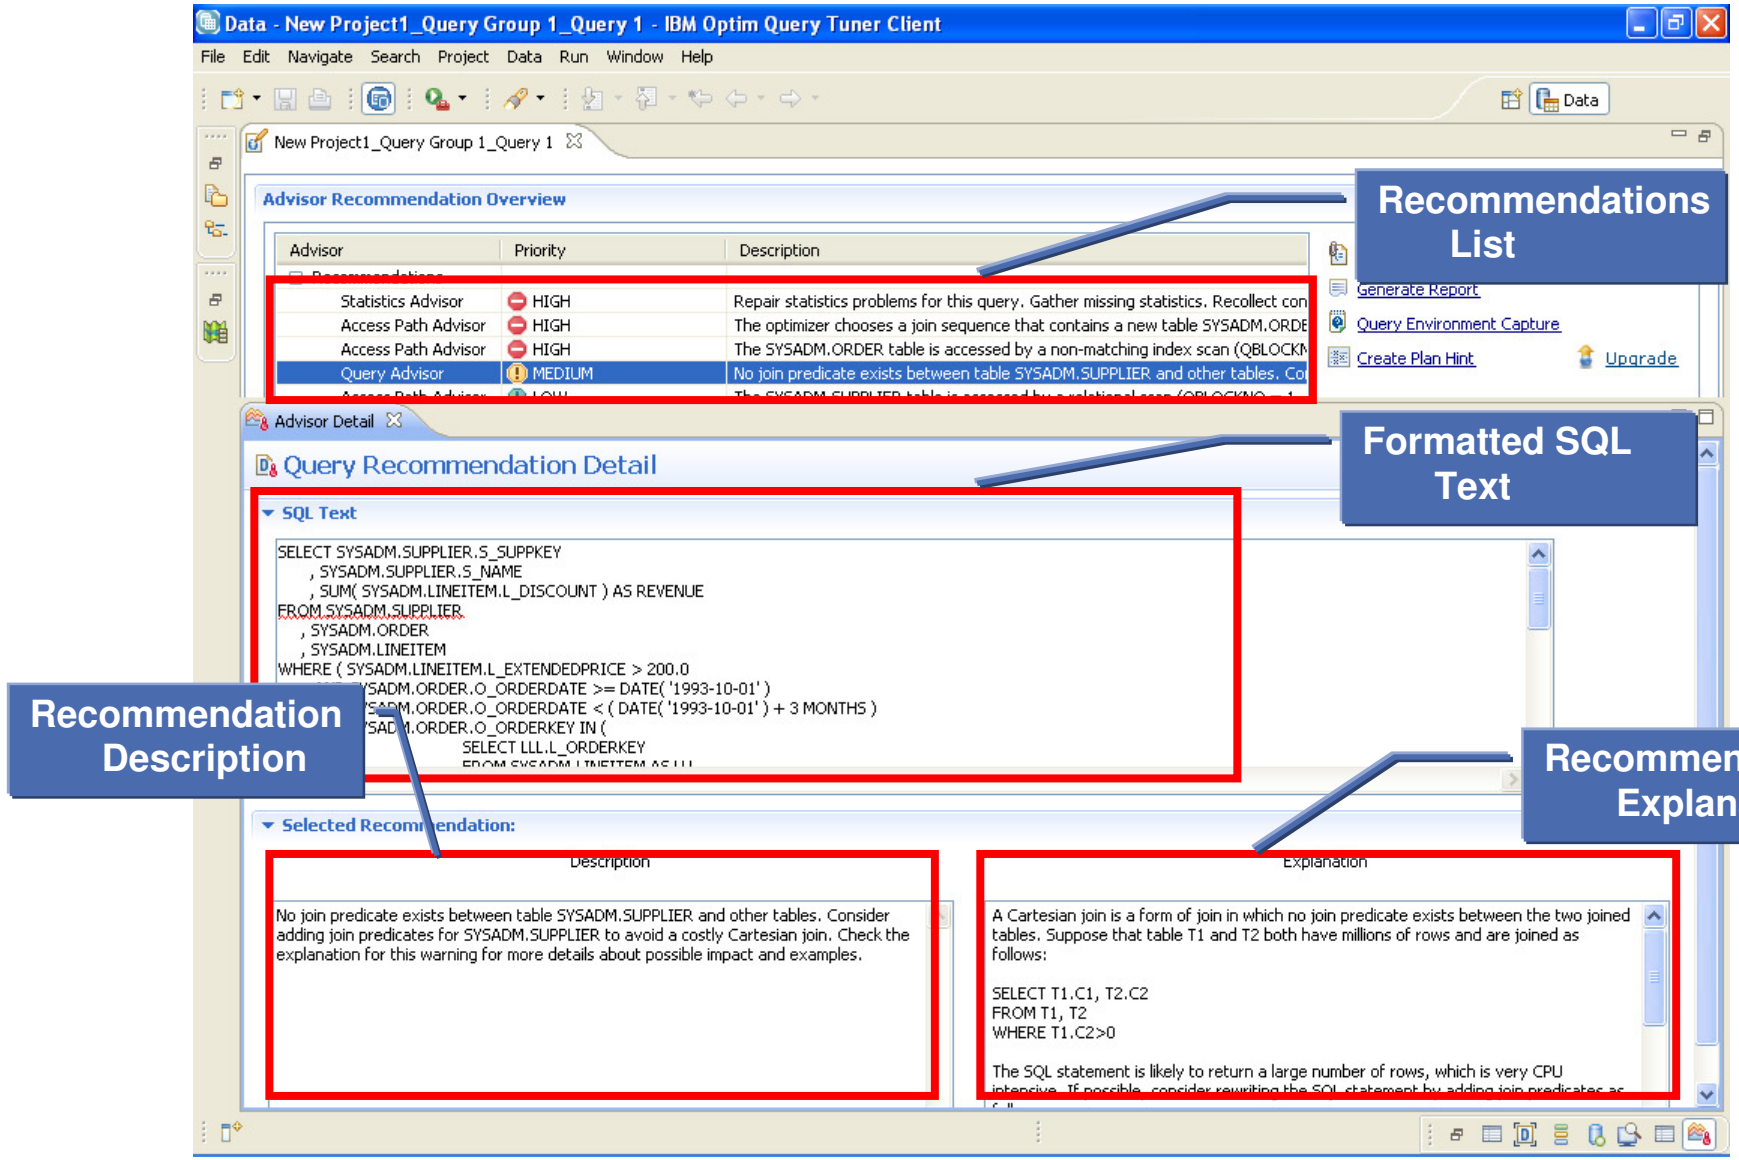

Index Advisor - Recommendations

- Query Tuner Index Advisor – with advice on indexes with estimated performance improvement below.

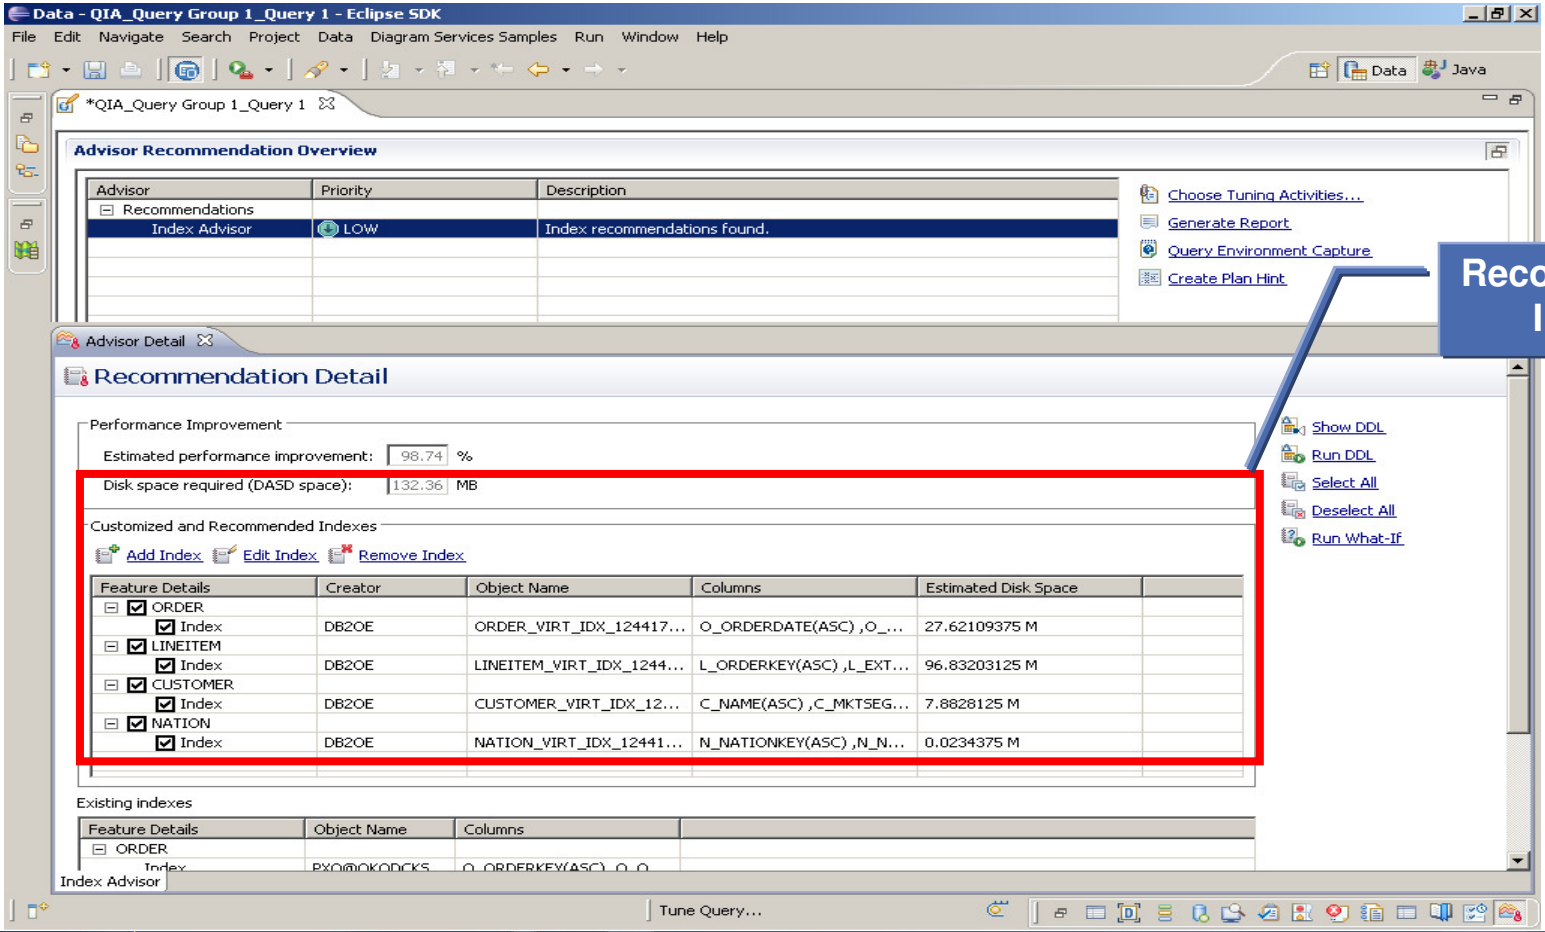

Access Path Graph Comparison

- **Single query APG Comparison of 2 SQL Tuning analysis results.**
- **Can be from the same project or different project**
- **Can be from same server or different server. For example: Tuning analysis APG comparison from DB2 z/OS and LUW**

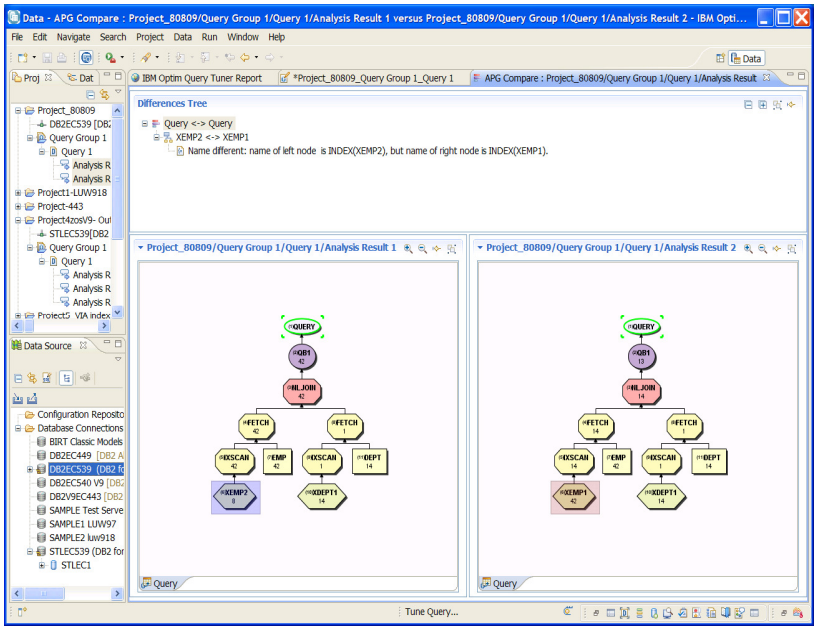

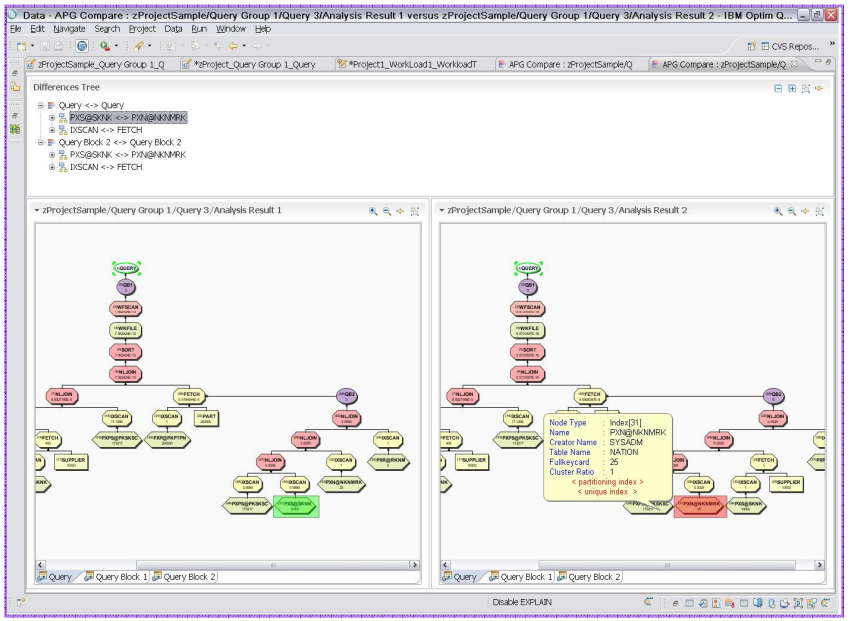

Access Path Graph Comparison

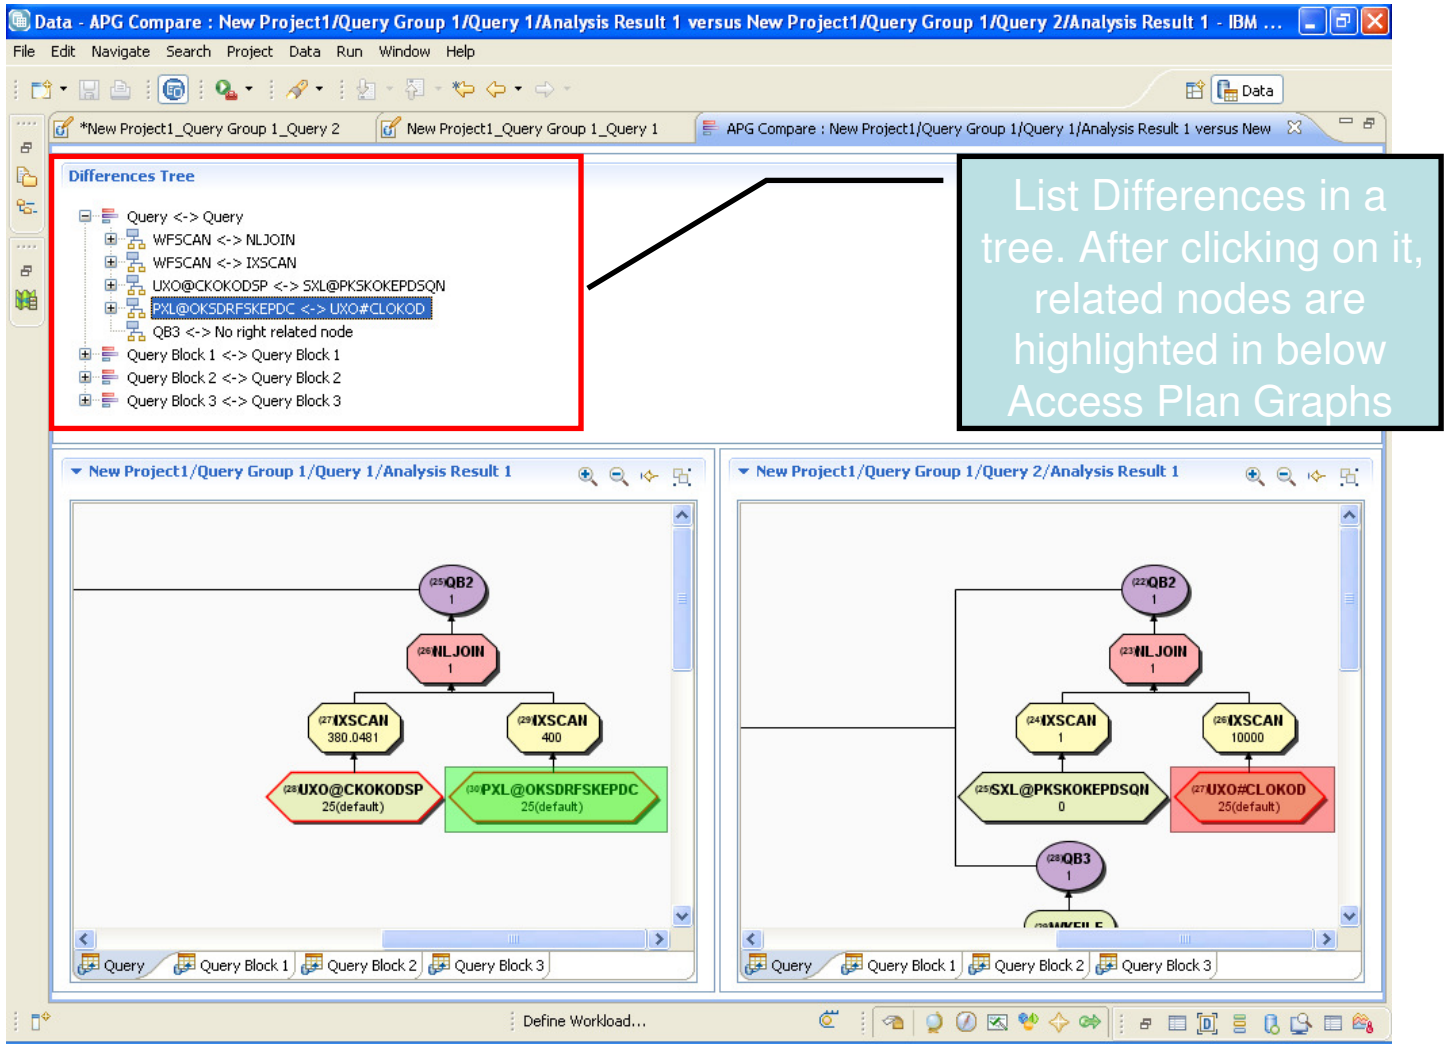

Visual Plan Hint

- A example Query with Join seq. change (emp->dept to dept-> emp)

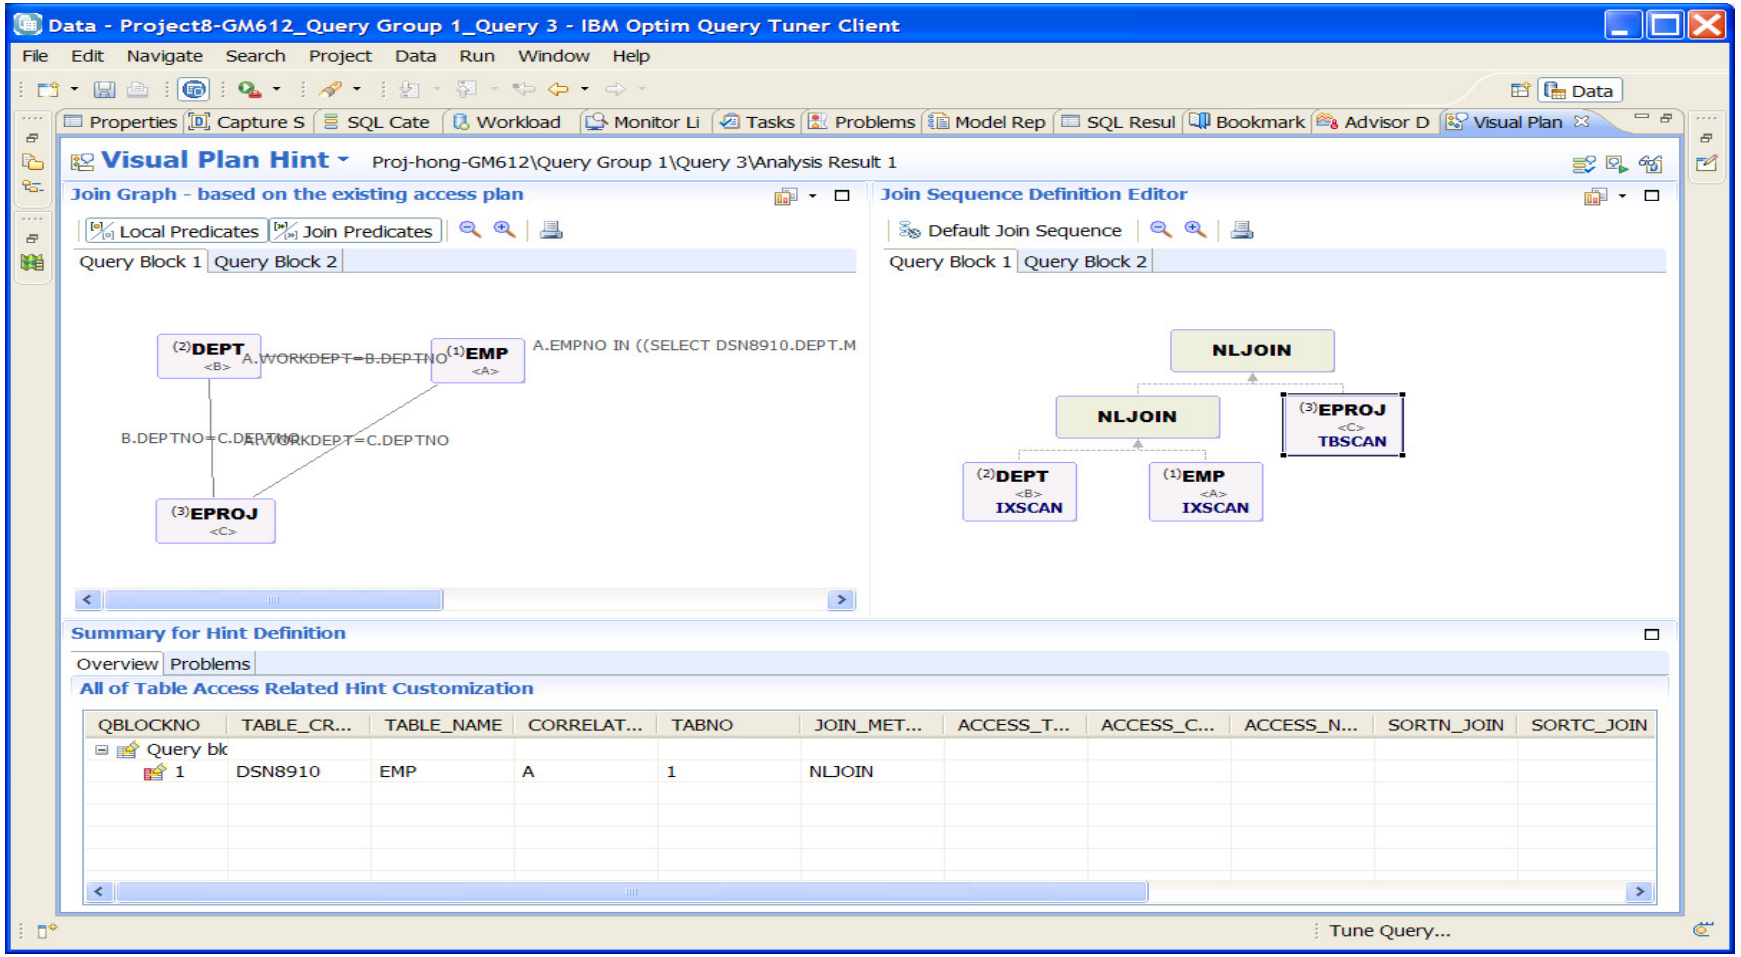

Workload Index Adviser -Recommendations

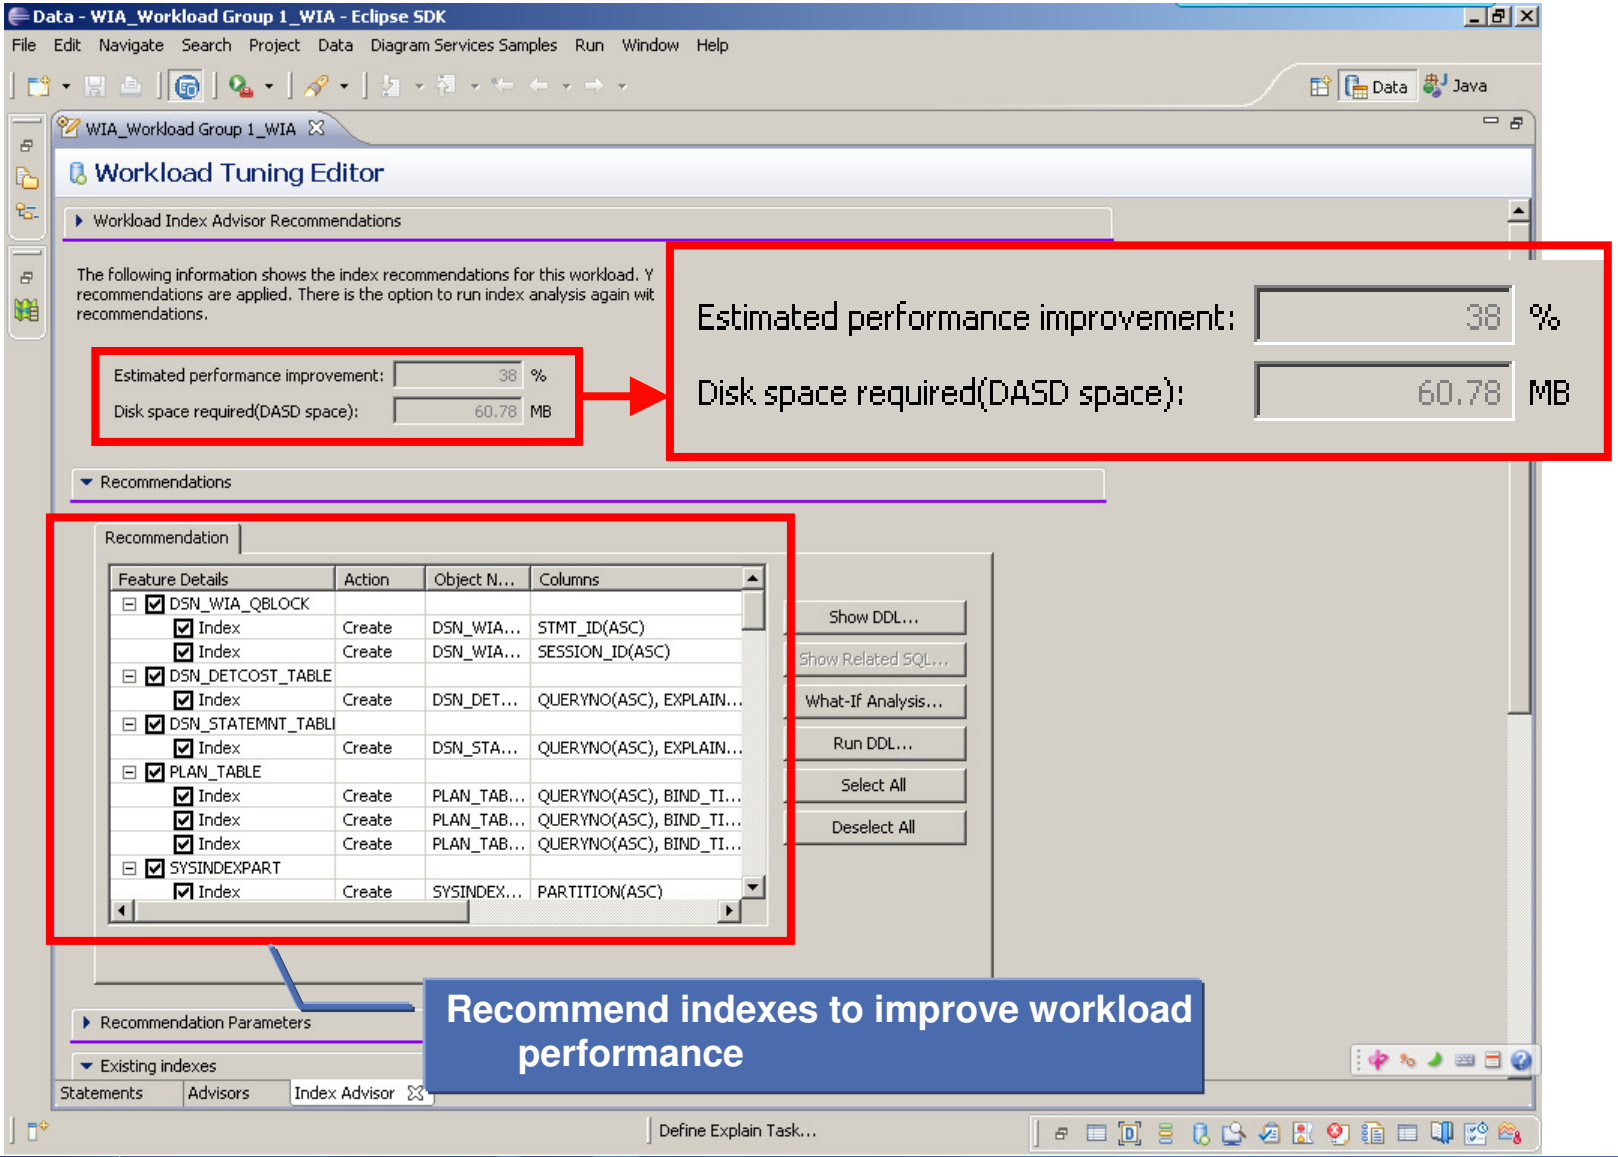

IBM Information Management

Optimizing Your Java Applications with Optim Tools

- Optim Development Studio, Optim pureQuery Runtime, and Optim Query Tuner

Optim pureQuery Runtime – every Java application benefits!

- **JDBC – acceleration for any JDBC application**
  - Convert dynamic SQL to static SQL
  - Replace problem queries without changing the source
  - Remove literals from SQL to get better statement cache hit ratios
- **Hibernate/OpenJPA/iBatis – acceleration for persistence layers**
  - Improved SQL “batch” performance
  - Auto-tuning of Hibernate and OpenJPA persistence options
- **SQL-friendly APIs for OO access to relational**
  - Object to relational mapping
  - APIs that can be tailored to return XML, JSON, arrays, etc.
- **Improved management, monitoring, problem determination**
  - Tracks SQL back to the Java class file and line number
  - Enables performance monitors to report by application name
- **Provides the foundation for improved developer tooling**
  - Syntax assist, code generation, performance reporting, etc.

What's so Great About DB2 Accounting for CICS Apps?

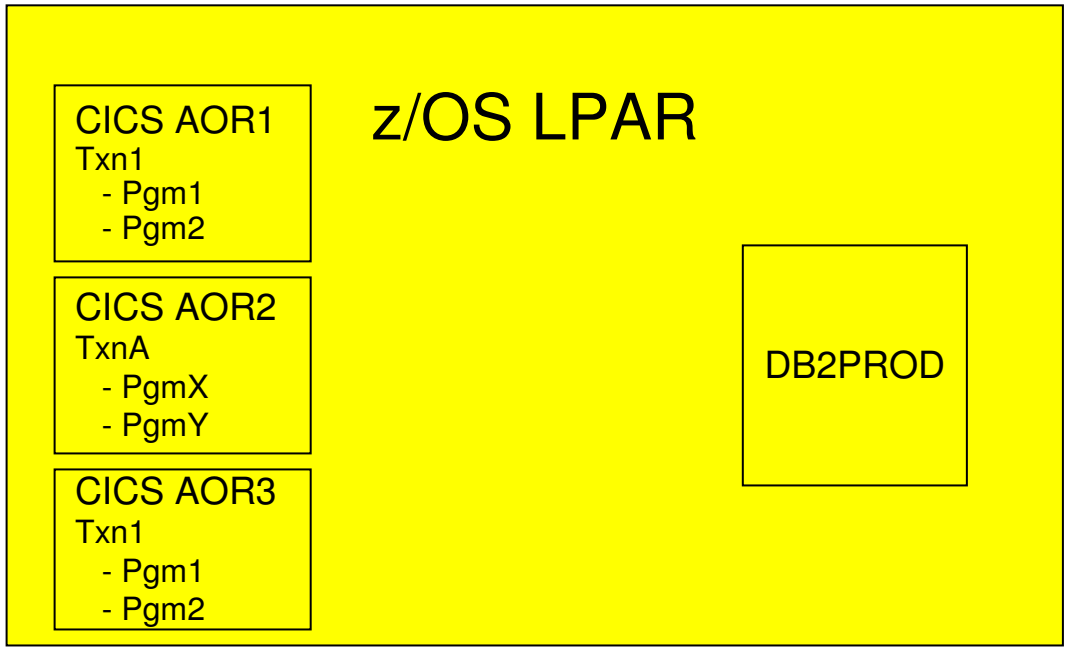

DB2 Accounting for CICS apps allows you to study performance data from many perspectives:

- By transaction (PLAN name)
- By program (package level accounting)
- By address space (AOR name)
- By end user ID (CICS thread reuse)

This flexibility makes it very easy to isolate performance problems, perform capacity planning exercises, analyze program changes for performance regression, compare one user's resource usage to another's, etc.

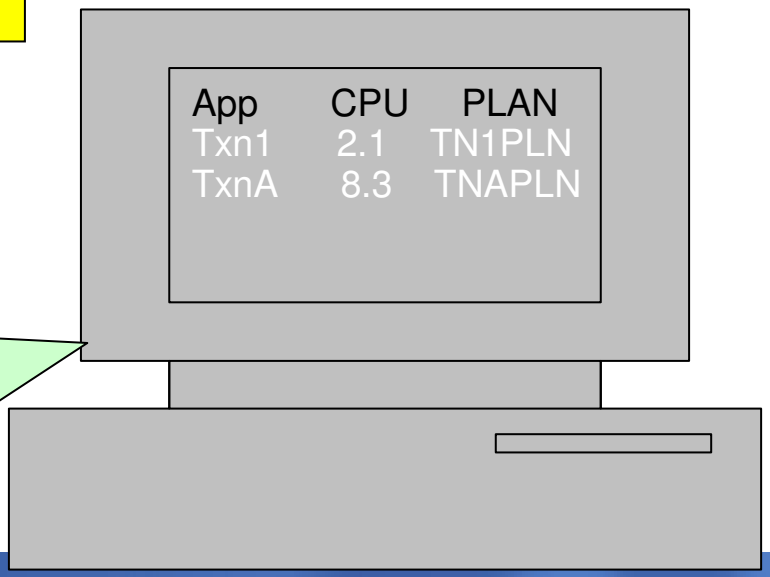

JDBC Performance Reporting and Problem Determination

– Before pureQuery

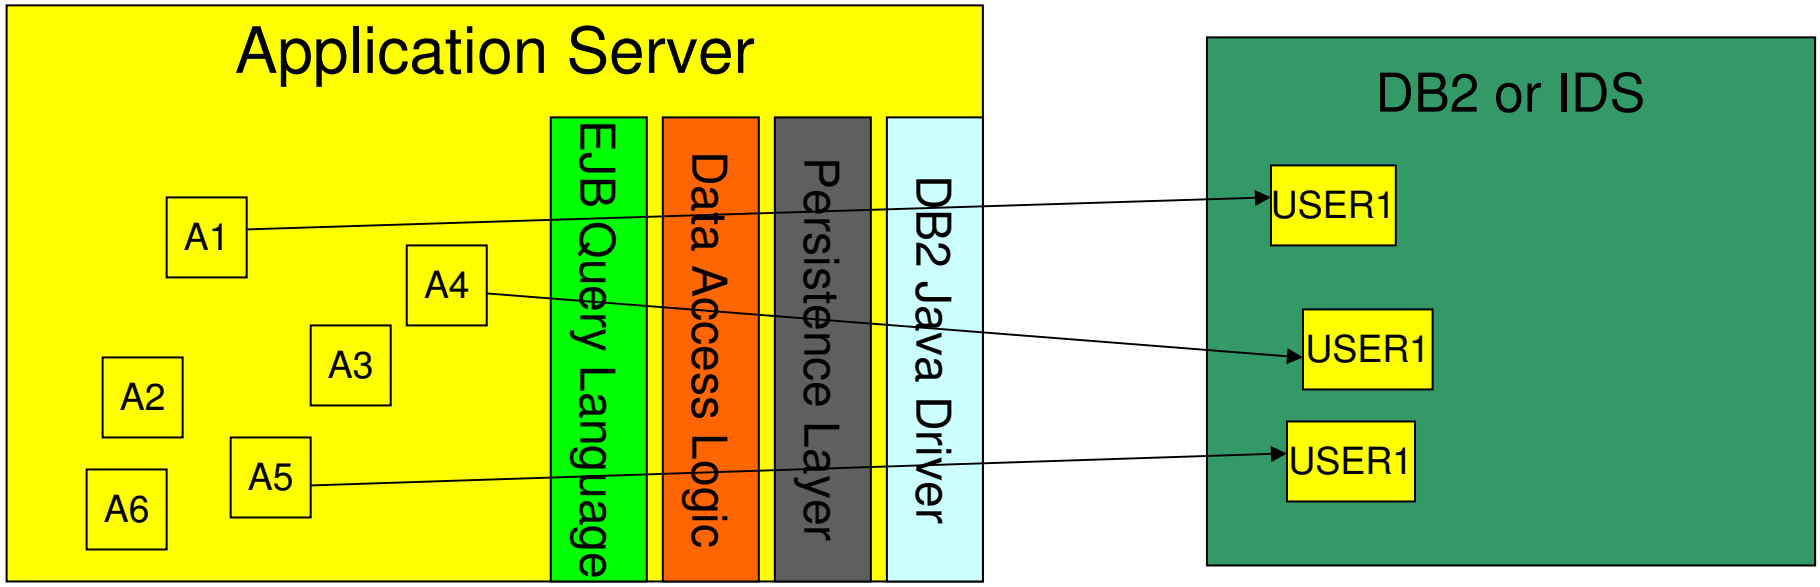

What is visible to the DBA?
- IP address of WAS app server
- Connection pooling userid for WAS
- app is running JDBC or CLI

What is not known by the DBA?
- which app 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?
- etc.

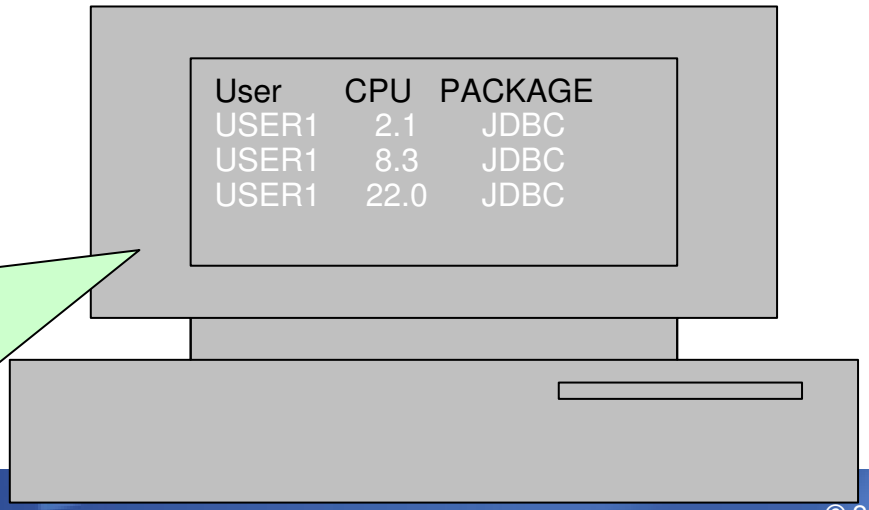

User	CPU	PACKAGE
USER1	2.1	JDBC
USER1	8.3	JDBC
USER1	22.0	JDBC

What's so Great About Optim pureQuery Accounting for WebSphere Applications?

z/OS LPAR

CICS AOR2
TxnA (PLANA)
- PgmX
- PgmY

Unix or Windows

WAS 21.22.3.4
TxnA (Set Client App=TxnA)
- ClassX
- ClassY

Optim pureQuery Accounting provides 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 Roles)

This flexibility makes it very easy to isolate performance problems, perform capacity planning exercises, analyze program changes for performance regression, compare one user's resource usage to another's, etc. just like CICS Accounting report

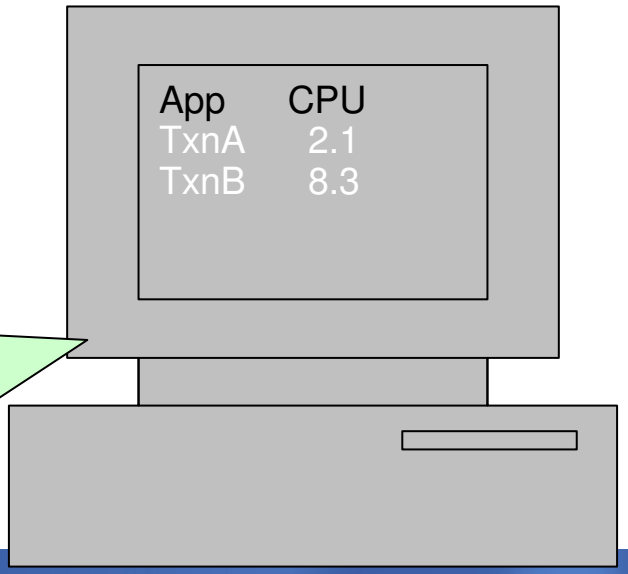

pureQuery API - Simplifying Problem Determination Scenario

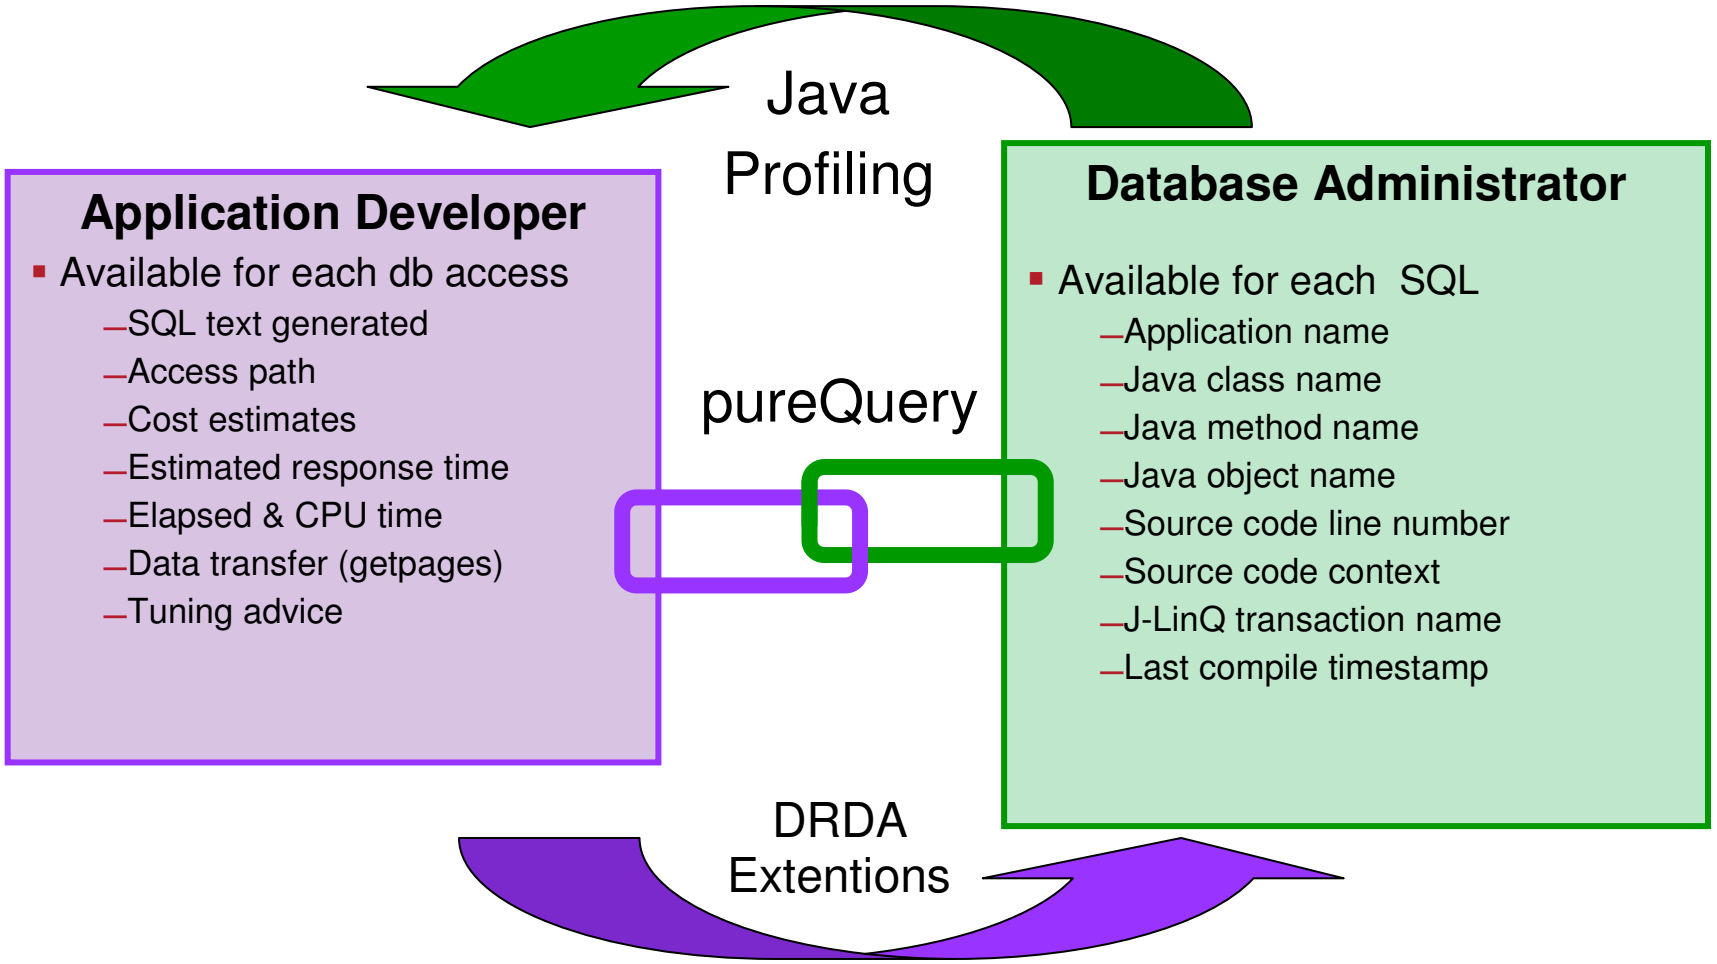

Simpler Development

pureQuery API's:

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

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

SQLJ:

```
#sql [con] { SELECT NAME, ADDRESS, PHONE_NUM INTO :name, :addr, :phone FROM EMP
        WHERE NAME=:name };
new Employee my_emp;
my_emp.setName(name);
my_emp.setHomeAddress(addr);
my_emp.setHomePhone(phone);
```

JDBC:

```
java.sql.PreparedStatement ps = con.prepareStatement(
    "SELECT NAME, ADDRESS, PHONE_NUM FROM EMP
        WHERE NAME=?");
ps.setString(1, name);
java.sql.ResultSet names = ps.executeQuery();
names.next();
new Employee my_emp;
my_emp.setName(names.getString(1));
my_emp.setHomeAddress(names.getString(2));
my_emp.setHomePhone(names.getString(3));
names.close();
```

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

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

Using pureQuery to Foster Collaboration and Produce Enterprise-ready Apps

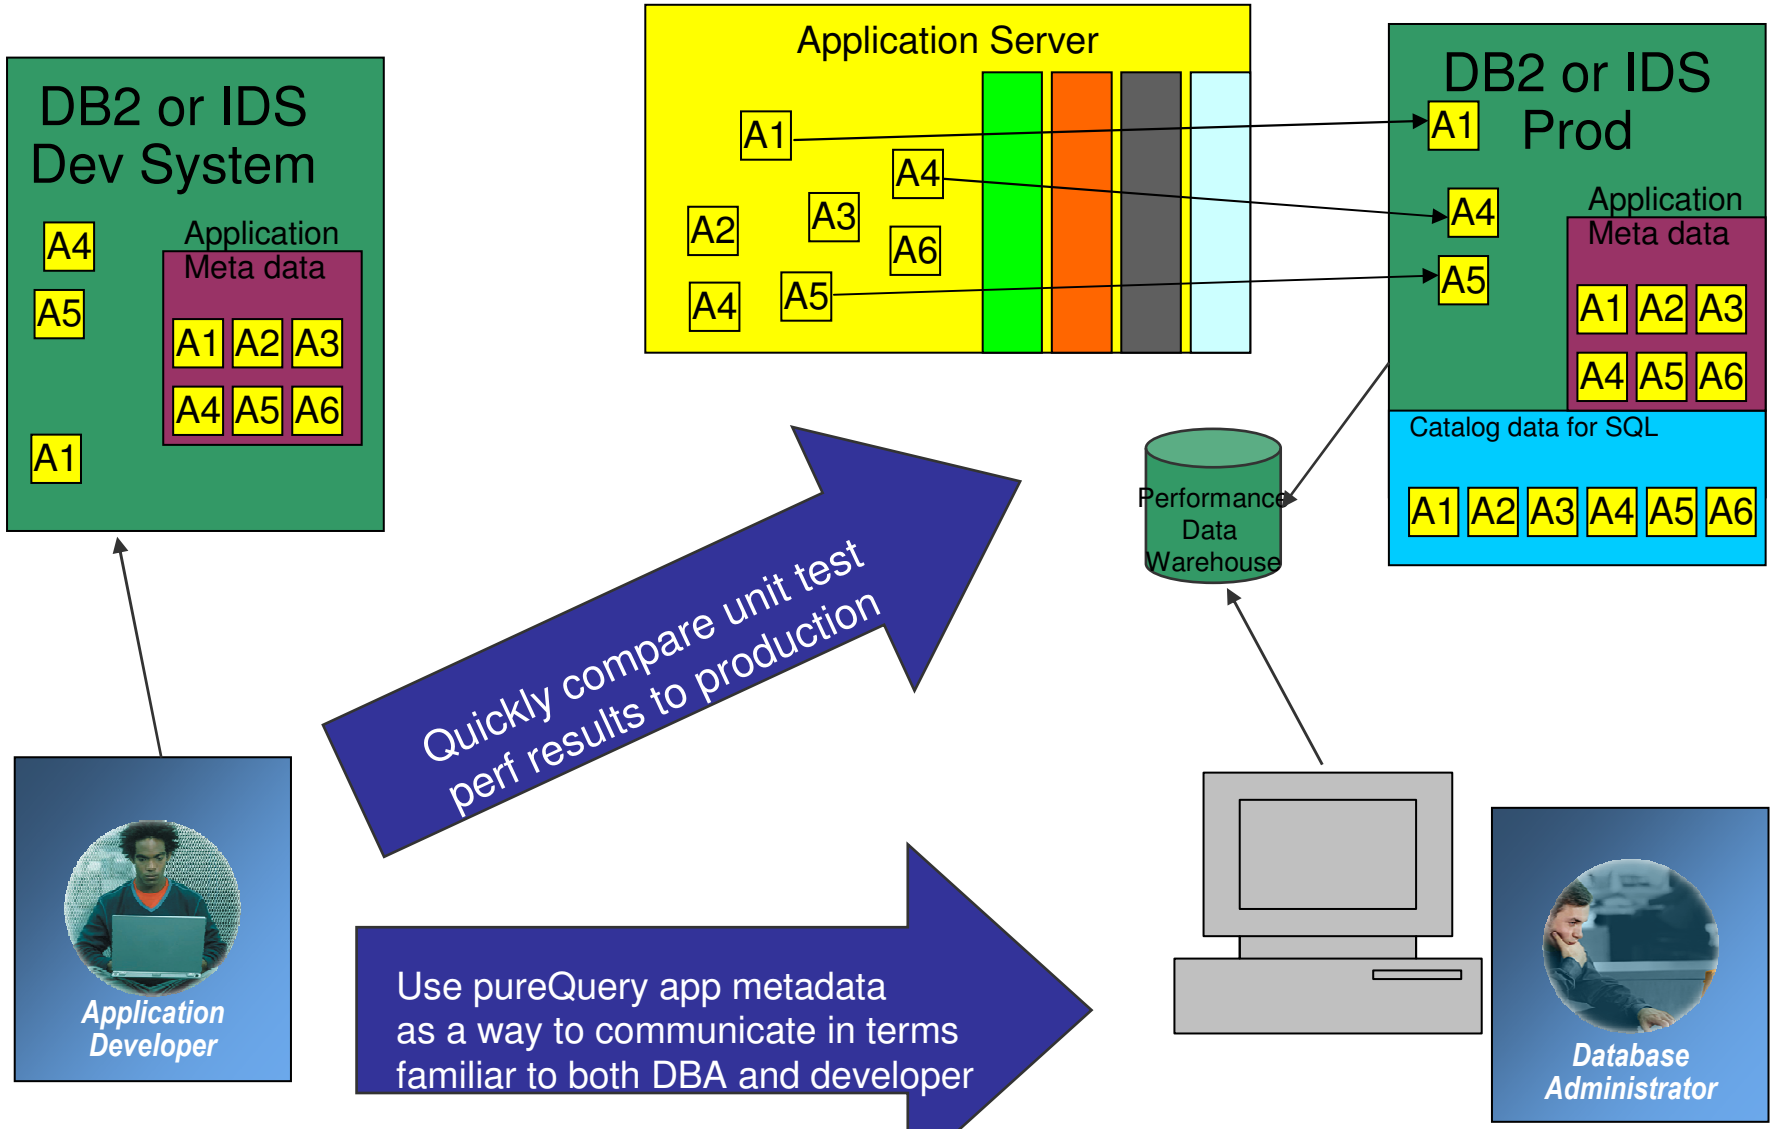

On-ramps to pureQuery

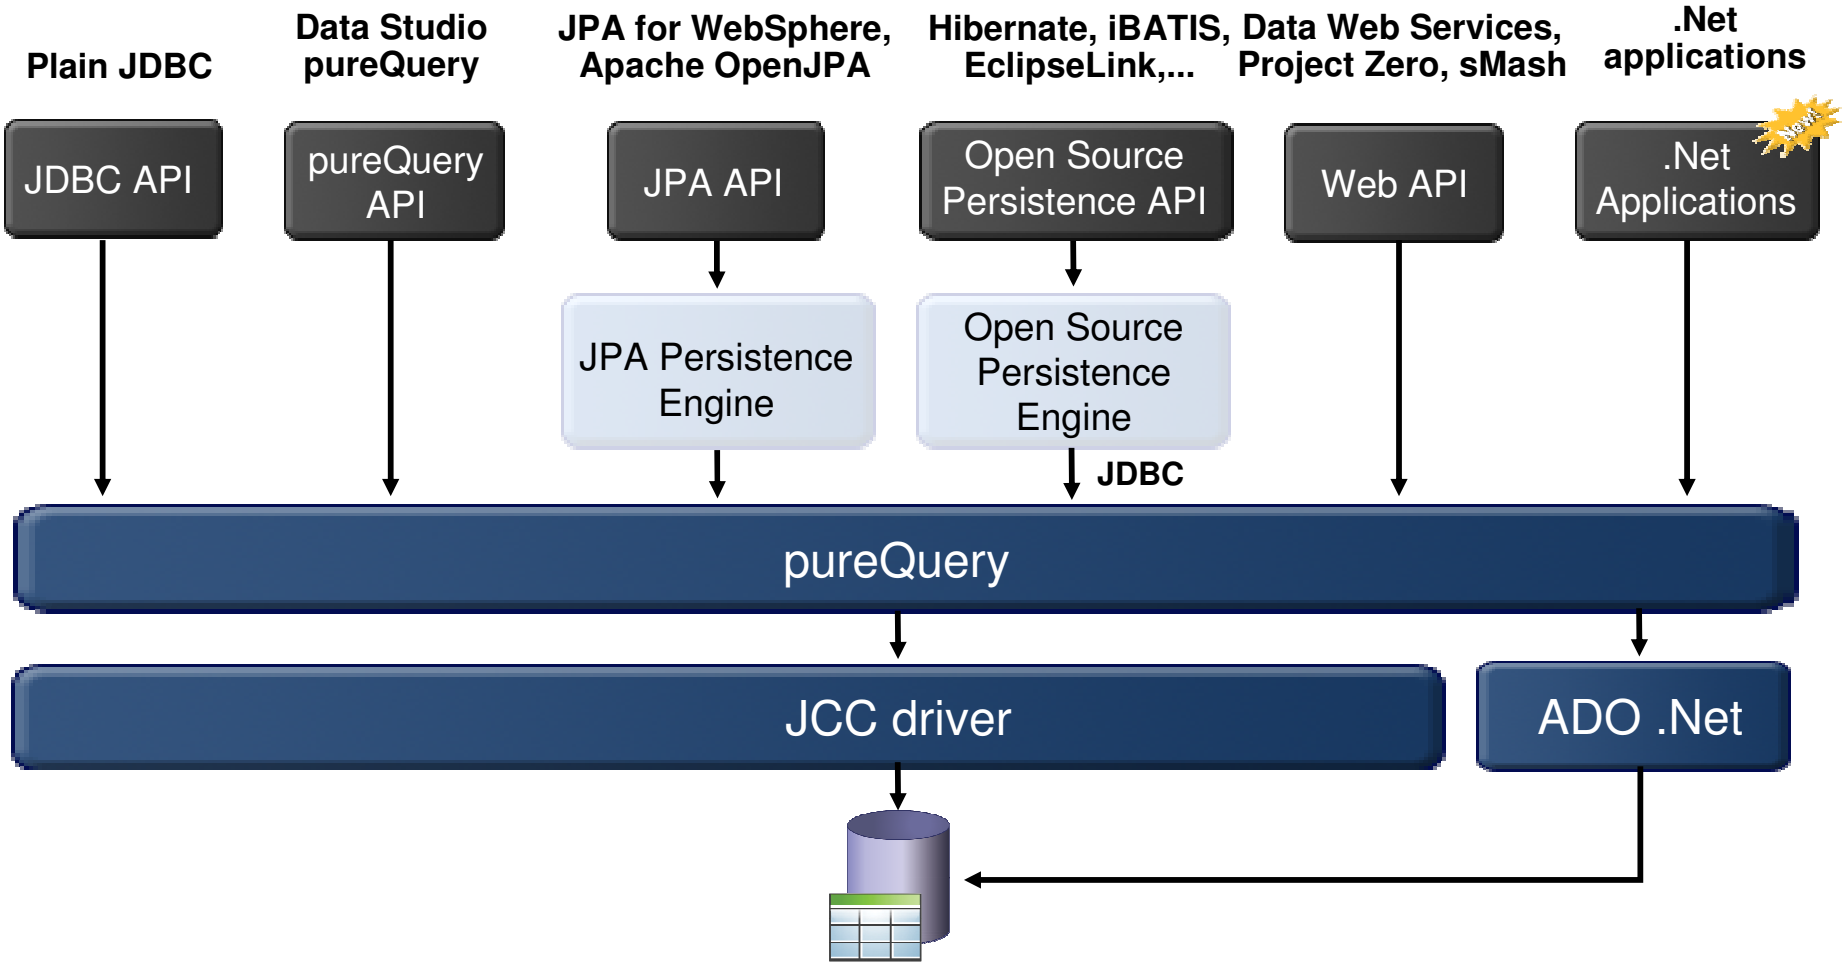

DB2, Informix, and Oracle now…more coming

Client Optimization

Improve Java data access performance for DB2 – without changing a line of code

- Captures SQL for Java applications
  - Custom-developed, framework-based, or packaged applications
- Bind the SQL for static execution without changing a line of code
  - New bind tooling included
- Delivers *static SQL execution* value to existing DB2 applications
  - Making response time predictable and stable by locking in the SQL access path pre-execution, rather than re-computing at access time
  - Limiting user access to tables by granting execute privileges on the query packages rather than access privileges on the table
  - Aiding forecasting accuracy and capacity planning by capturing additional workload information based on package statistics
  - Drive down CPU cycles to increase overall capability
- Choose between dynamic or static execution at deployment time, rather than development time

What Is Heterogeneous Batching?

Heterogenous Batching – multiple operations across different tables all execute as one batch

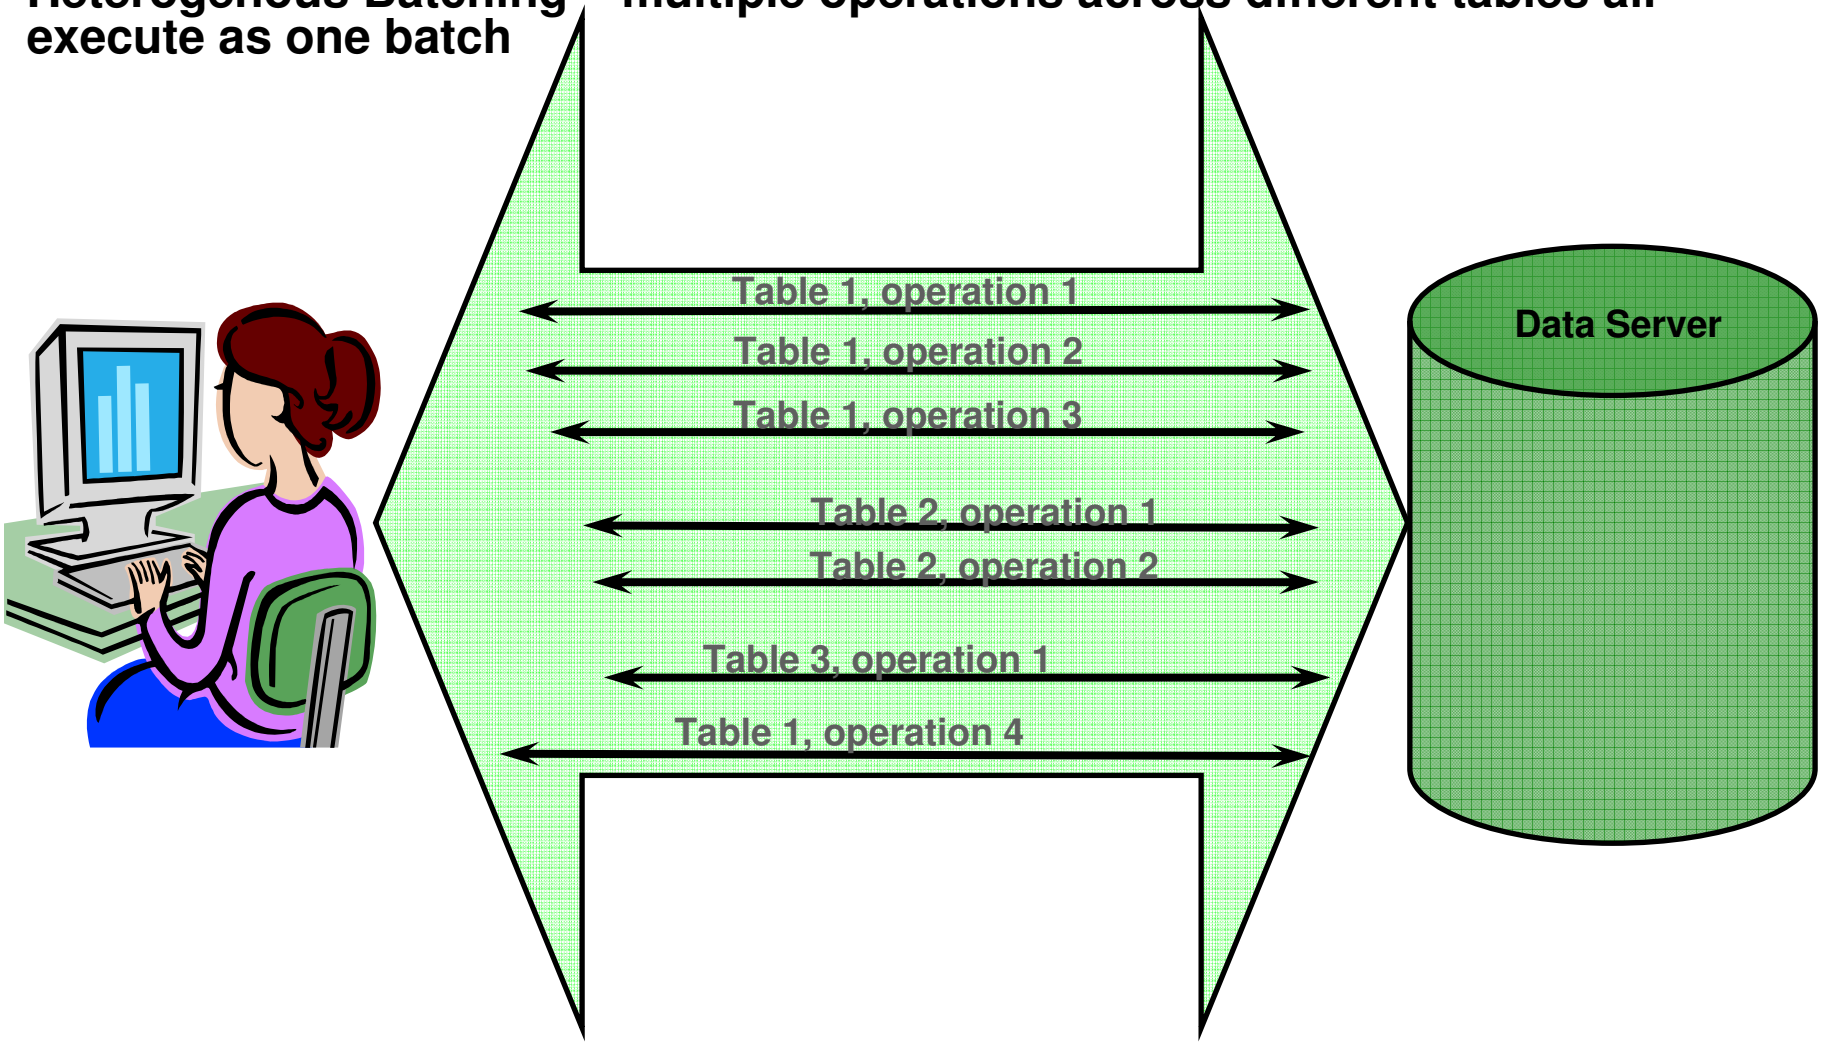

JDBC Batching v/s pureQuery Heterogeneous Batching

- JDBC batching used by Hibernate Batcher is currently limited
  - Cannot batch entities that map to multiple tables
    - Primary and Secondary tables.
    - Inheritance Join and Table per class strategies
  - Cannot batch different operations against same table
    - Field level updates
    - Insert, update
  - Cannot batch different entities
  - Each batch is a message to the database
- pureQuery heterogeneous batching plug-in for Hibernate on the other hand
  - Can batch entities that map to multiple tables
  - Can batch different operations against the same table
  - Can batch different entities into a single batch
  - Combines insert, deletes, updates into single batch

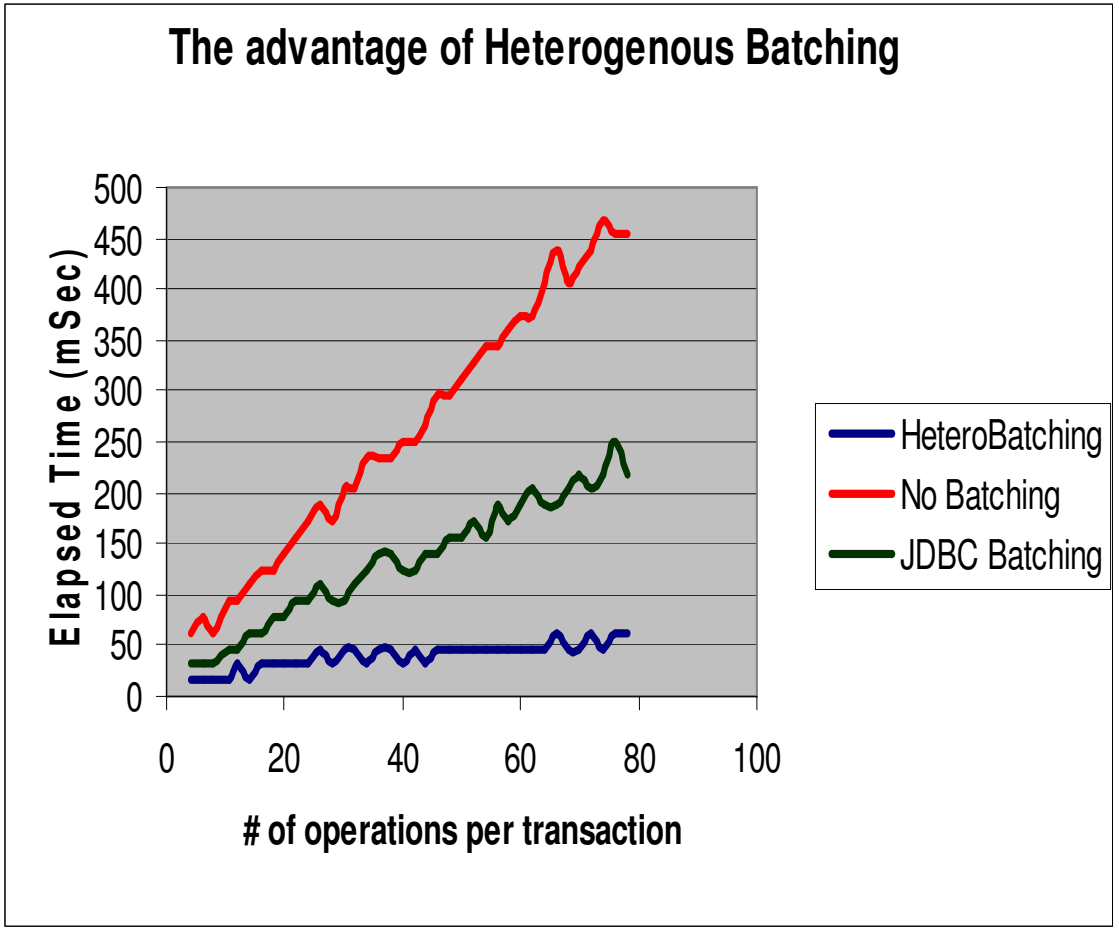

*** Preliminary findings based on validation with a test designed to demonstrate heterogeneous batching differences. This is not intended to be a formal benchmark.**

pureQuery – Stripping Literals from SQL

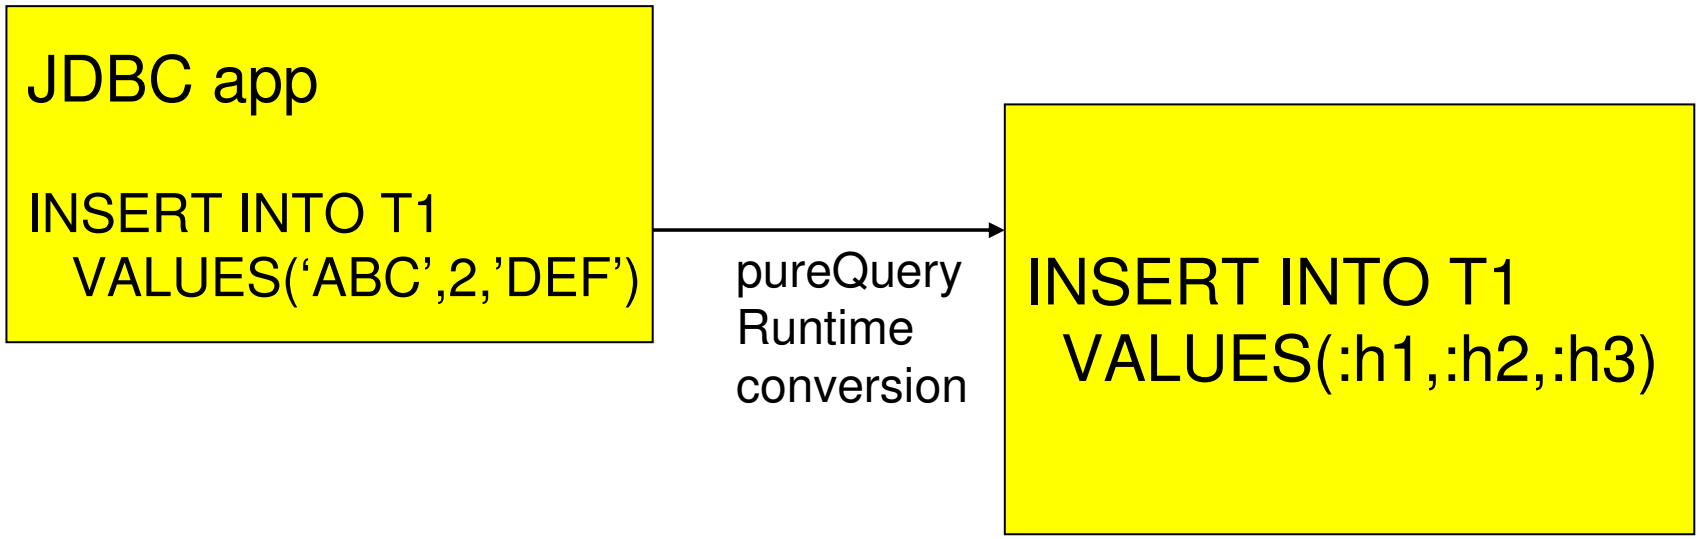

- pureQuery can identify statements that use no parameter markers, and strip the literals out at runtime
- significant performance gains:
  - less CPU cost at PREPARE
  - better use of dynamic statement cache

Optim pureQuery Runtime for z/OS

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

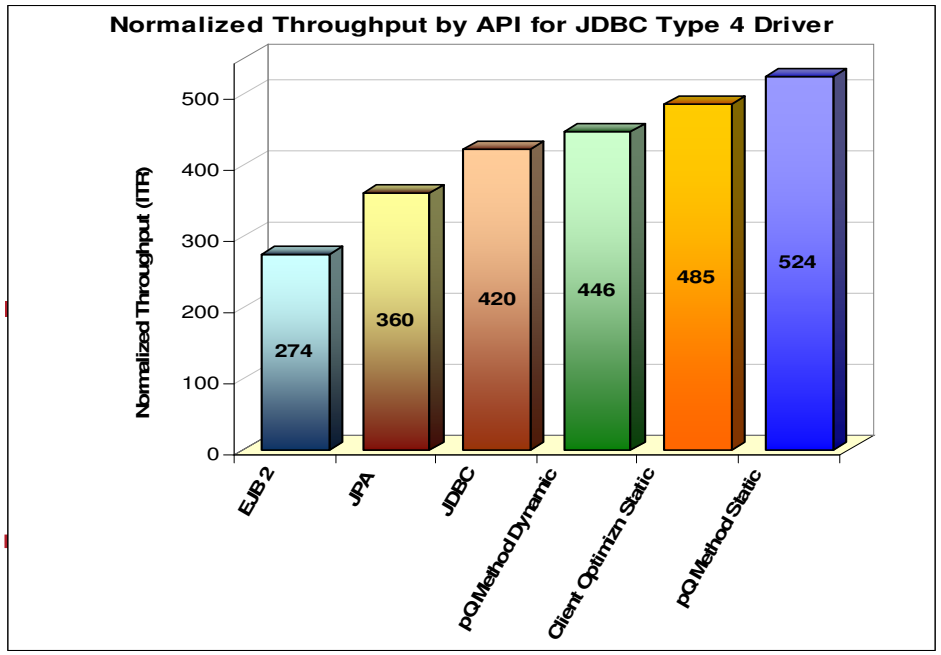

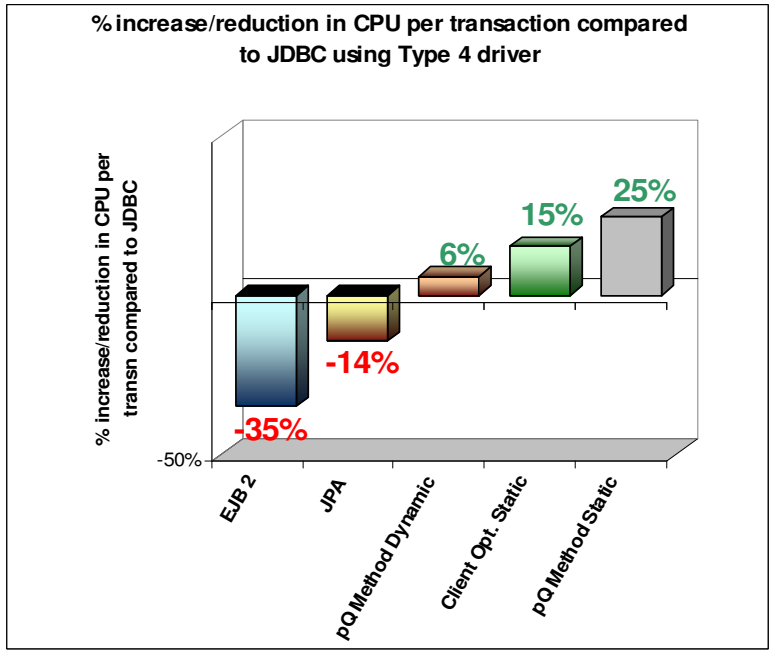

- IRWW – an OLTP workload, **Type 4 driver**
- Cache hit ratio between 70 and 85%
- **15% - 25% reduction on CPU per txn** over dynamic JDBC

Throughput Increase with .NET

- Same IRWW OLTP application used for the Java tests but in .NET
- Application access DB2 for z/OS via Windows Application Server (IIS)
- Throughput during static execution increased by 159% over dynamic SQL execution assuming a 79% statement cache hit ratio

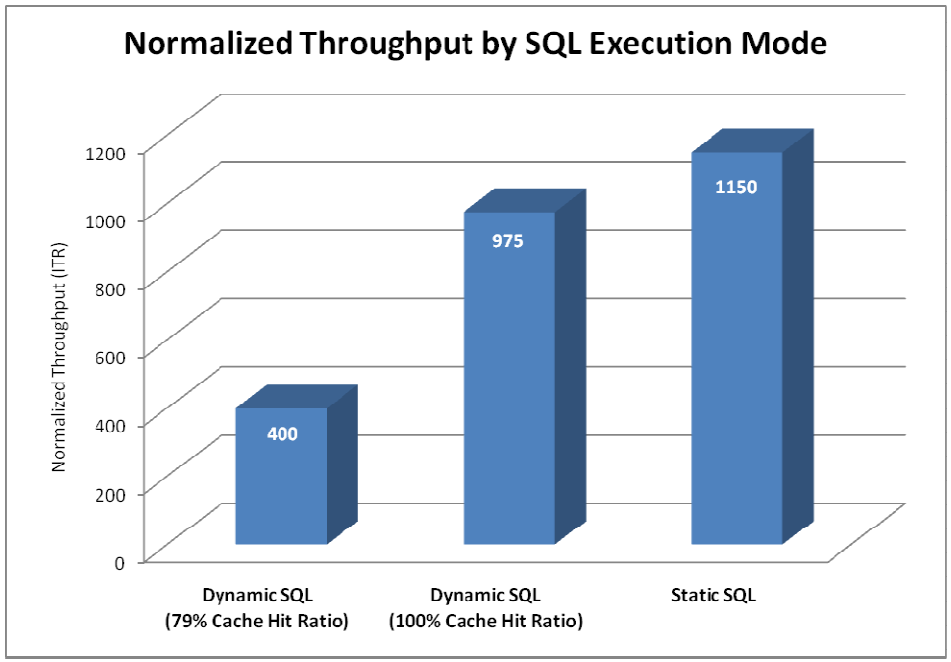

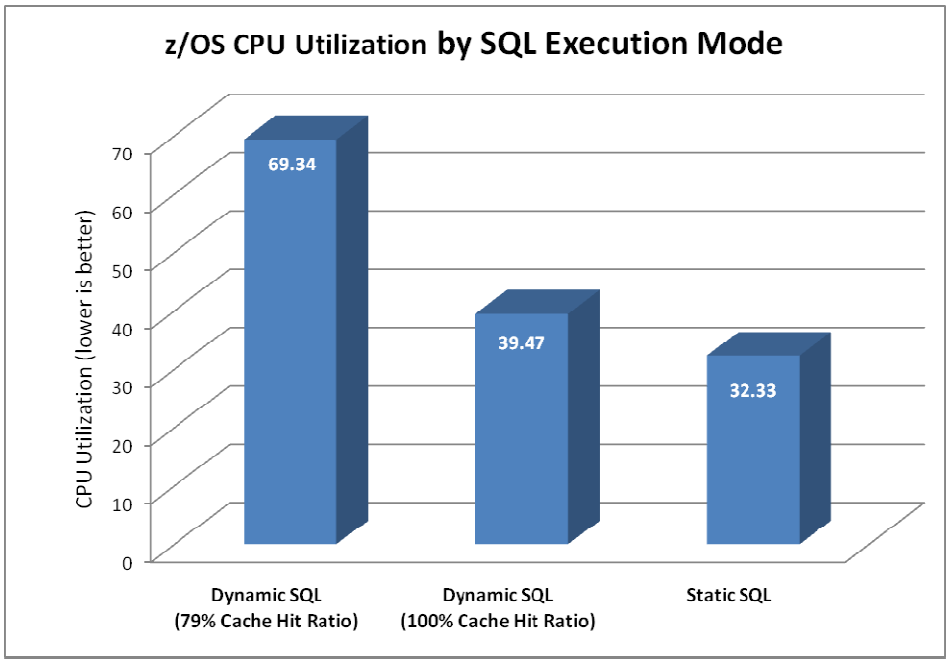

Optim Development Studio – SQL Outline

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 column used within the application?”
- Use with modern Java frameworks e.g. Hibernate, Spring, iBatis, OpenJPA

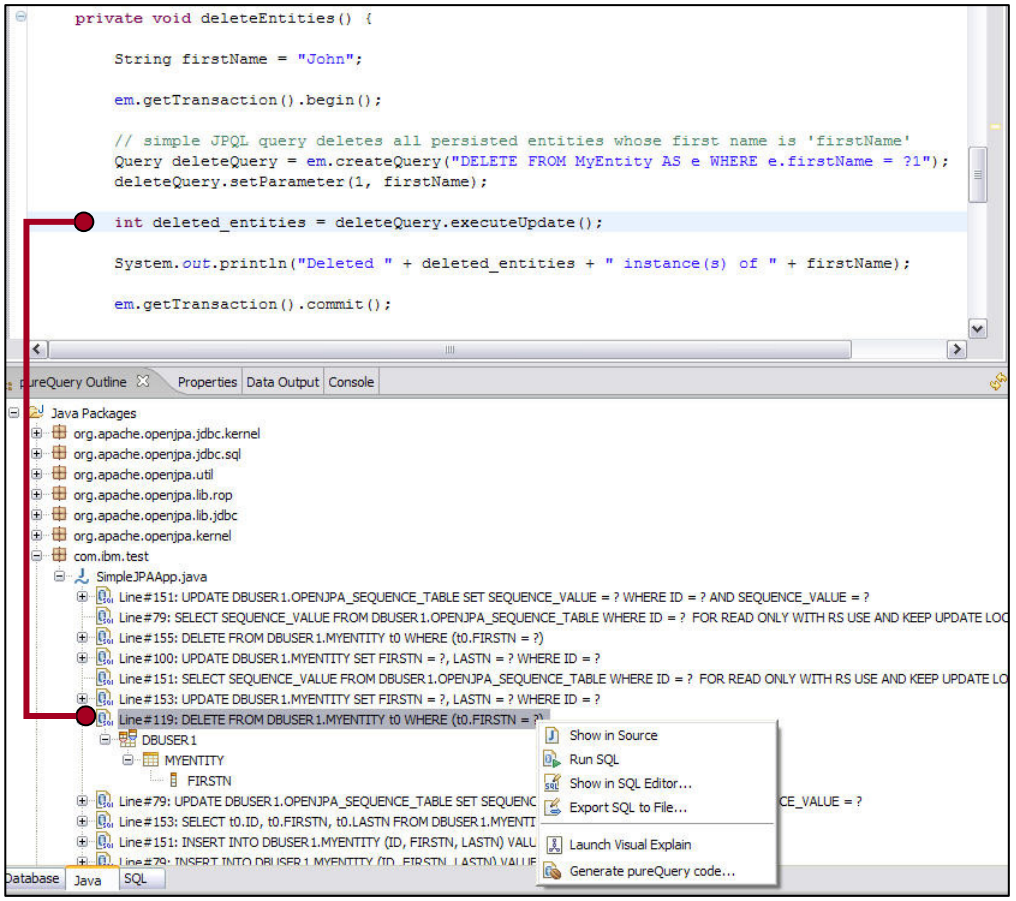

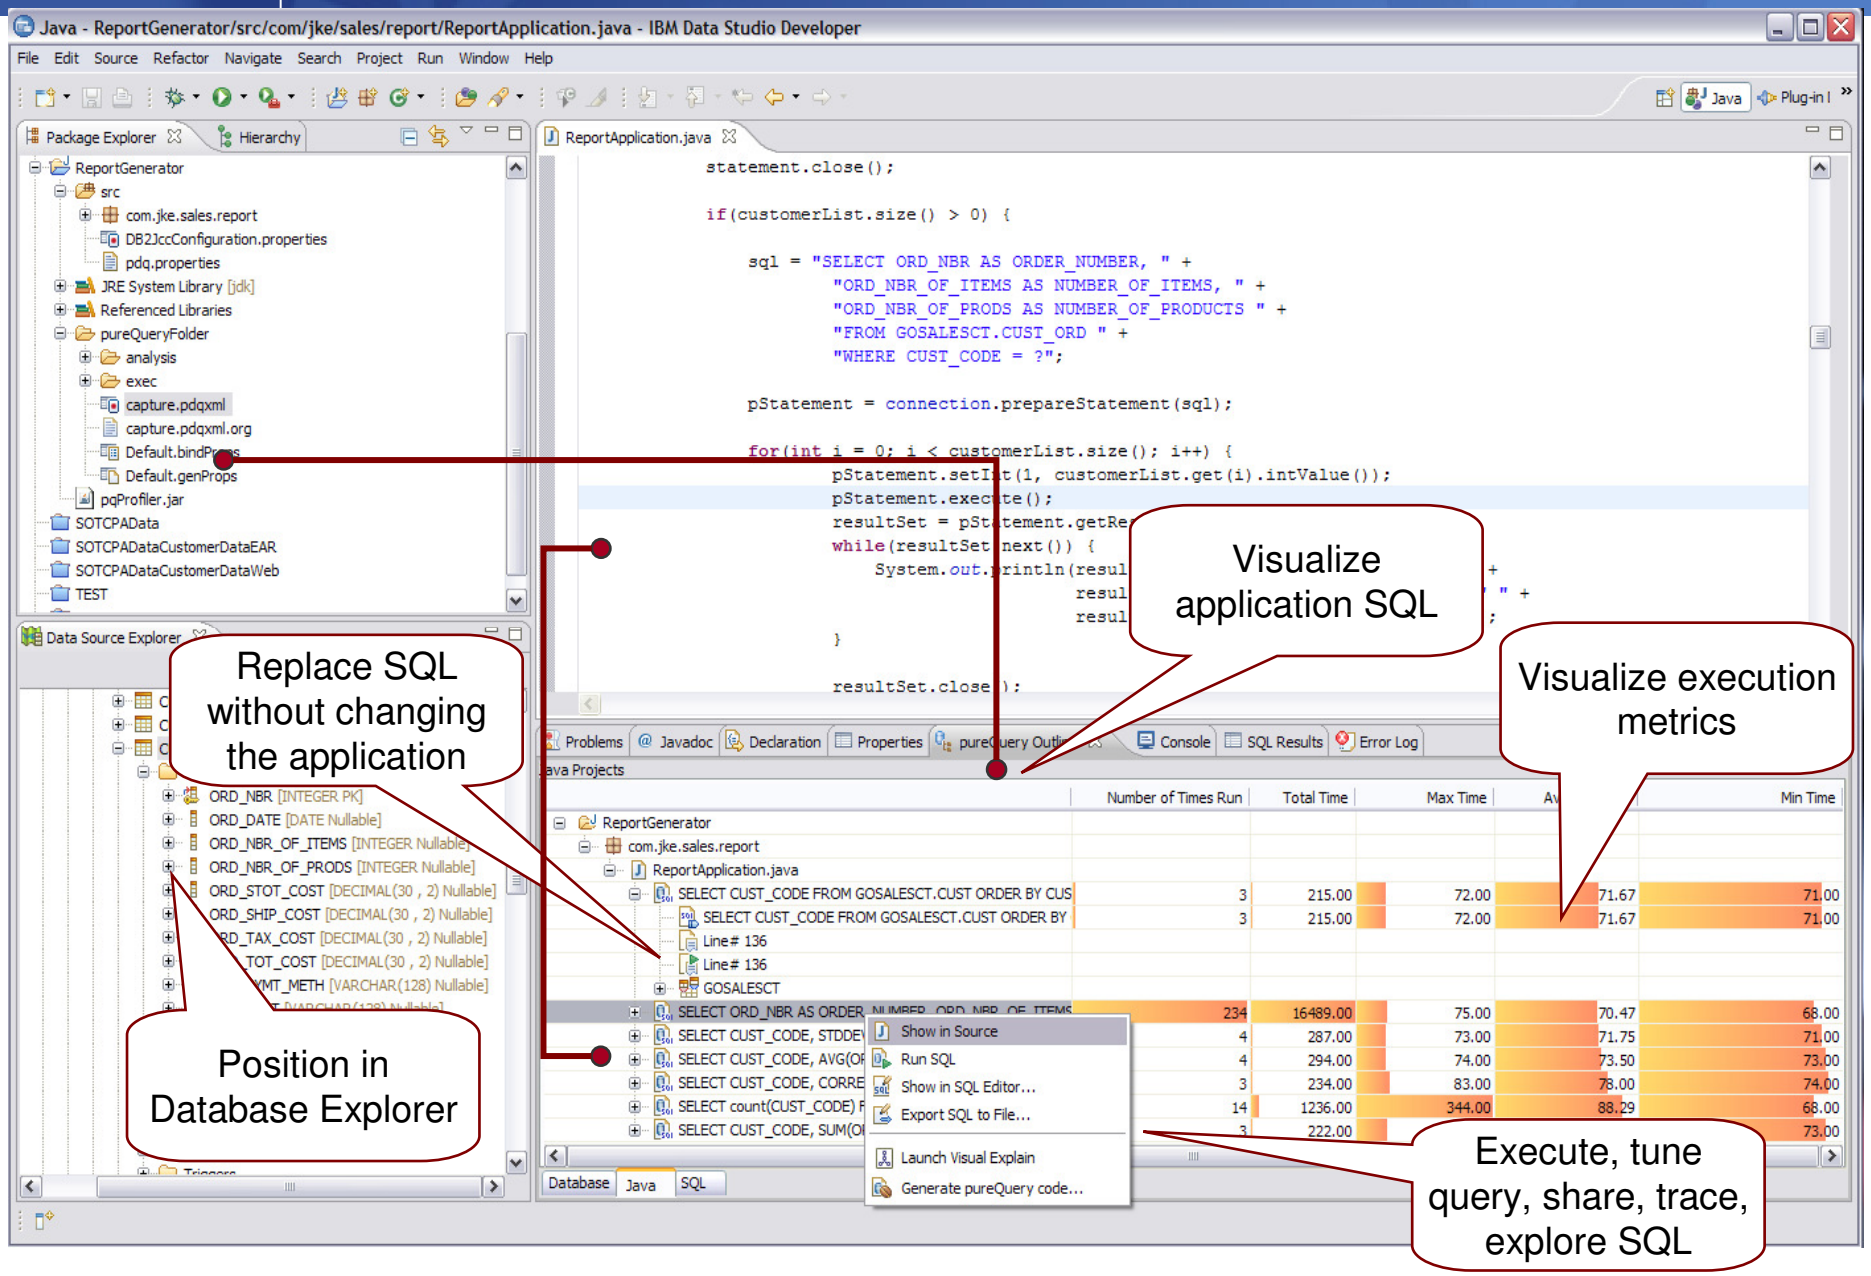

Optim Development Studio

- *Analyze Use of Sensitive Data in Applications*

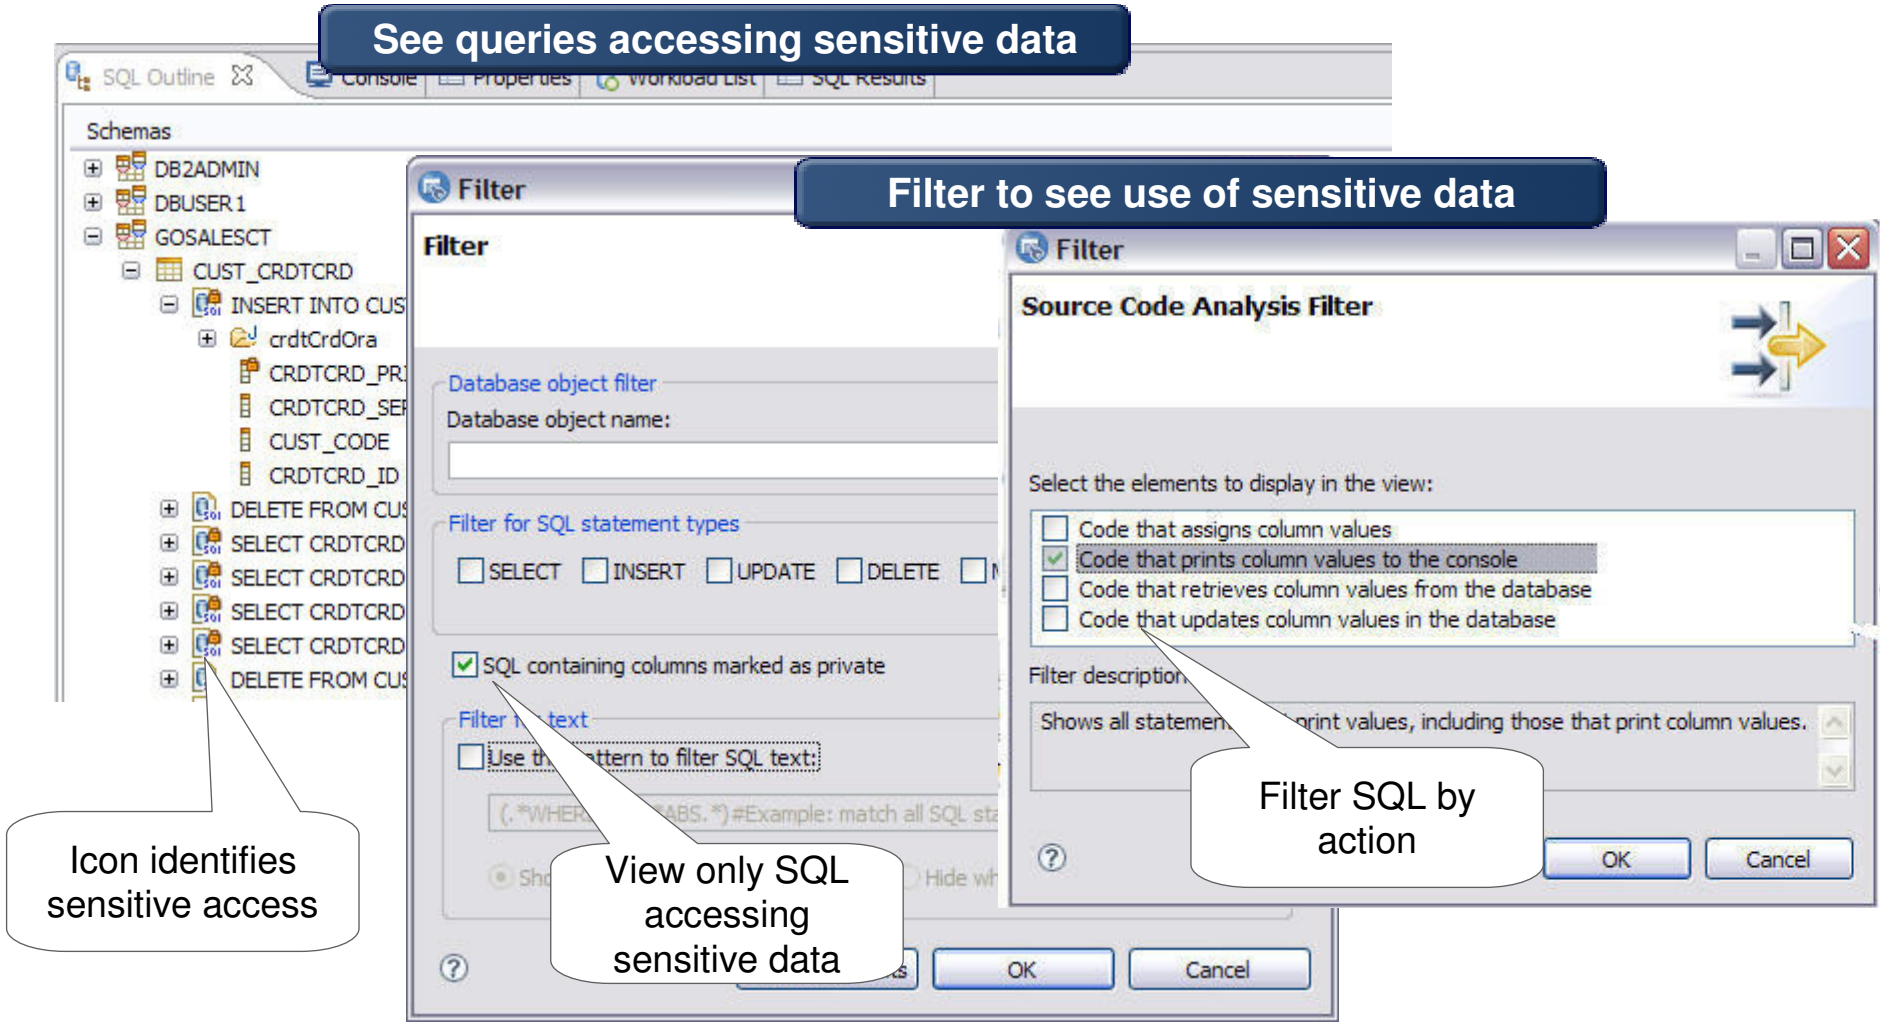

More Visibility, Productivity, and Control of Application SQL

- Capture SQL with pureQuery
  - Capture performance and application metadata, works with:

HIBERNATE iBATIS spring source **Java Persistence Architecture**

- Optimize
  - Review and share SQL
    - Visualize hotspots
    - Analyze impacts from schema changes
    - Trace SQL to originating source
  - Optimize and replace SQL
    - ***using Optim Query Tuner to get advice***
  - Create approved SQL list

- Deploy
  - Configure execution properties
  - Optionally bind for static execution

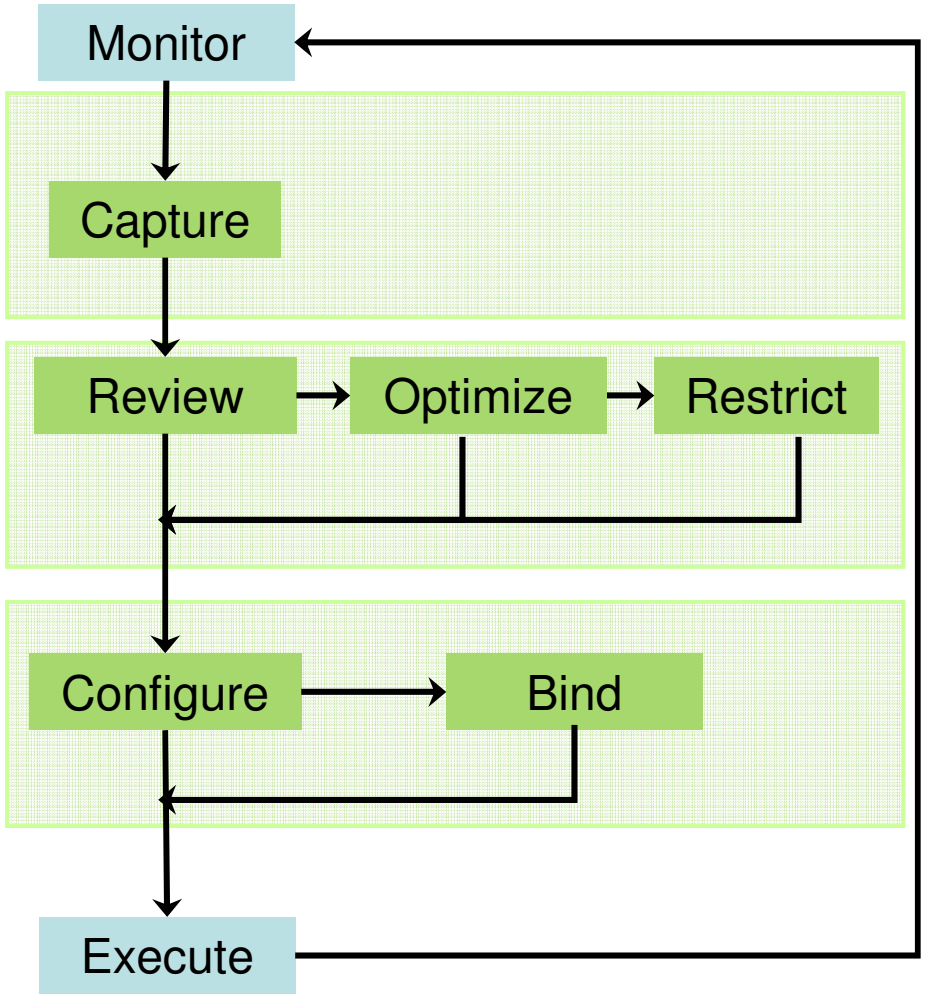

Optim Database Administrator
–Deploy without Disruption

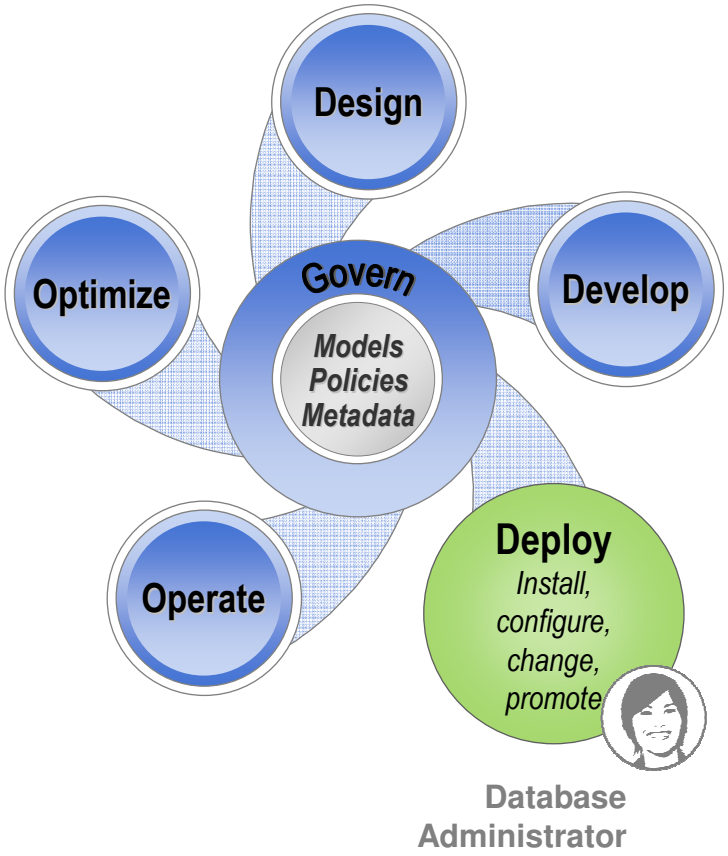

Optim Database Administrator…

Improves DBA productivity and reduces application outages by automating and simplifying complex DB2 structural changes including change-in-place as well as database migration scenarios.

- Models, automates and deploys complex schema changes
- Identifies dependencies and analyzes impact to mitigate deployment risk
- Preserves data, dependent objects, privileges, and application binding
- Synchronizes, copies, clones, or merges database schema definitions from the source to the target
- Documents changes for collaboration and audit
- Enables undo or restart -- if

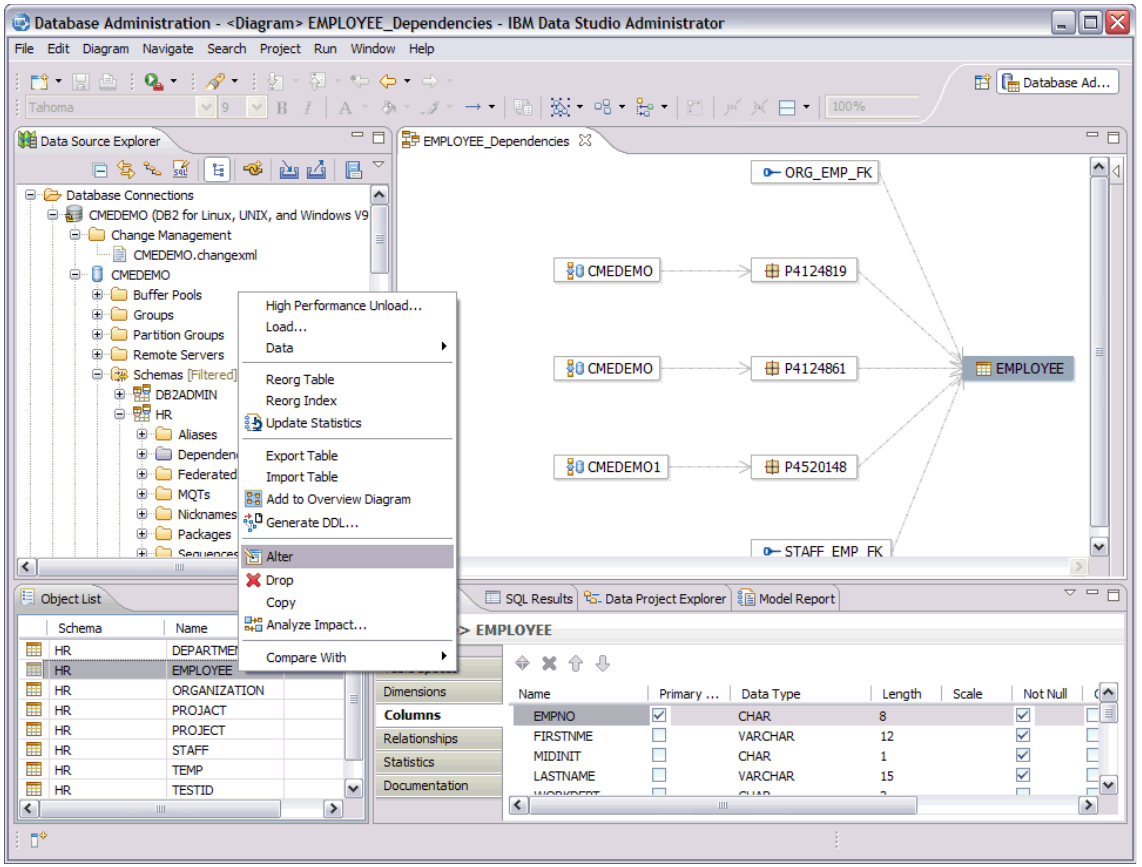

Get a Broader Perspective on Database performance

- *DB2 Performance Expert and Extended Insight Solutions*

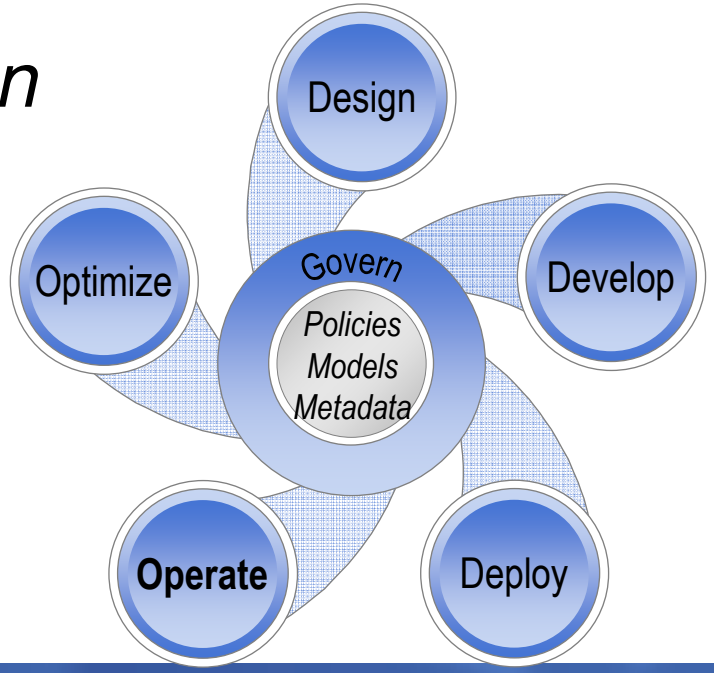

Get a Broader Perspective on DB Performance

- Toughest issue for Web applications
  - Problem diagnosis and resolution/Where is the time spent in application?

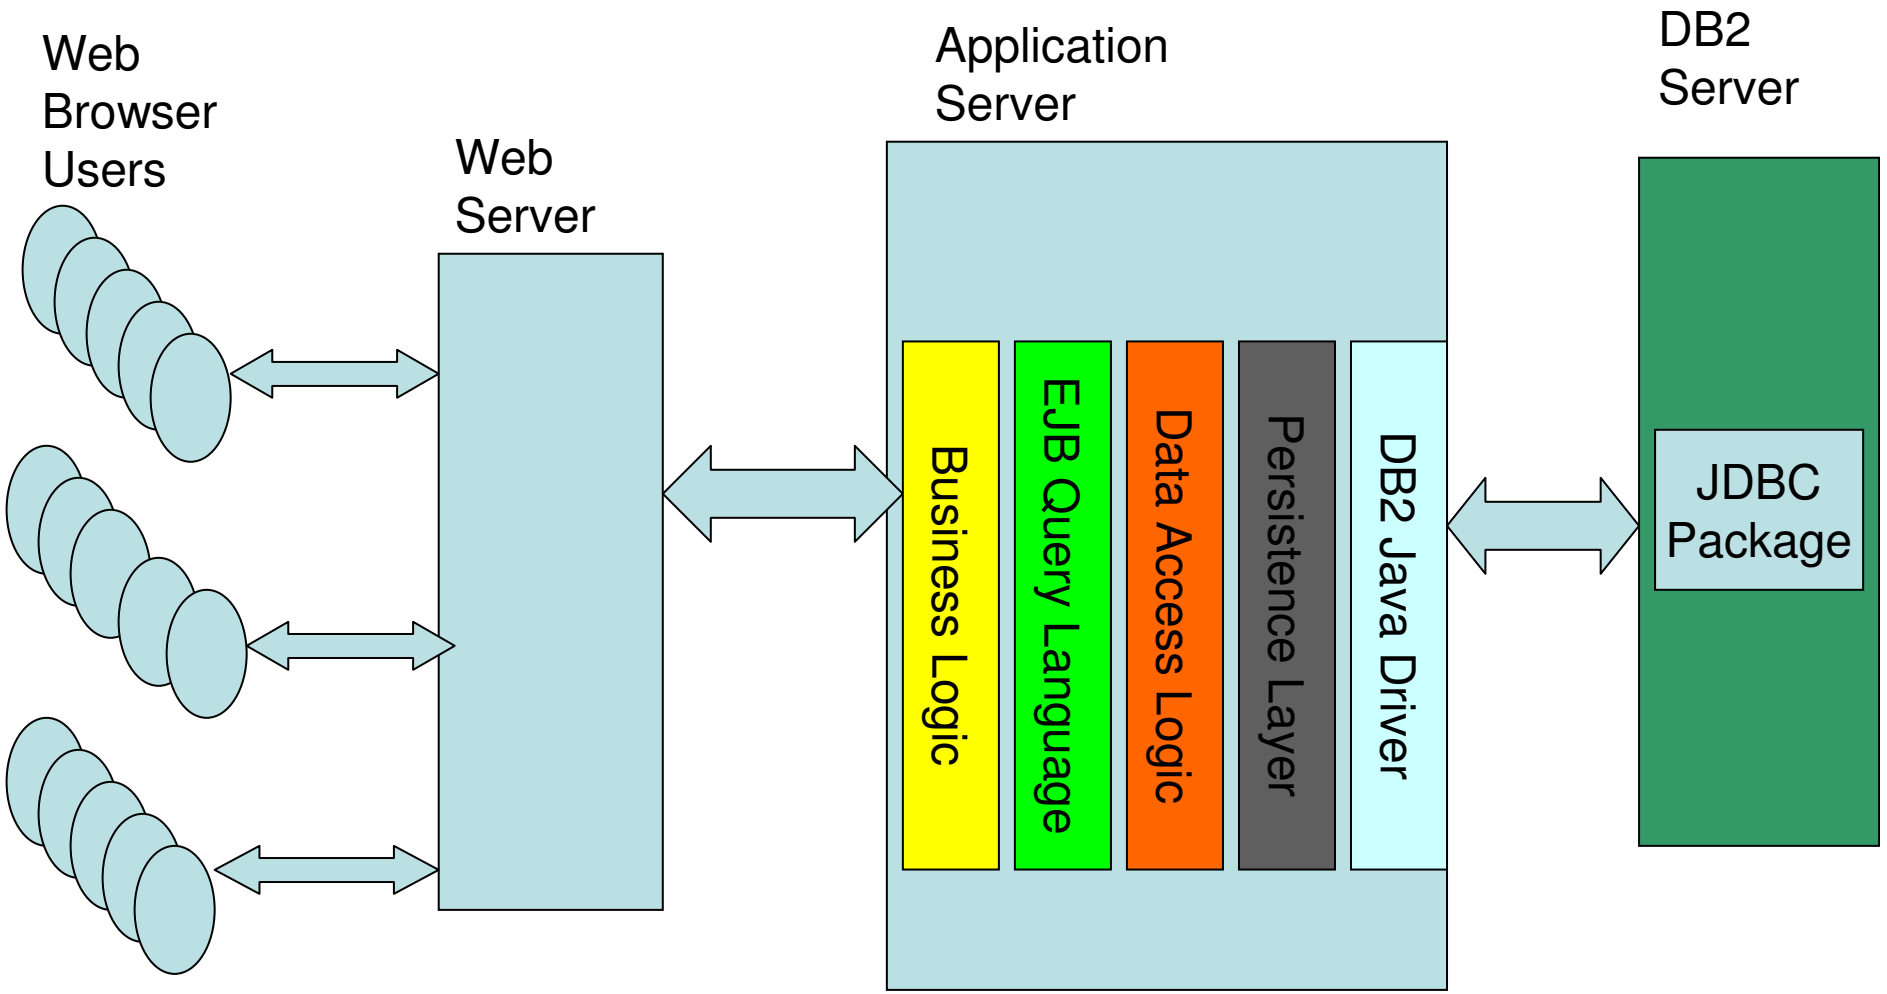

Where is my DB application spending its time?

- Customer Job Roles – A Barrier to a "Holistic View"

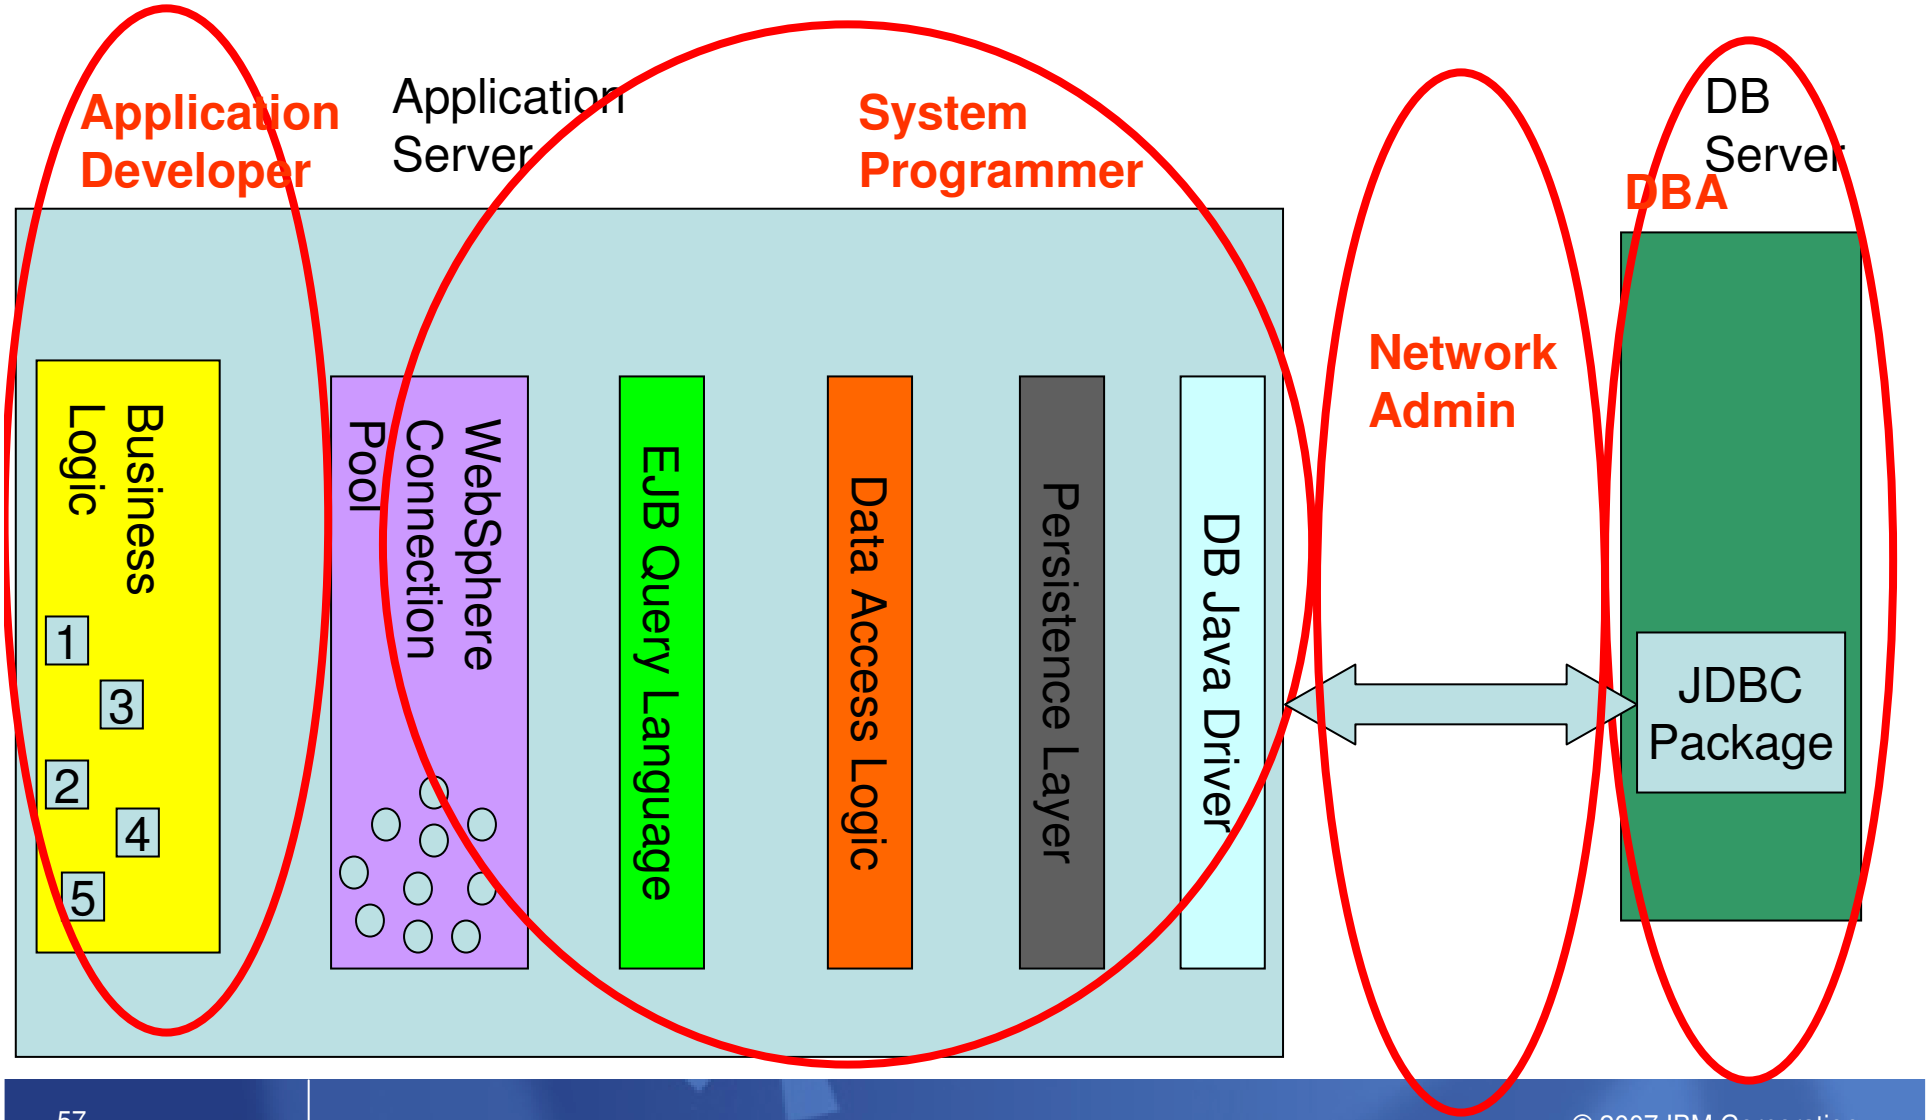

How do we plan to help?

- Show me what my application is seeing
- Let me figure out where in the software/hardware stack my problem is
- Is it really my problem, or someone elses ?
- Include database related information from WebSphere

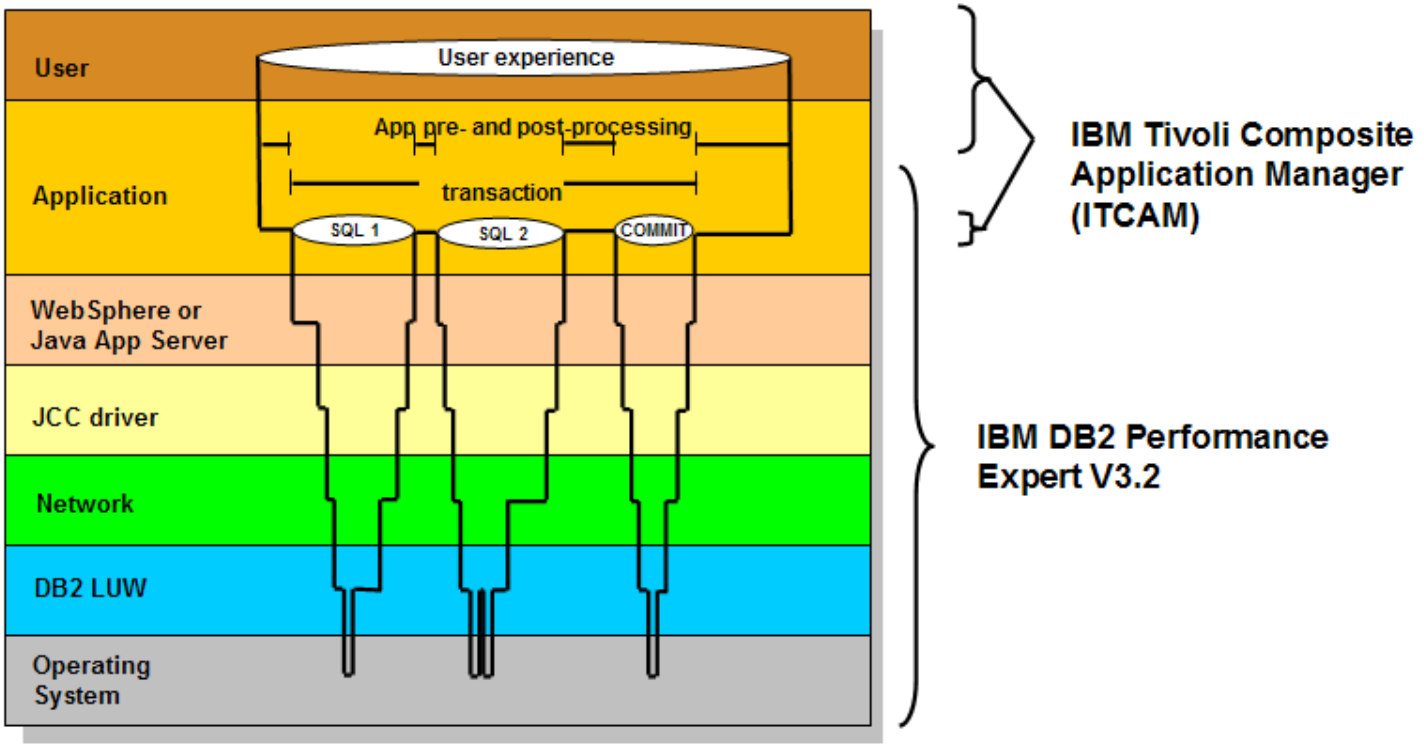

Manage to Service Level Agreements

DB2 Performance Expert and Extended Insight Feature

- A single point for monitoring, analysis and control of multiple DB2 instances to improve service level management
  - *With monitor information like time spent in application, in connection pooling, in the driver, or in the network*
- Monitor DB2 for Linux, UNIX, and Windows servers
  - Application monitoring
  - WLM monitoring
  - Engine monitoring
  - OS monitoring
  - Partition analysis
  - Realtime and historical data
  - Dashboards and alerts

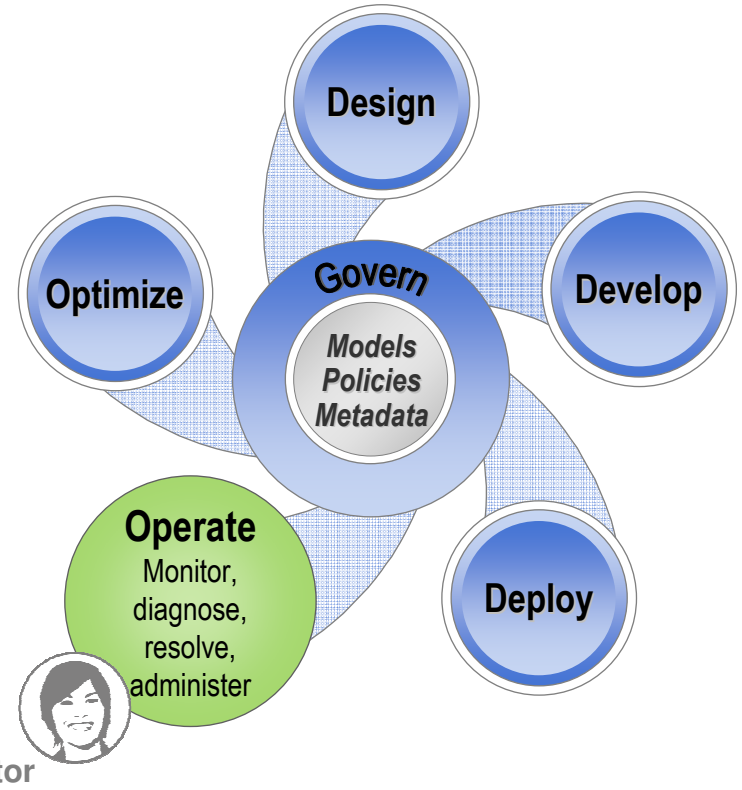

Diagnose overall application response time –
a example of the end-to-end monitoring

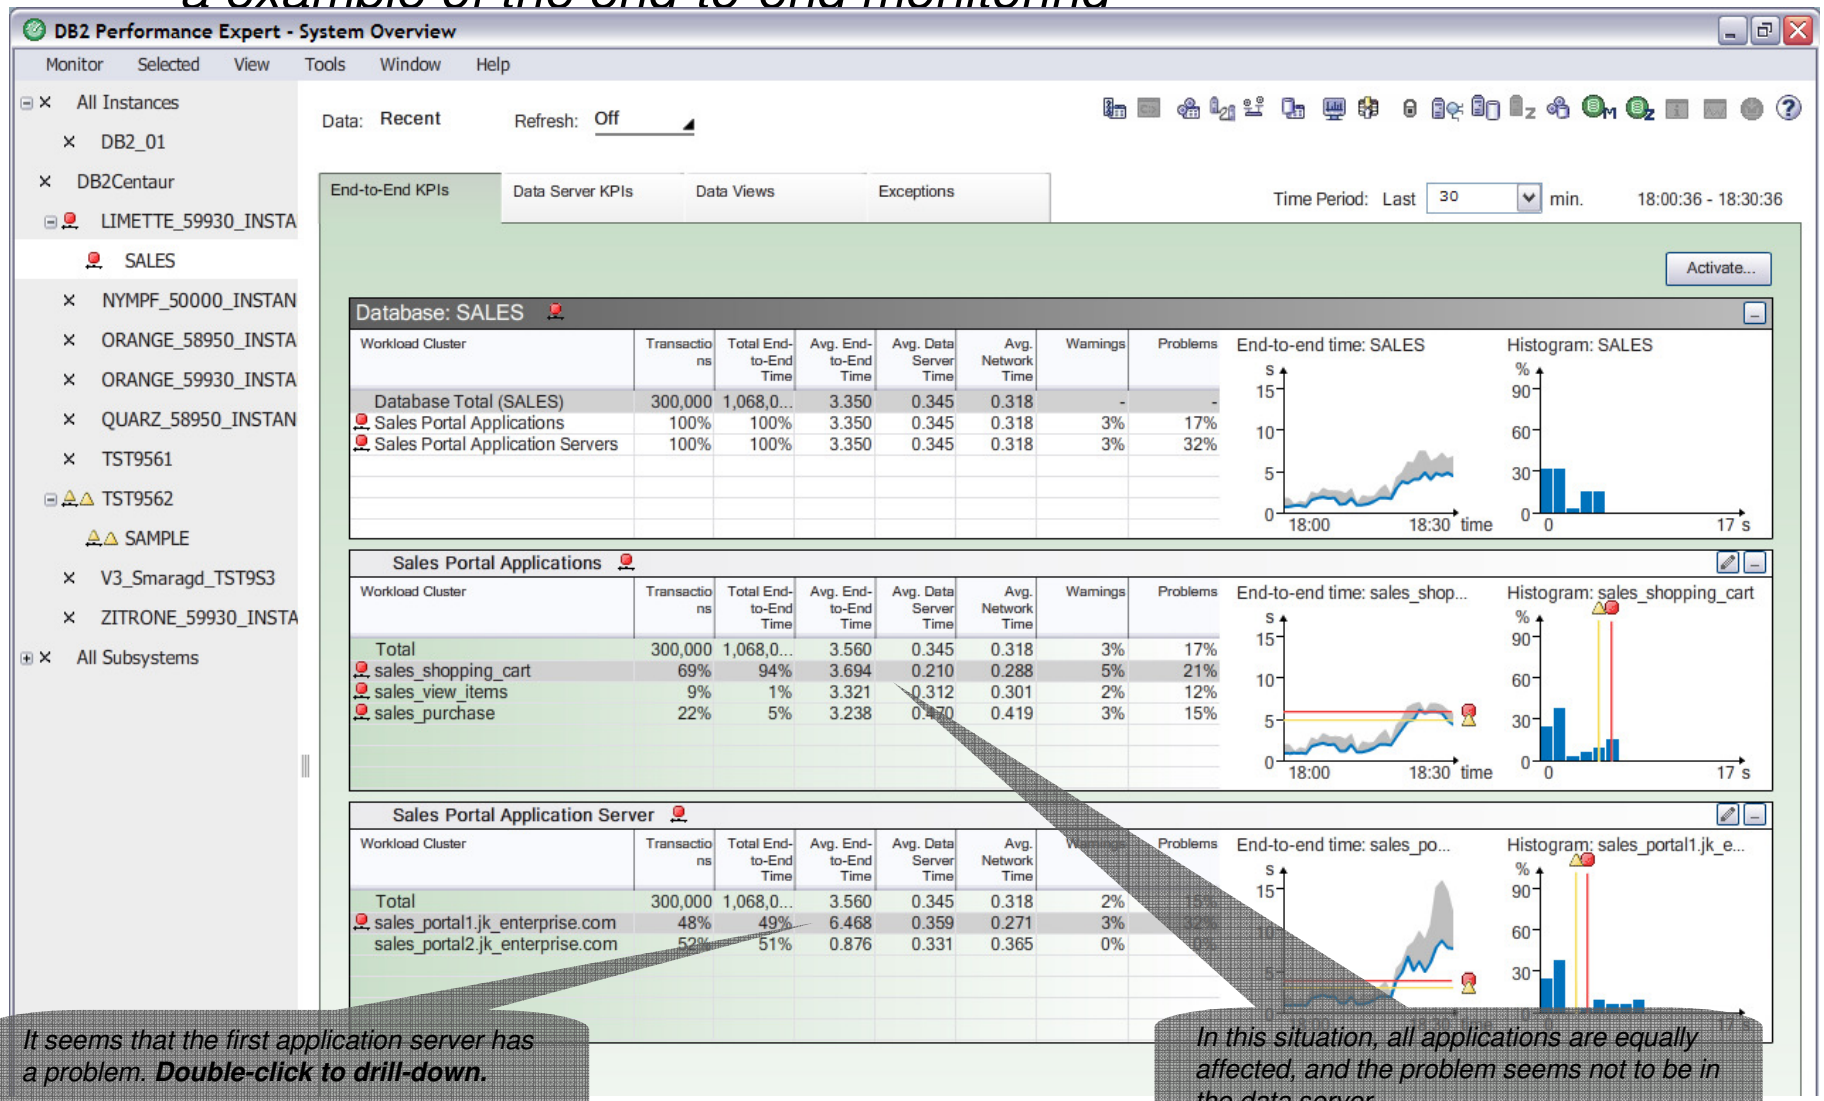

It seems that the first application server has a problem. ***Double-click to drill-down.***

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

Diagnose Application Server

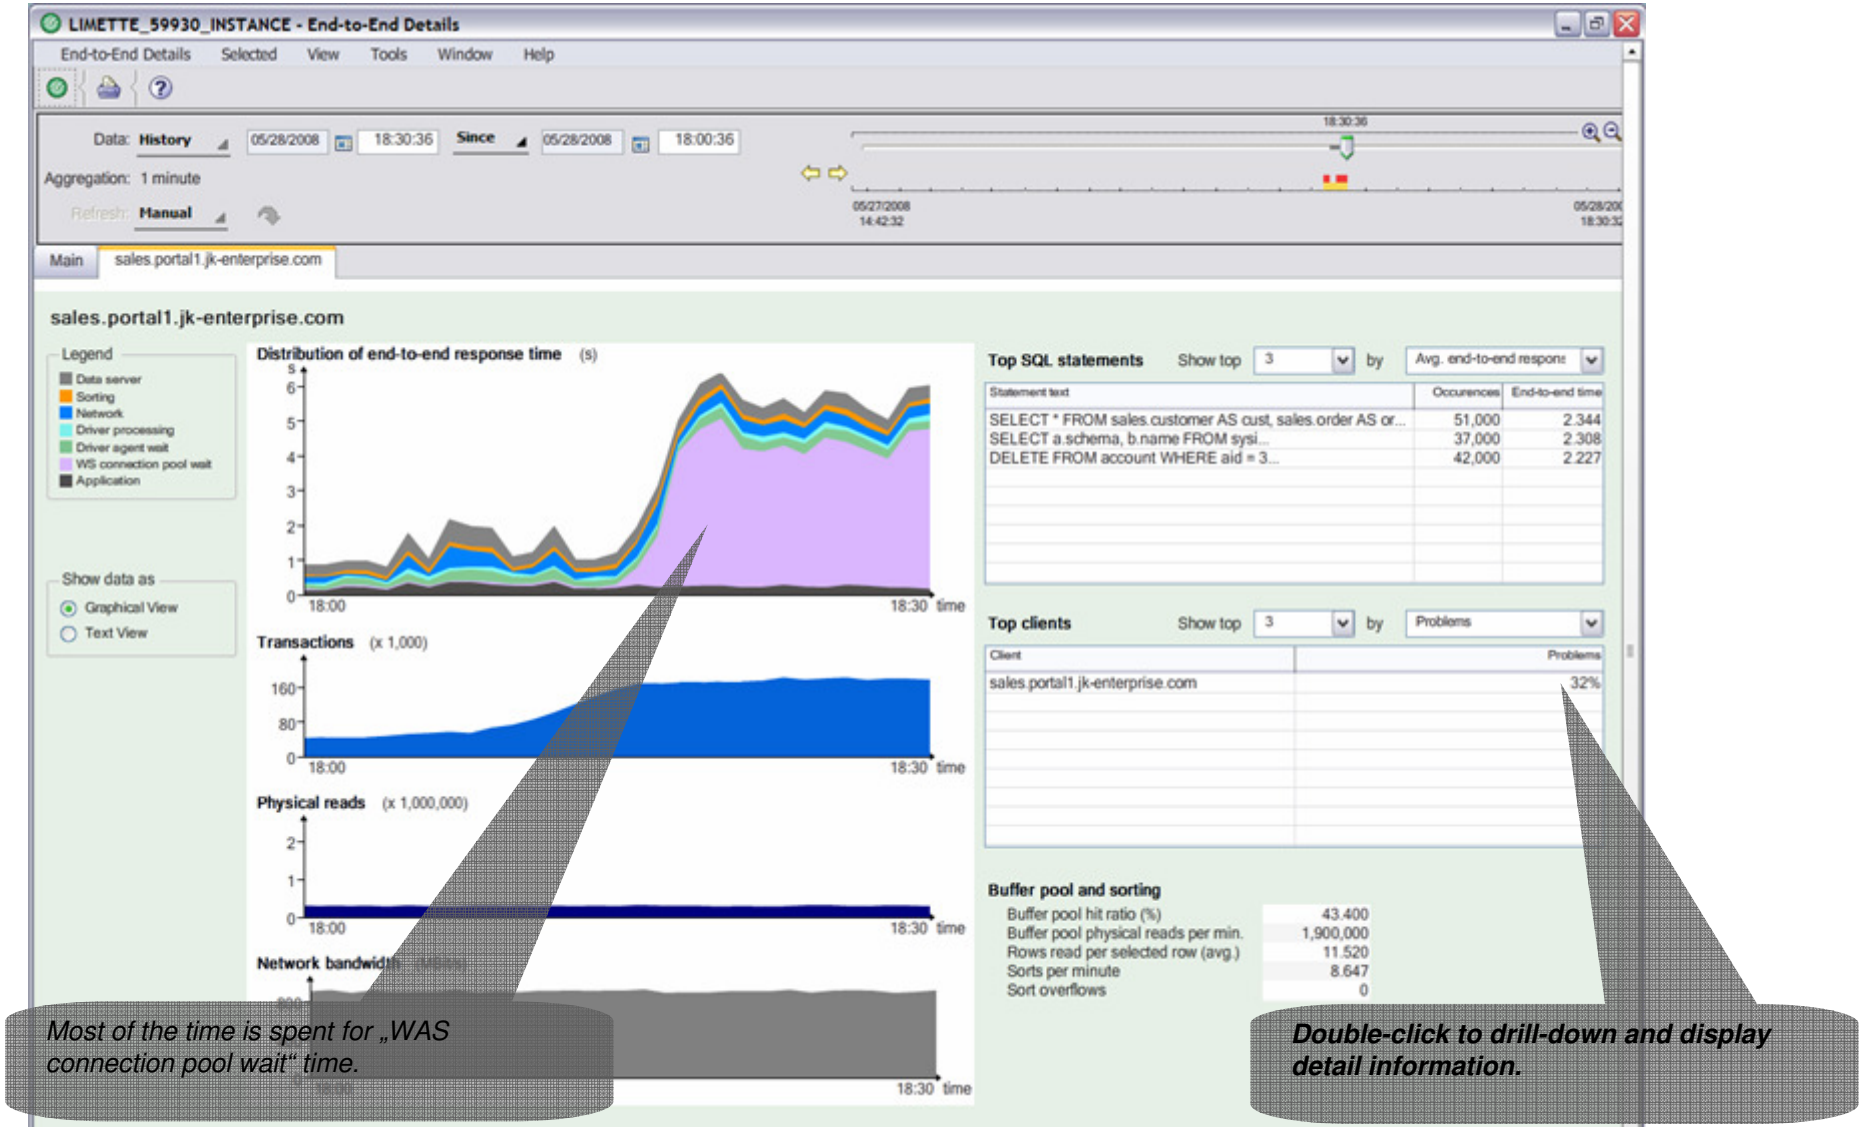

Diagnose on connection Pool ...

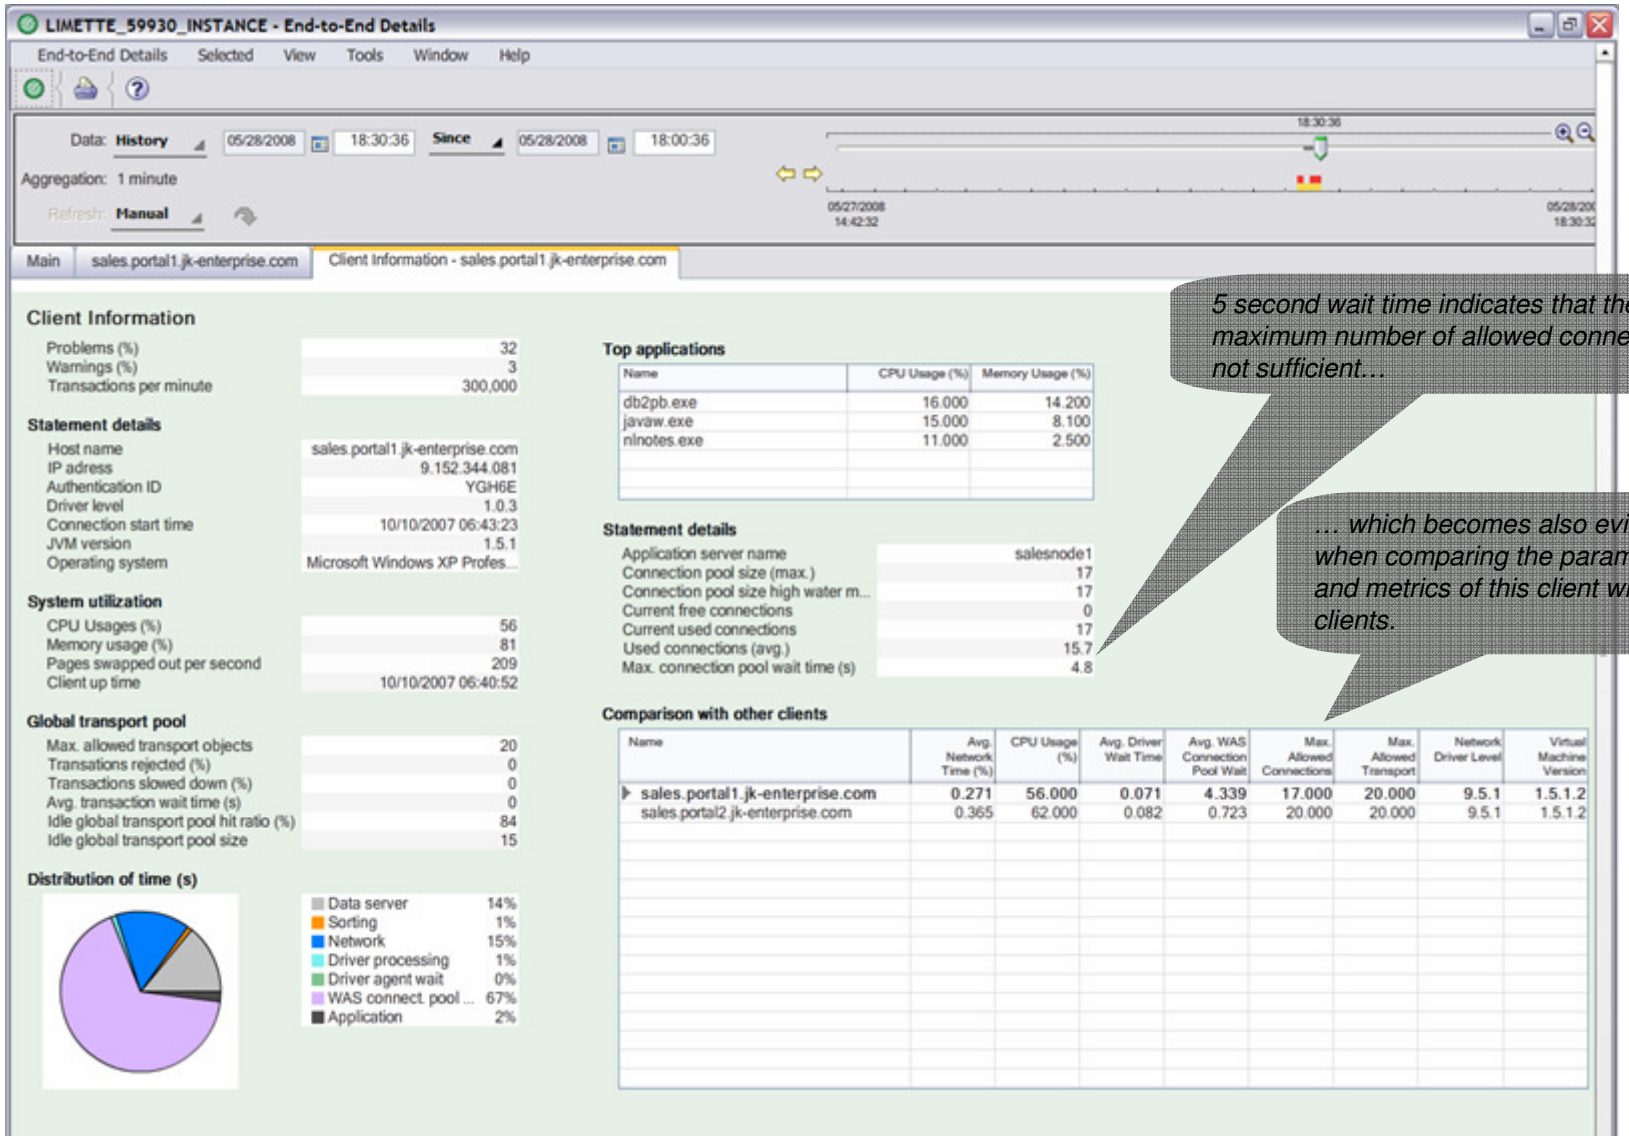

Optim Database Performance Manager future directions -- Associate SQL with Java Source

Heat Chart | Alerts | Dashboard | SLAs | In-flight analysis

▶ Database: Accounting

0:60 | 02:37 Nov, 13 | 09:37 | 3 Hours | 12:37

TOP 3 currently running SQL Statements DS Proc

Statement text	schema	E2E elapsed ▼	occurrences	sort time	phys. I/O
SELECT TIME FROM UNIVERSE	SAP3	132.13	1323	123.32	1.303
SELECT SALARY FROM PAYMEN					
DELETE FROM ACCOUNT WHERE					

TOP by
- DS elapsed
- DS CPU time
- cal I/O
- me

Statement information X

Stmt text SELECT TIME FROM UNIVERSE — Analyze

Application	
DS user ID	KARN
Client IP addr / hostname	TPKARN.de.ibm.com
Client user ID	KARN
Client workstation name	TPKARN
Client application name	Jawaw.exe
Client accounting	N/A
application name	Online banking
application contact	hkarn@de.bm.com
package	West.OLBank
class	Account
method	Transfer()
source line	314

Time distribution

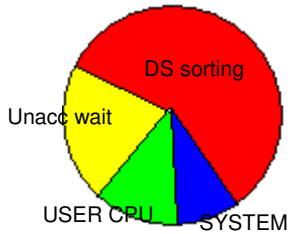

Resource usage	
Query cost estimates	18.456
Buffer Pools	
Data – hit ratio (%)	43.4%
Data – physical reads / min	4323
Index – hit ratio (%)	54.2%
Index – physical reads / min	3214

Statement elapsed time	
Current	132.13 sec
last day	239.40 sec
last week	15.60 sec

Stop SQL | Force application

Database Support by Product (as of today)

	DB2 for z/OS	DB2 for LUW	DB2 for i	IDS	Oracle	SQL Server	Sybase	MySQL
Data Studio	✔	✔	✔	✔				alphaworks derivative
Data Studio Administration Console	✔	✔		OpenAdmin Tool				
InfoSphere Data Architect	✔	✔	✔	✔	✔	✔	✔	✔
Optim Development Studio	✔	✔	✔	✔	✔			
Optim pureQuery Runtime	✔	✔	✔	✔	✔			
Optim Query Tuner	✔ + Workload	✔						
Optim Database Administrator	DB2 Admin Tool/Object Compare	✔						
Optim Test Data Management	✔	✔	✔	✔	✔	✔	✔	
Optim Data Privacy	✔	✔	✔	✔	✔	✔	✔	
Optim Data Growth	✔	✔	✔	✔	✔	✔	✔	
DB2 Performance Expert	Omegamon	✔						
DB2 PE Extended Insight		✔						
Database Encryption Expert	EE for DB2 and IMS	✔		✔	Vormetrics	Vormetrics	Vormetrics	

IBM Optim Solutions

- Optim Solutions Page:
  - http://www.ibm.com/software/data/optim/
- IBM Integrated Data Management (Optim and Data Studio):
  - http://www.ibm.com/developerworks/spaces/optim
    - Tutorials
    - Downloads
    - Forums / Blogs
    - Join the community!
- New demo: **Optim solutions for accelerating Java database access**
  - https://www.ibm.com/developerworks/offers/lp/demos/summary/im-optimsolutionsforjava.html

Disclaimer

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.