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

Bryan F. Smith bfsmith@us.ibm.com
IBM

Session: Xxx
Friday, March 4, 2011: 11:00 AM-12:00 PM
ACC, Room 201A

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, .NET, and CLI 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

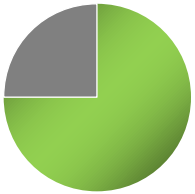

75% of CIO's believe they can strengthen their competitive advantage by better using and managing enterprise data.

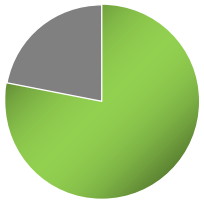

78% of CIO's want to improve the way they use and manage their data.

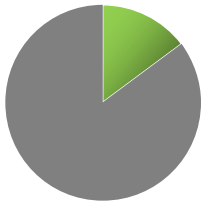

…but only 15% believe that their data is currently comprehensively well managed.

- ***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

Source: Accenture CIO Data Management Survey 2007. n=167 CIOs

Introducing Integrated Data Management

SHARE
Technology • Connections • Results

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

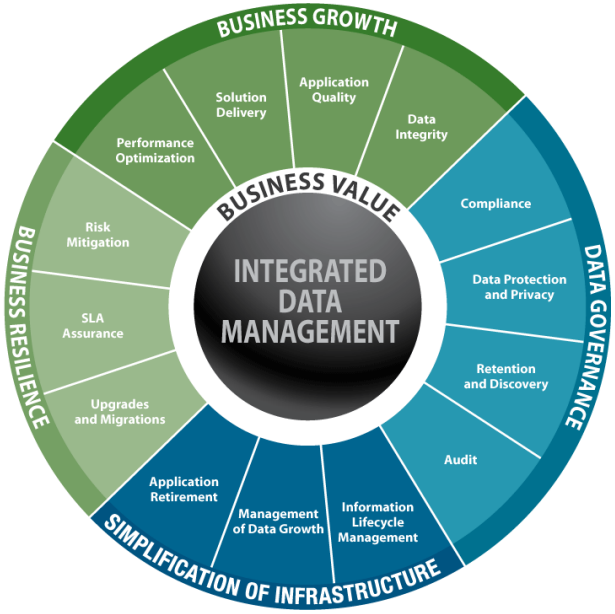

- ***Grow the business, without growing costs***
  - Develop and deploy business critical applications faster
  - Mitigate compliance risks with model-driven data governance
  - Prevent runaway infrastructure spending
  - Improve performance – of work teams, databases, applications, and business units

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

Integrated Data Management

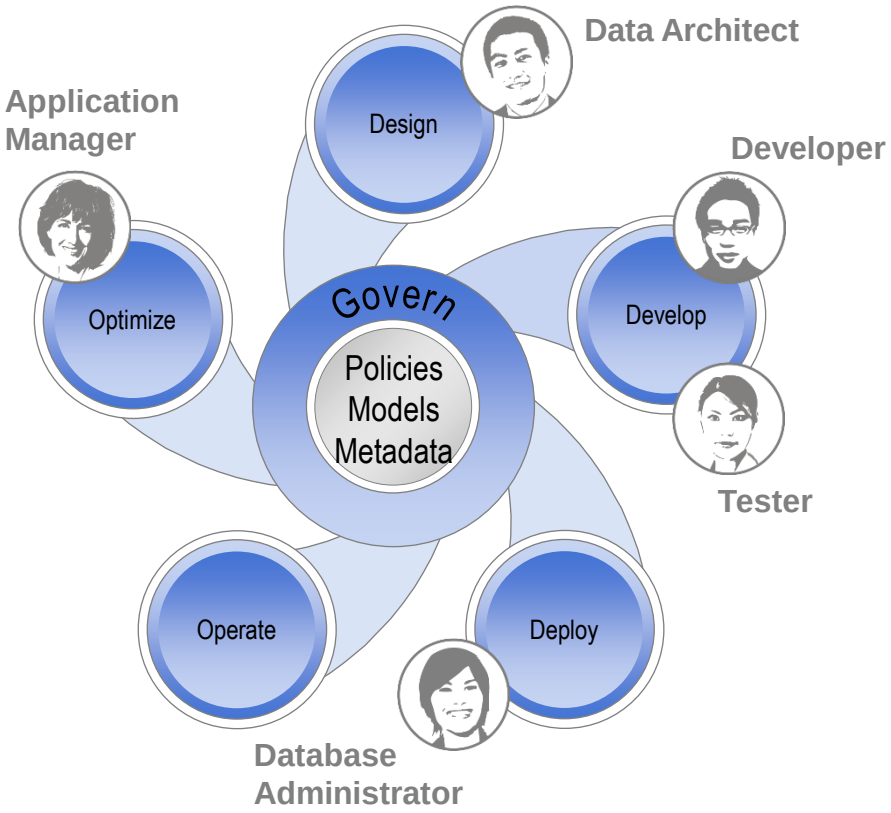

- **Deliver increasing value across the lifecycle**, *from requirements to retirement*
- **Facilitate collaboration and efficiency across roles,** *via shared artifacts automation and consistent interfaces*
- **Increase ability to meet service level agreements,** *improving problem isolation, performance 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

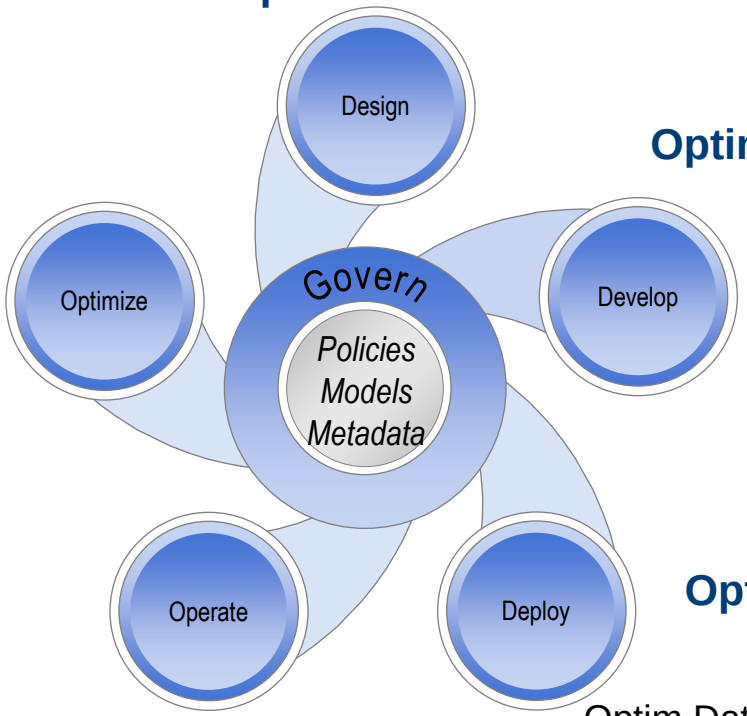

Optim Data Growth Solutions

Optim Query Workload Tuner

DB2 Performance Management

DB2 Performance Expert and Extended Insight

DB2 Automation Tool

DB2 Recovery Solution

DB2 Utilities Suite

DB2 Change Management Solution

Optim Development Studio

Optim Query Tuner

Optim Test Data Management

Optim Data Privacy Solutions

Optim pureQuery Runtime

Optim Database Administrator

DB2 Audit Management Expert
Database Encryption

Integrated Data Management User Interface Directions

SHARE
Technology • Connections • Results

Data Studio and Optim

z/OS Tools

Eclipse UI

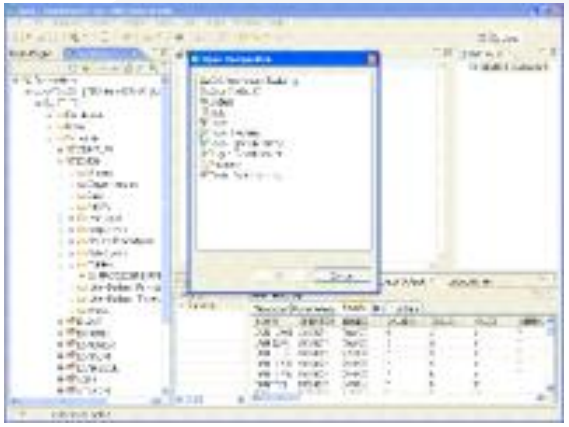

Web UI

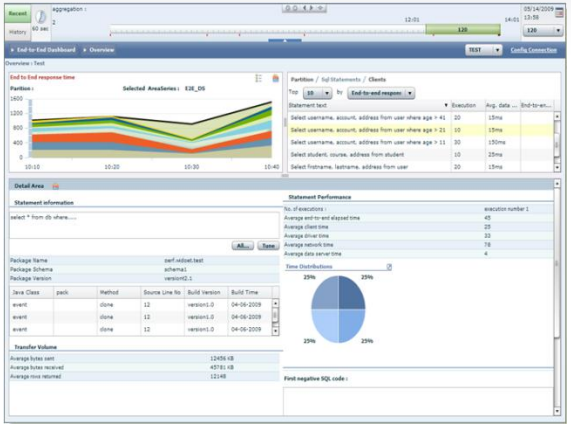

3270

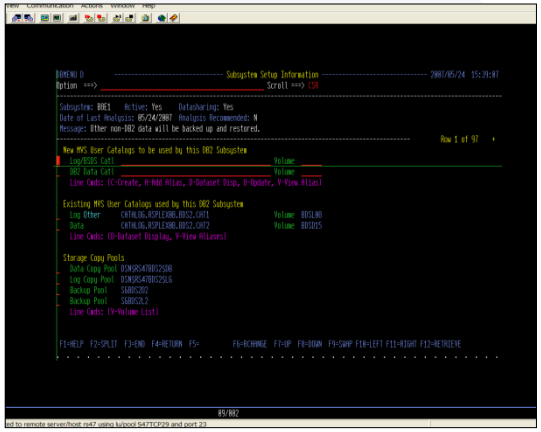

Develop and Deploy

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

Operate

- Monitoring
- Automation Management
- Reporting
- Data Governance

z/OS Database Administration

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

DB2 for z/OS Tools Portfolio

SHARE
Technology · Connections · Results

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

IBM Data Studio V2.2

No-charge

- 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
- Query Tuning (stats advisor, query formatter, etc.)
- Security Access Control
- Connection Management integration with Kerberos and LDAP
- Data Web Services
- IDS Server Support

Data Studio is no longer used as a brand

Data Studio: Data Modeling / App & DB Dev

A Consistent & Productive work environment

SHARE
Technology • Connections • Results

Create, Alter, Drop, Browse and Filter database objects

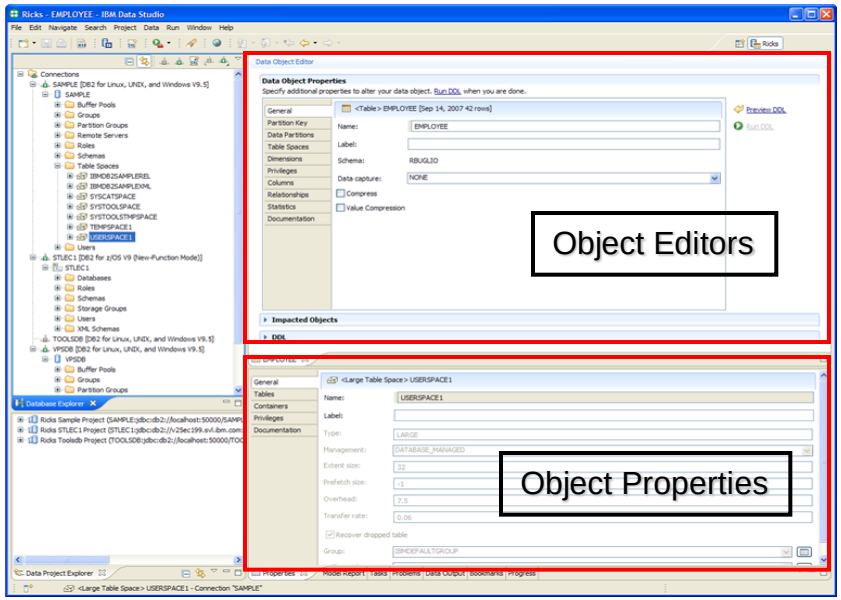

- Express yourself with optimal queries
  - Content assistance for database objects
  - Rapid interactive end-user feedback
  - Extensible templates
  - Multiple SQL statement testing
  - SQL assistance and XQuery assistance

ER Diagramming

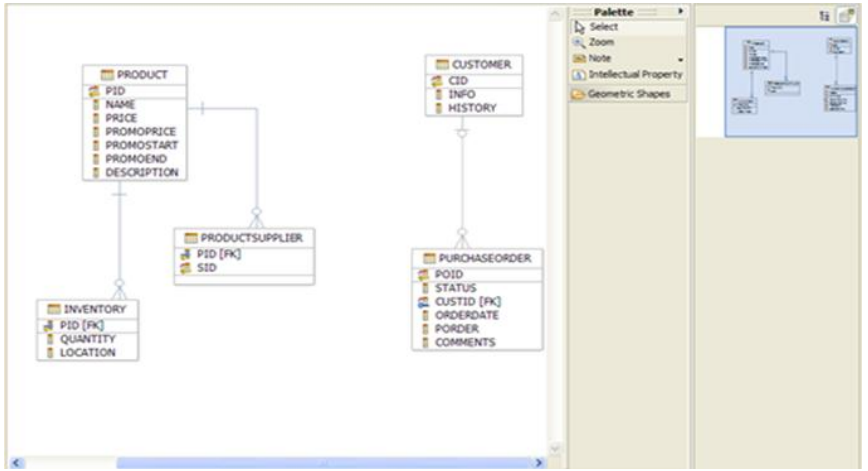

Integrated Query Editor

```
*s1.sql   *x2.sql
-- <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 IDE

- SQL content assist

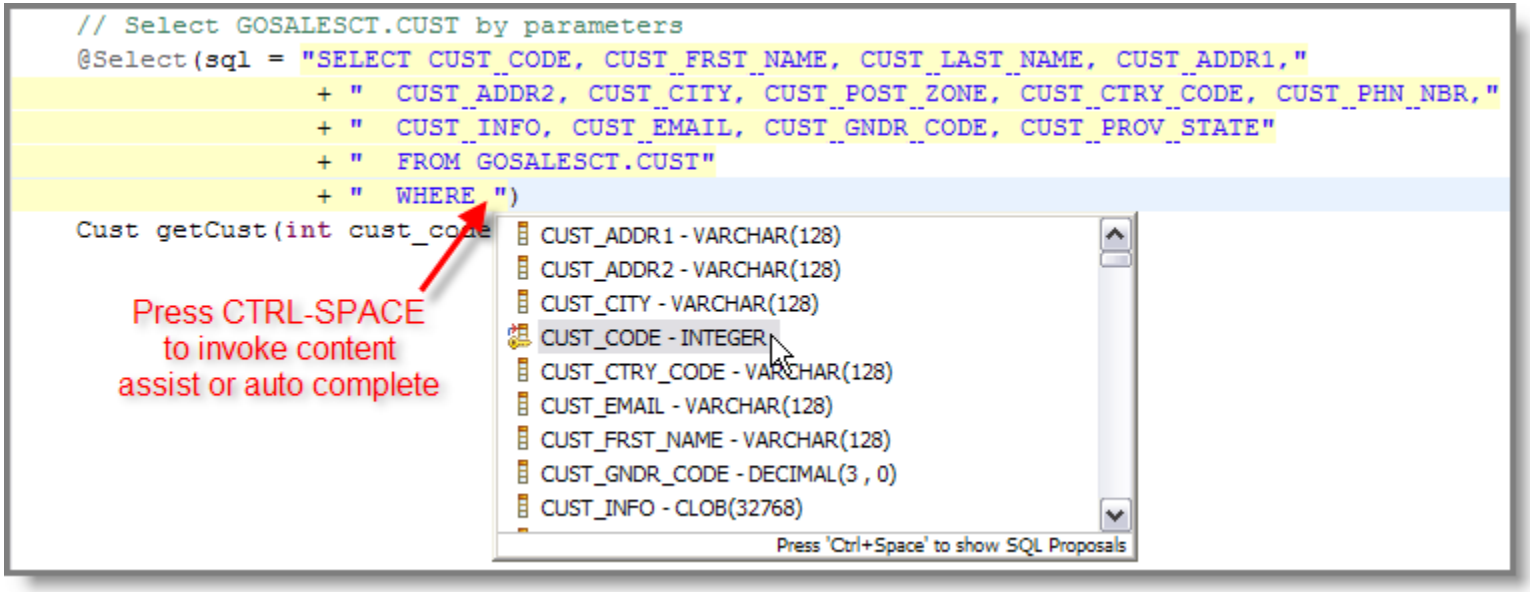

- SQL validation

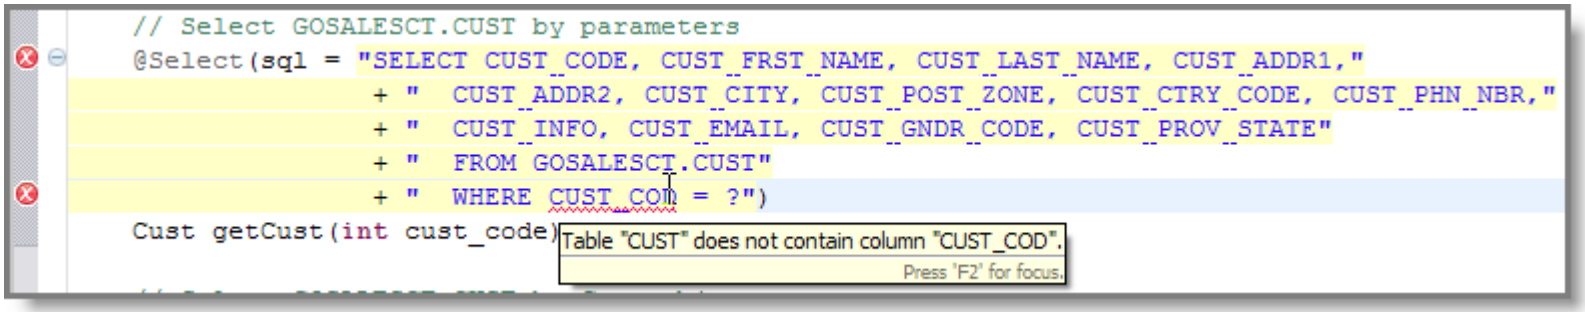

Unleash SQL in your IDE

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

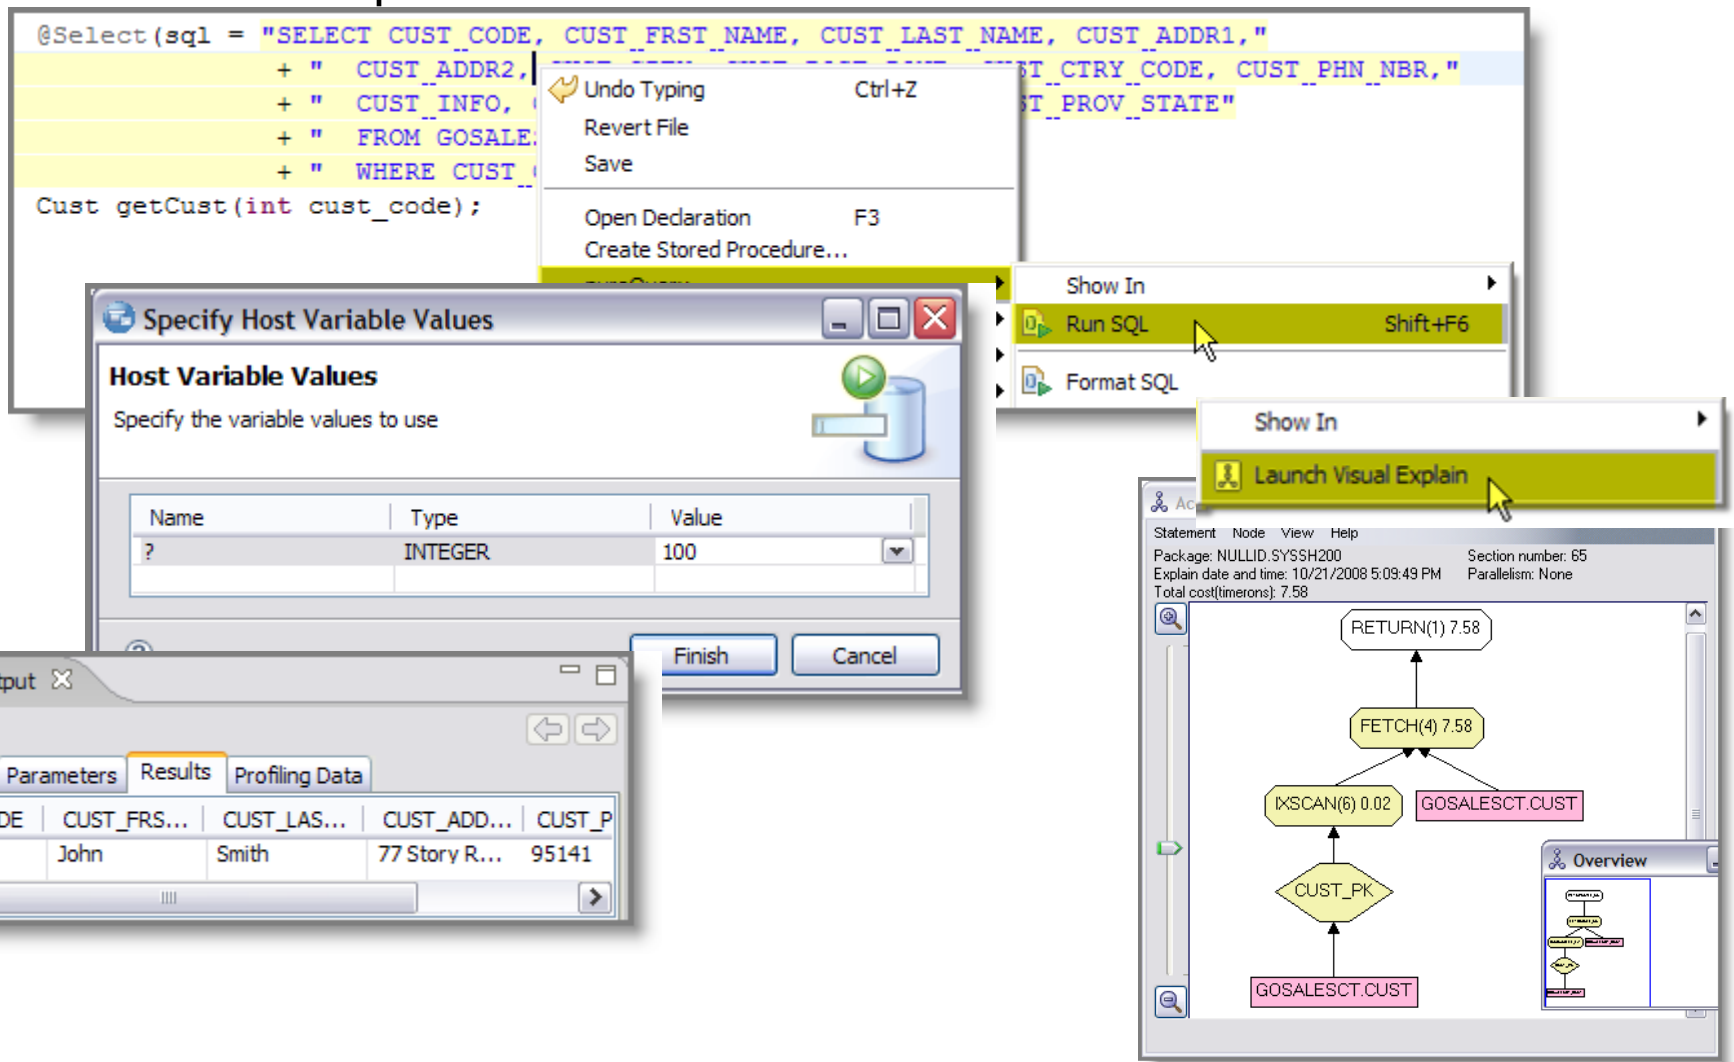

Stored Procedure Support
SQL and Java Debugger

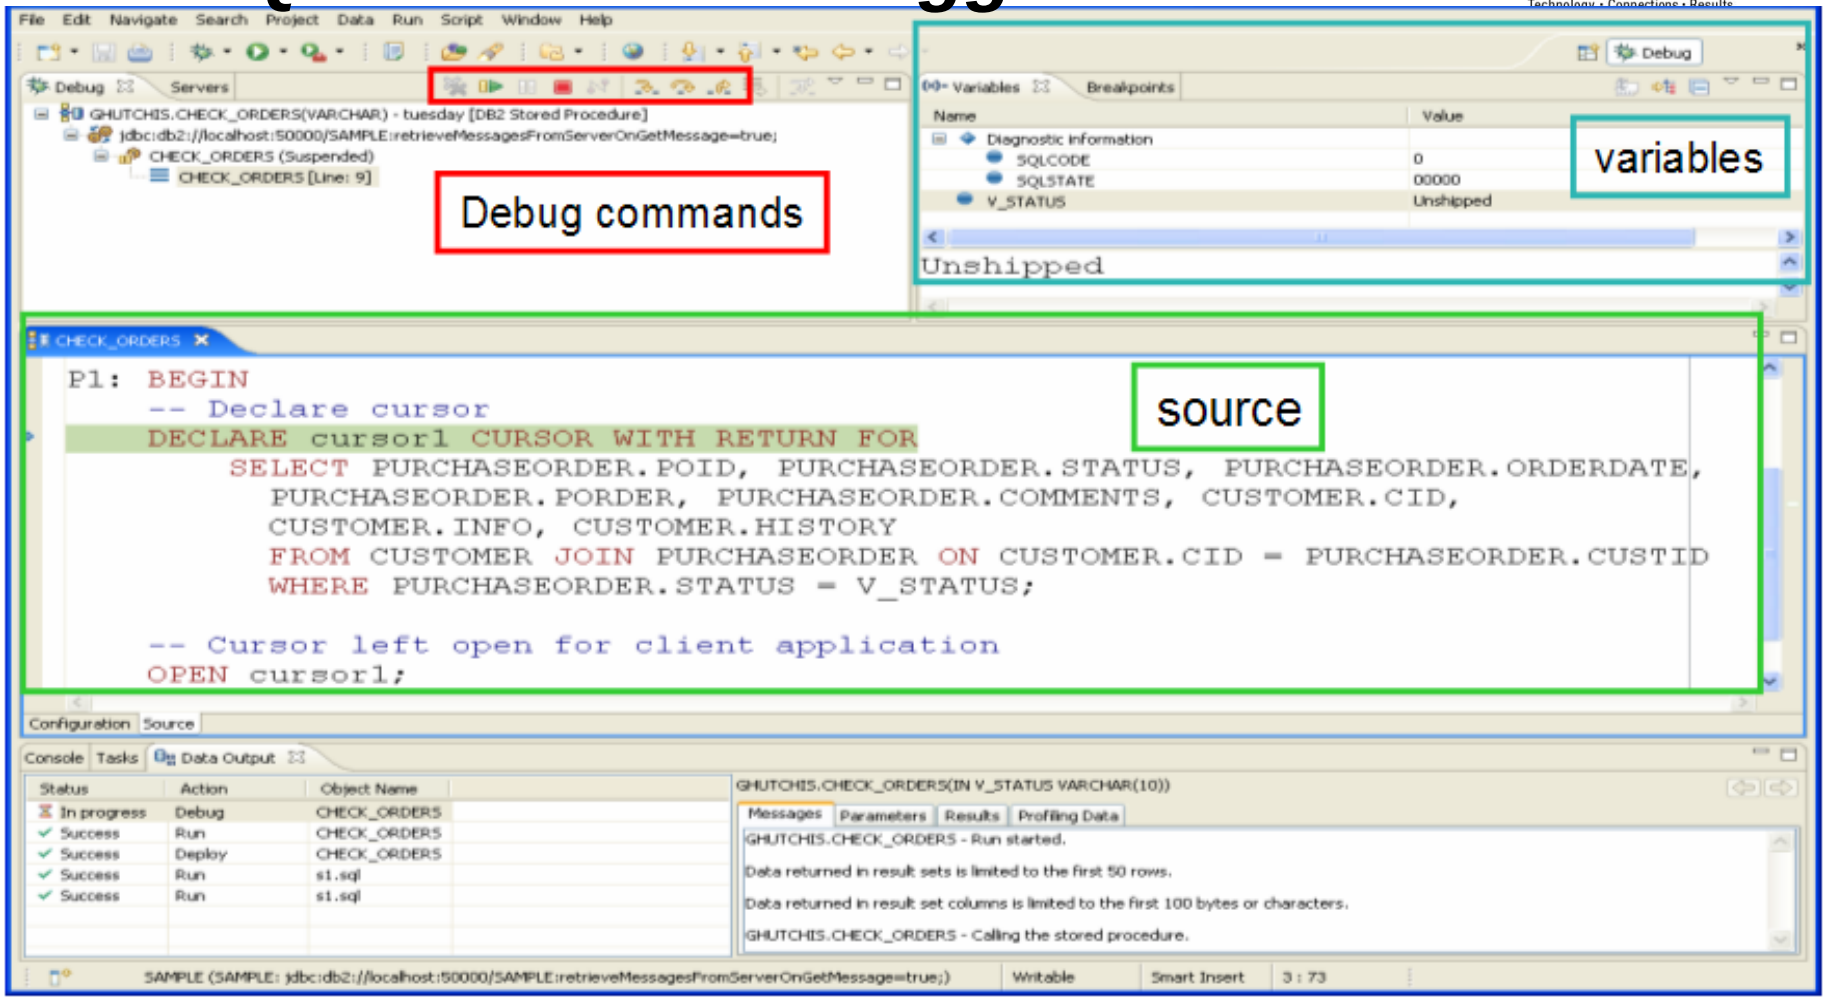

Integrated editor and debugger for SQL PL and Java (JDBC or SQLJ)

Data Management

Data Distribution Editor

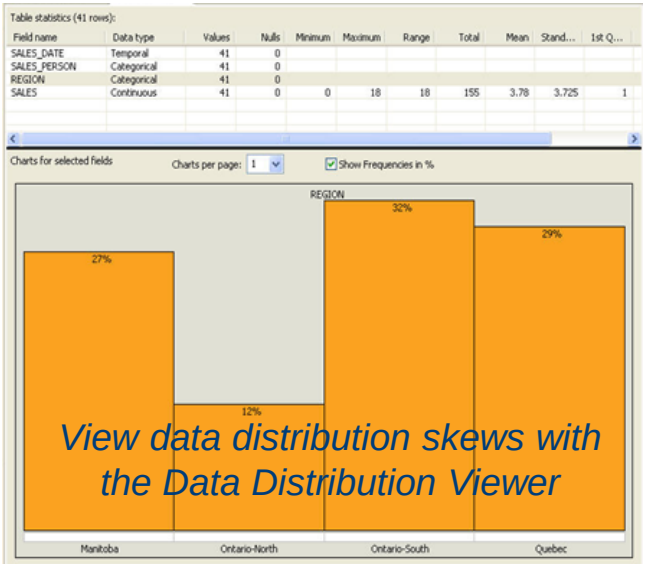

View data distribution skews with the Data Distribution Viewer

Edit Data

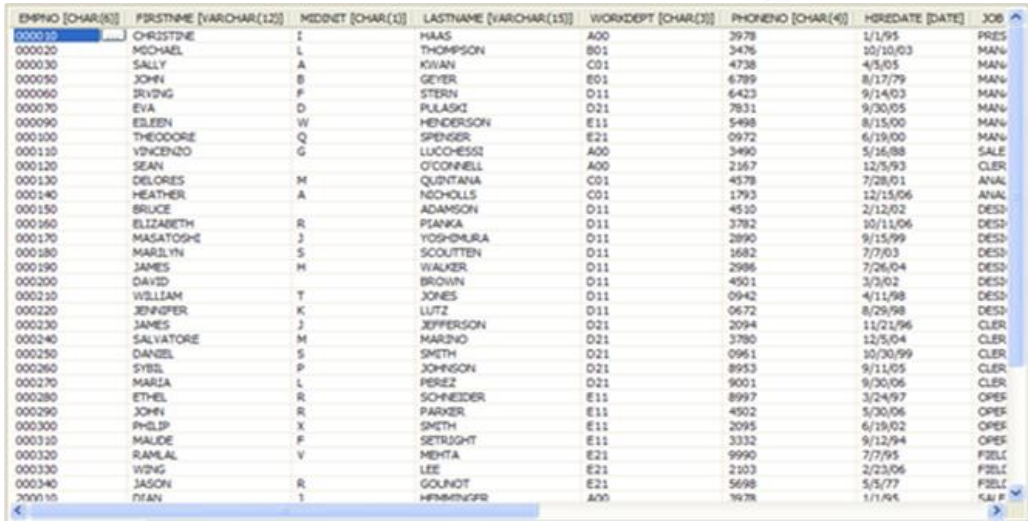

Roles, Users, Privileges

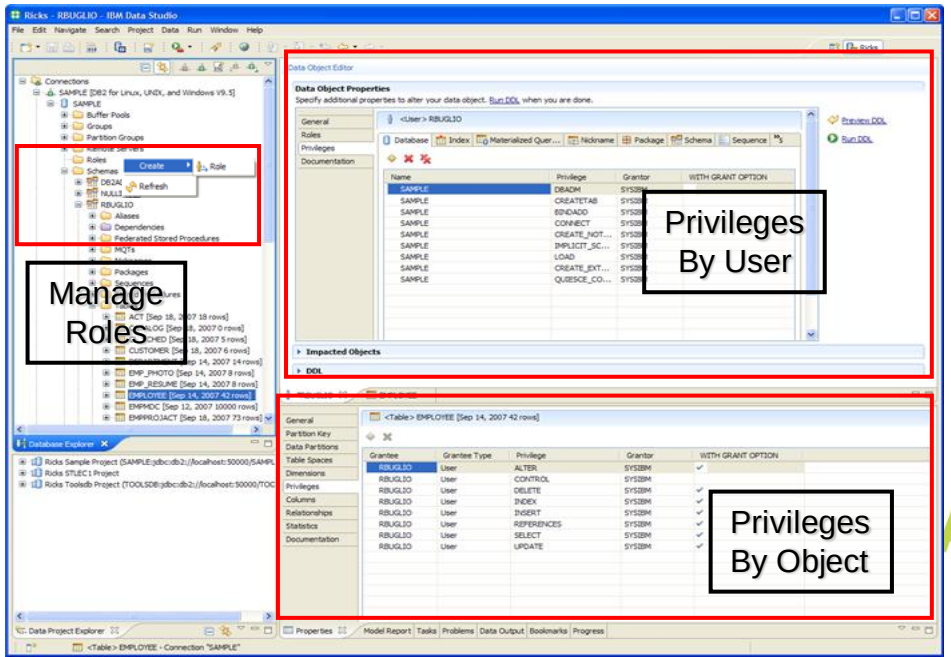

IBM Data Studio

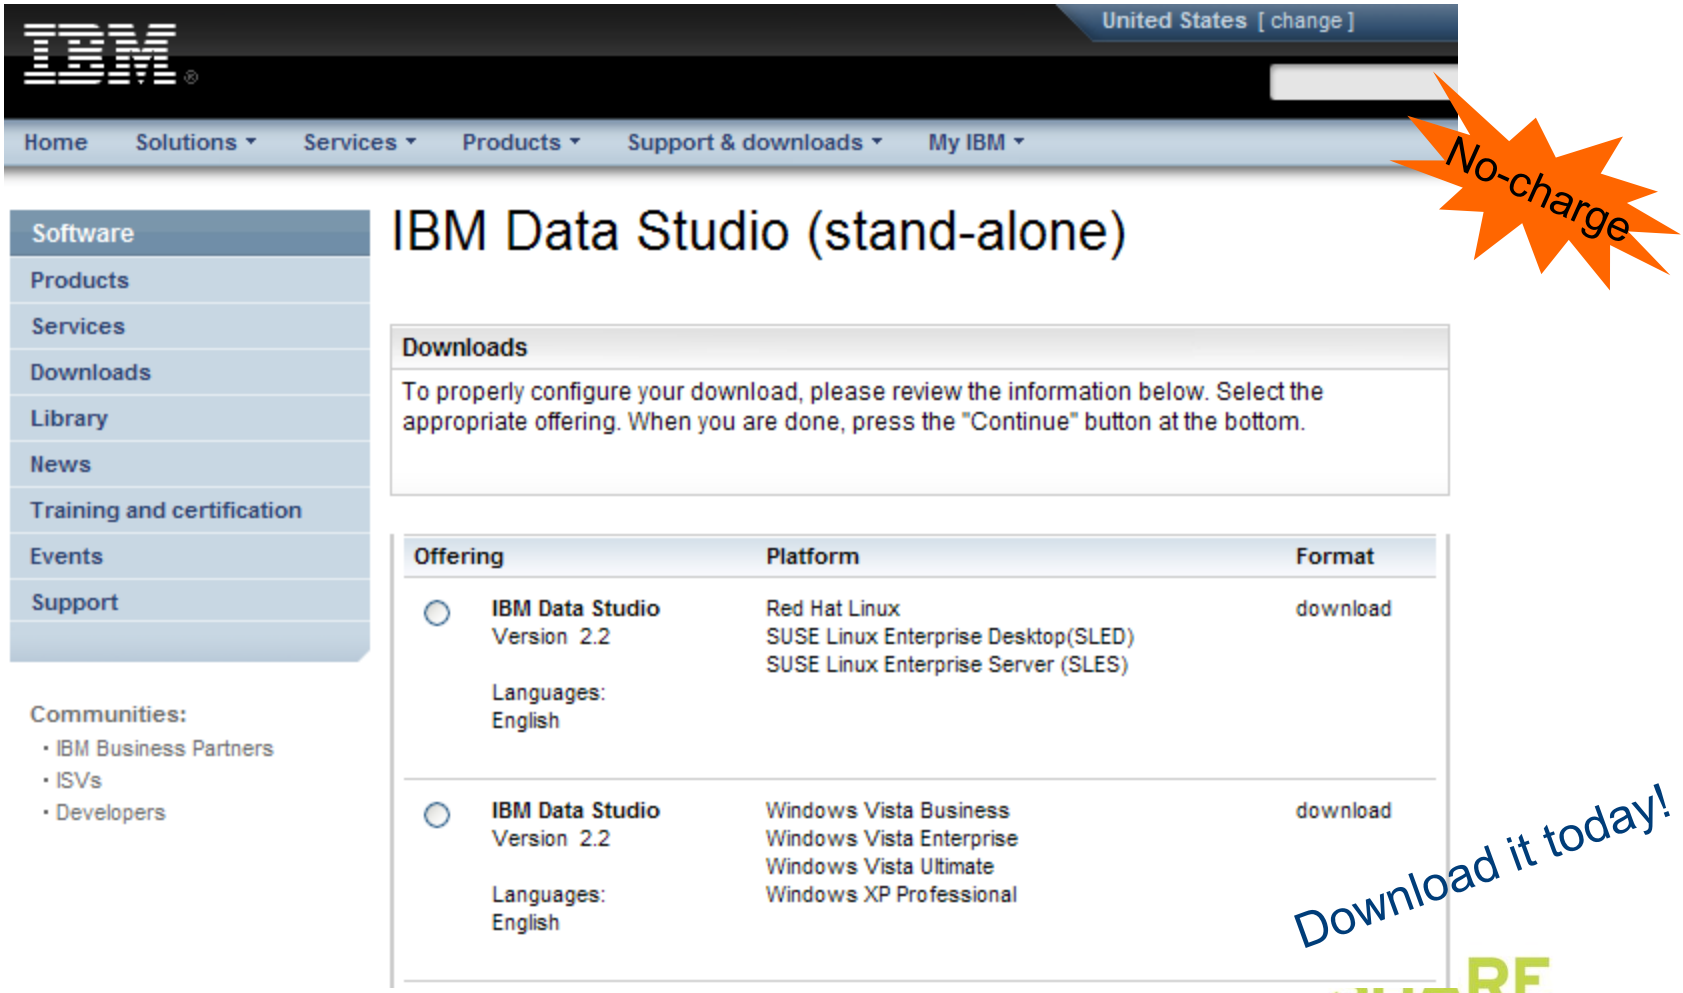

IBM Optim Development Studio

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

UNIVAR

"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, .NET, and CLI applications can be managed like CICS/IMS/COBOL 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, .NET, and CLI 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

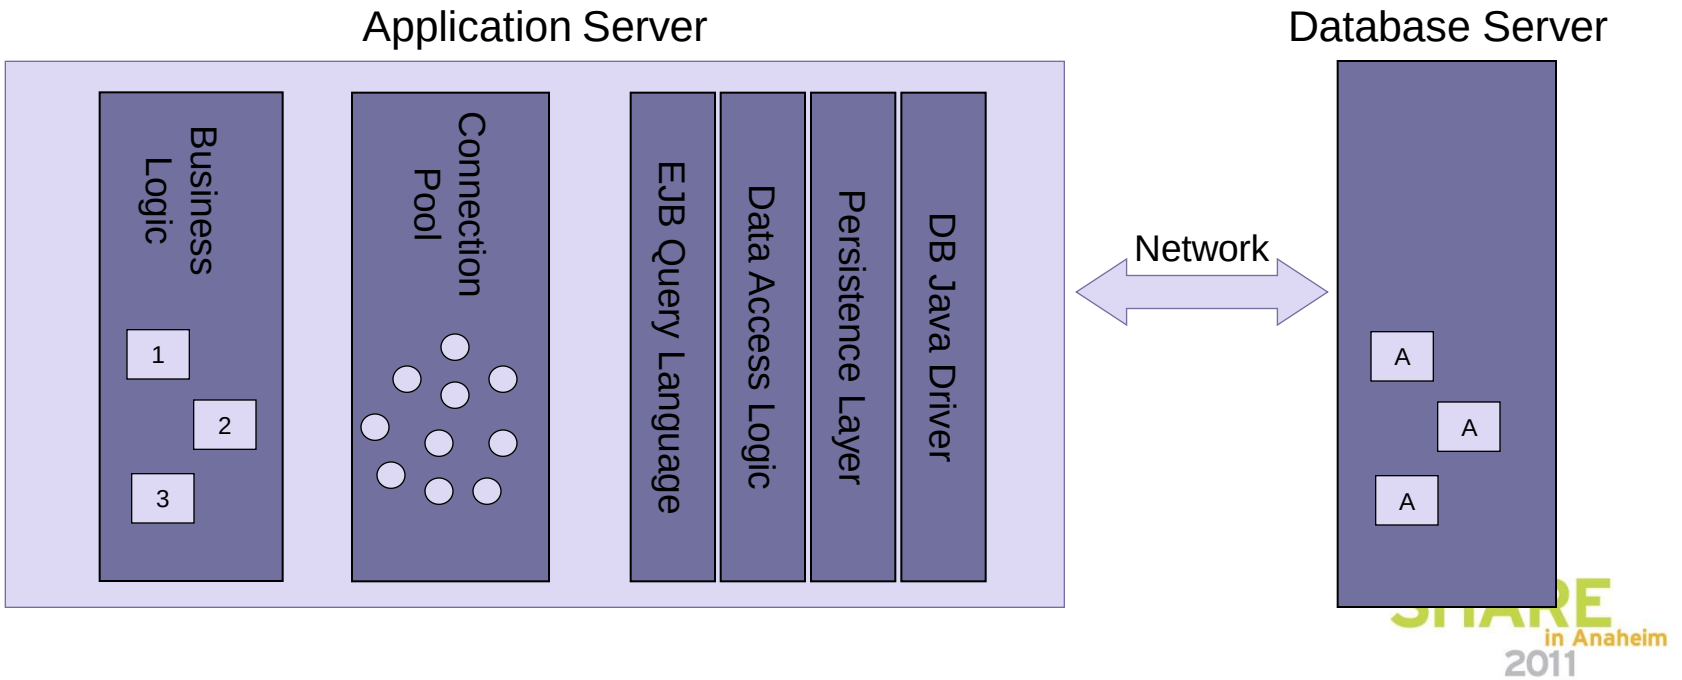

Introducing pureQuery

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

Java and Microsoft .NET

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

SHARE
Technology • Connections • Results

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

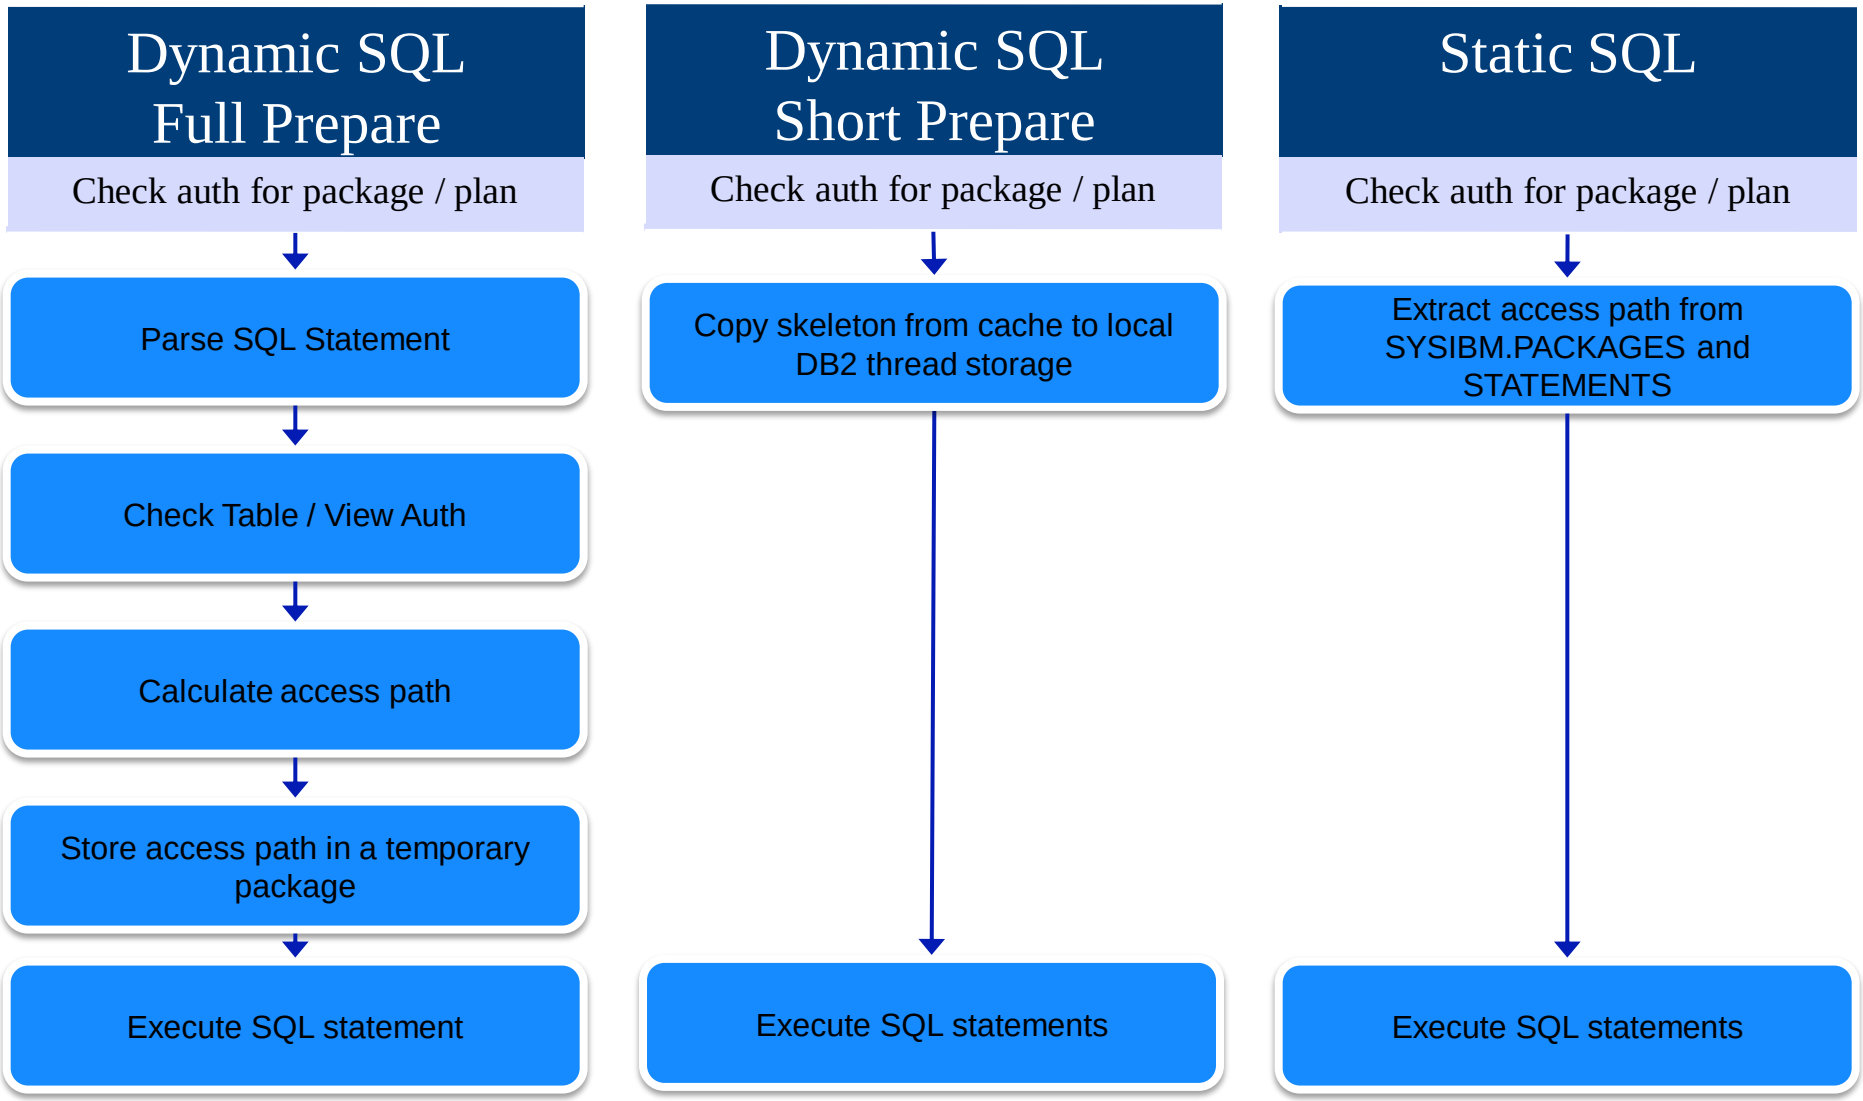

Static Execution Advantages

SHARE
Technology • Connections • Results

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

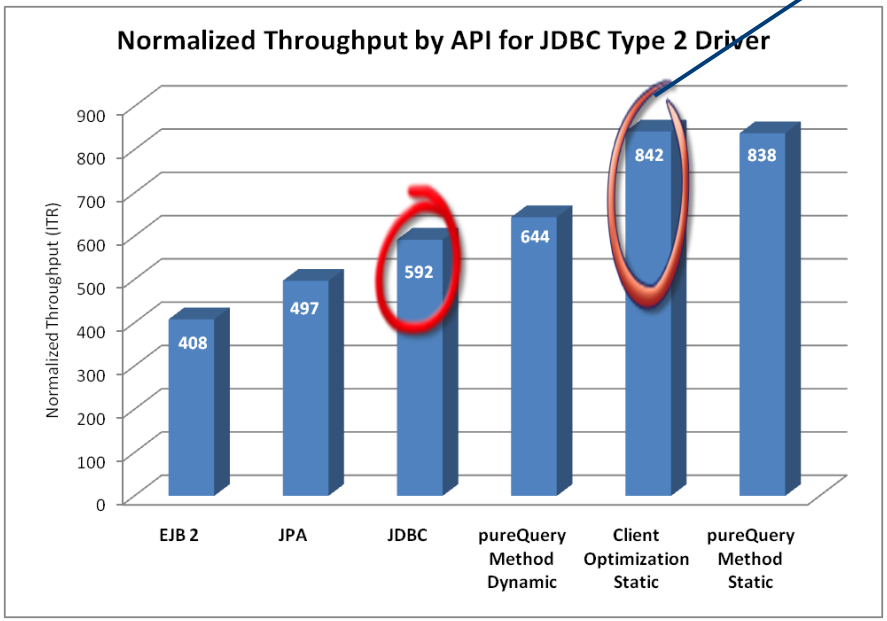

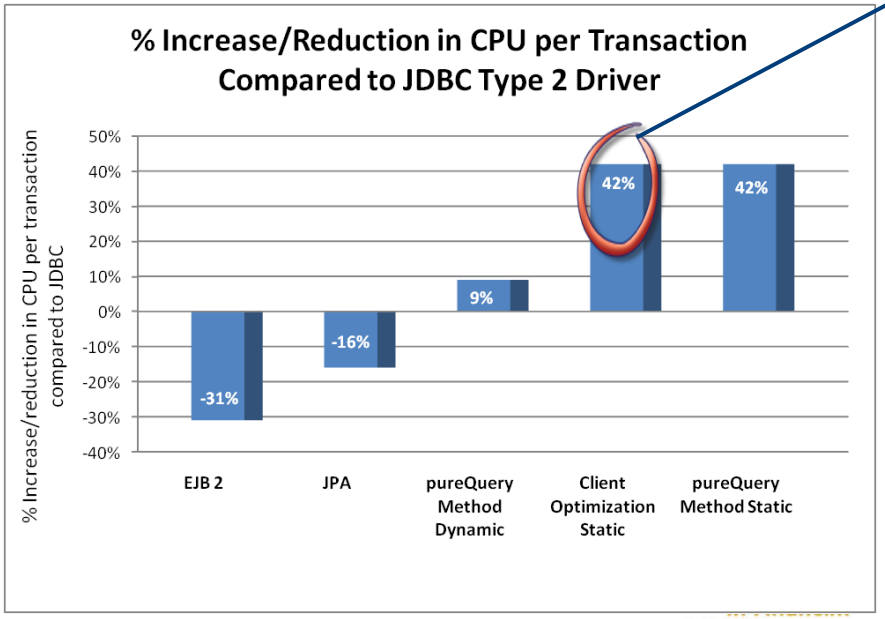

Java Database Access and pureQuery

Many on-ramps for new and existing applications

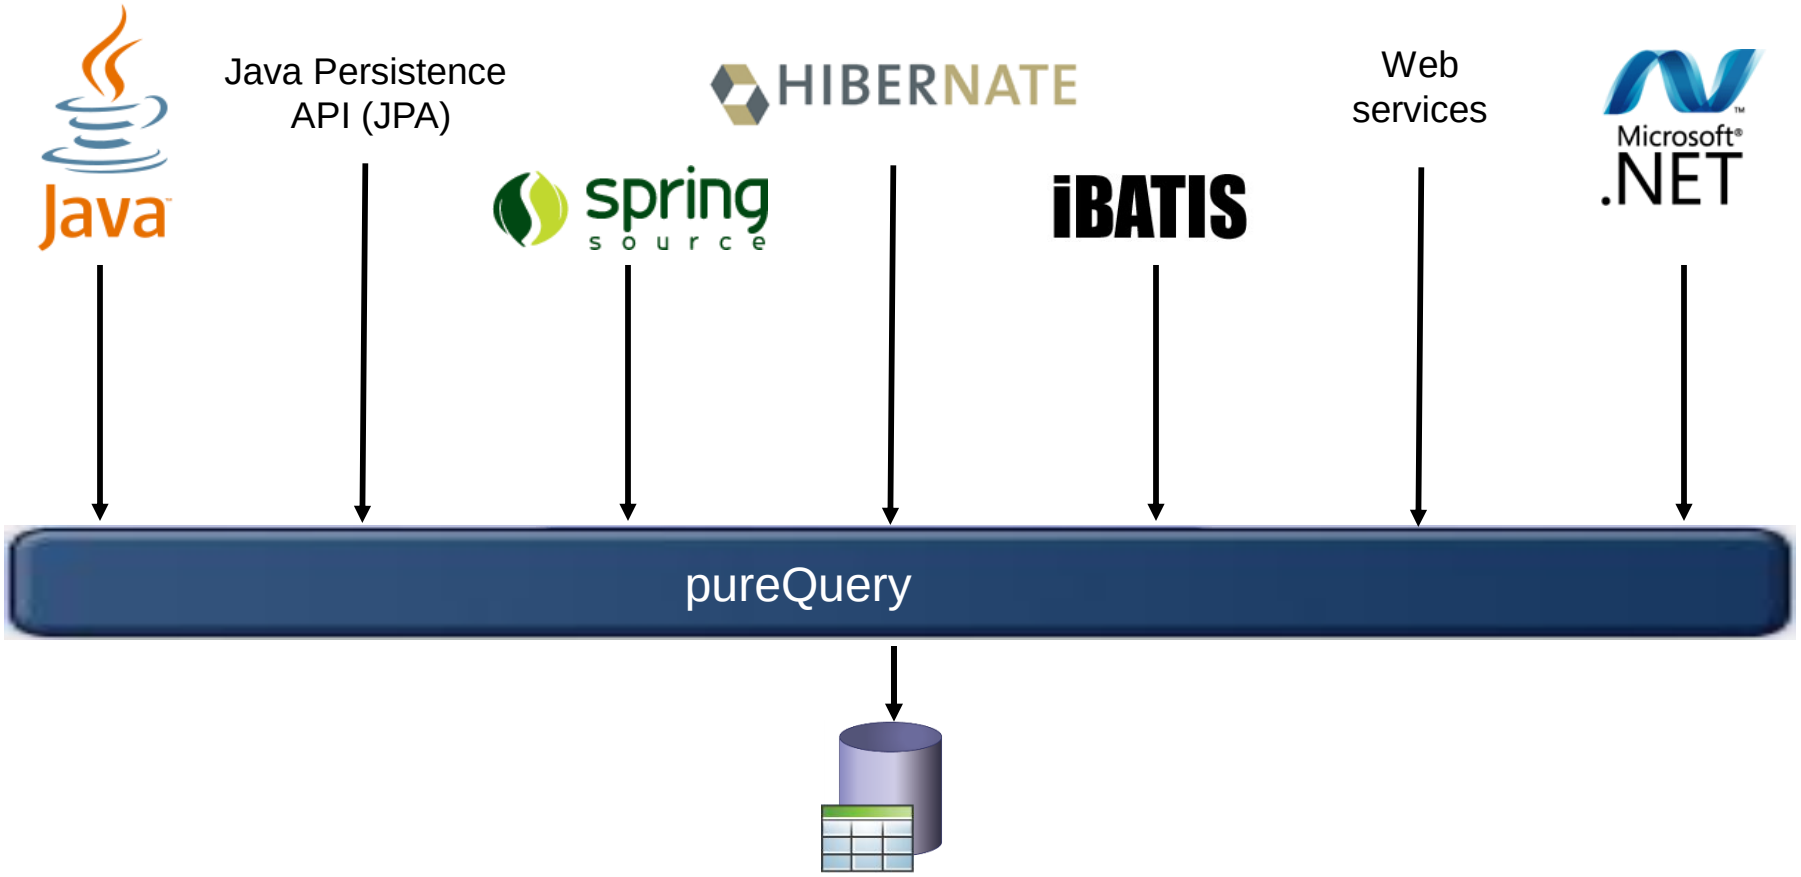

DB2 (z/OS, i & LUW), Informix, and Oracle now
More coming

Deploying with pureQuery Runtime

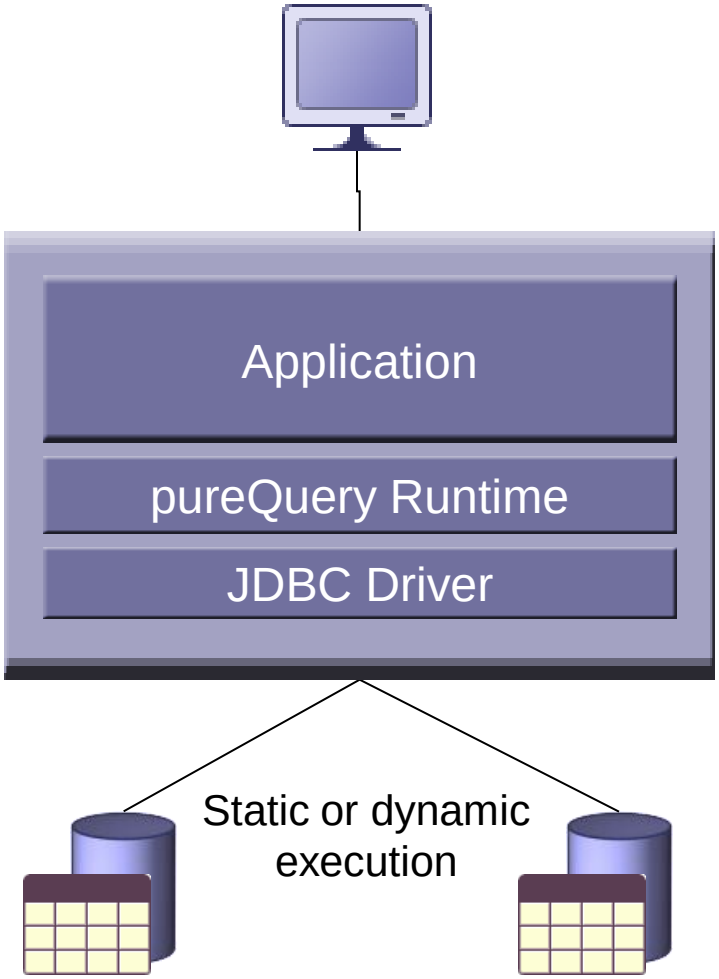

Application tier:

- 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

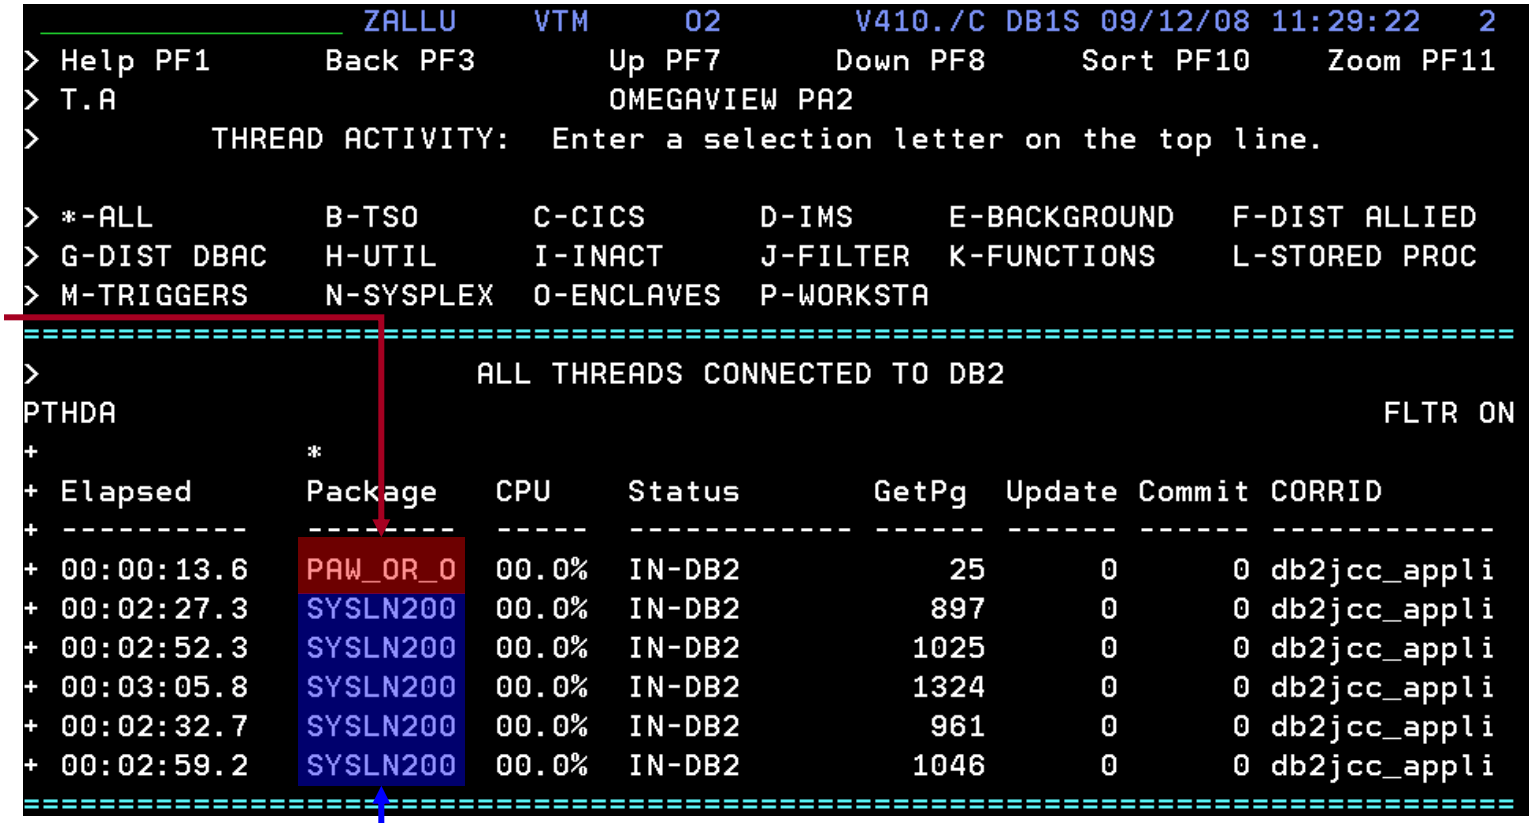

Static SQL for security – Administering table privileges

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

```
GRANT SELECT
ON TABLE PAYROLL
TO DEPT_D47
```

PAYROLL

NAME	POSITION	SALARY	…

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

```
GRANT EXECUTE
ON PACKAGE
POSITION_REPORT
TO DEPT_D47
```

PACKAGE

BIND

SELECT NAME, POSITION FROM PAYROLL…

```
GRANT SELECT
ON PAYROLL
TO BIND_ADMIN
```

PAYROLL

NAME	POSITION	SALARY	…

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

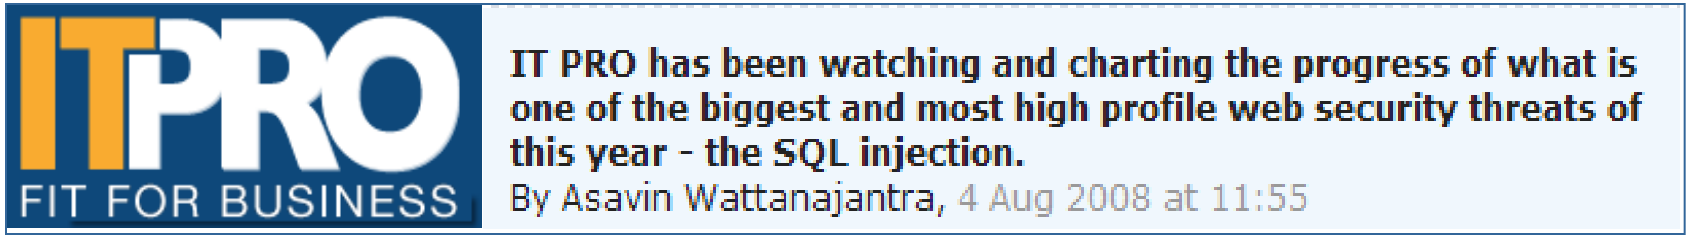

Capture Application SQL: At Development or Later

Capture → Review → Optimize → Revise → Restrict

Three methods

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

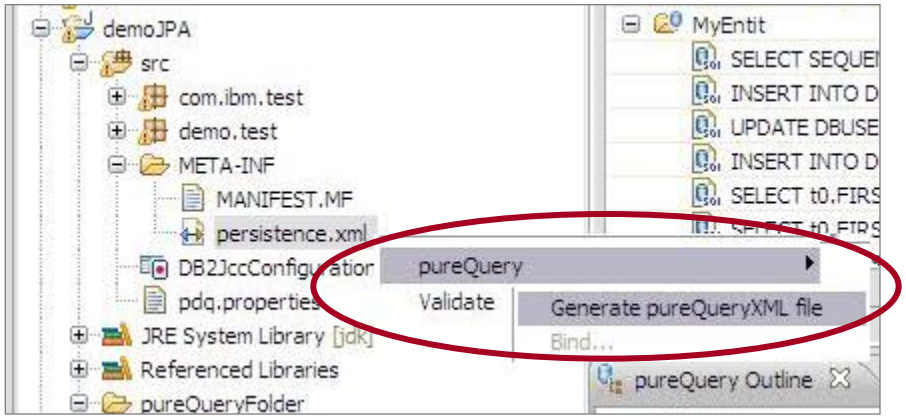

IBM Optim Development Studio

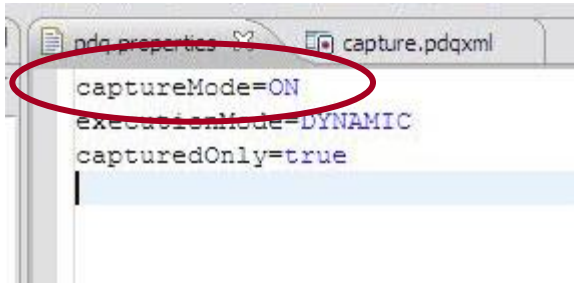

Visualize Application and SQL Metadata

Capture → **Review** → Optimize → Revise → Restrict

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

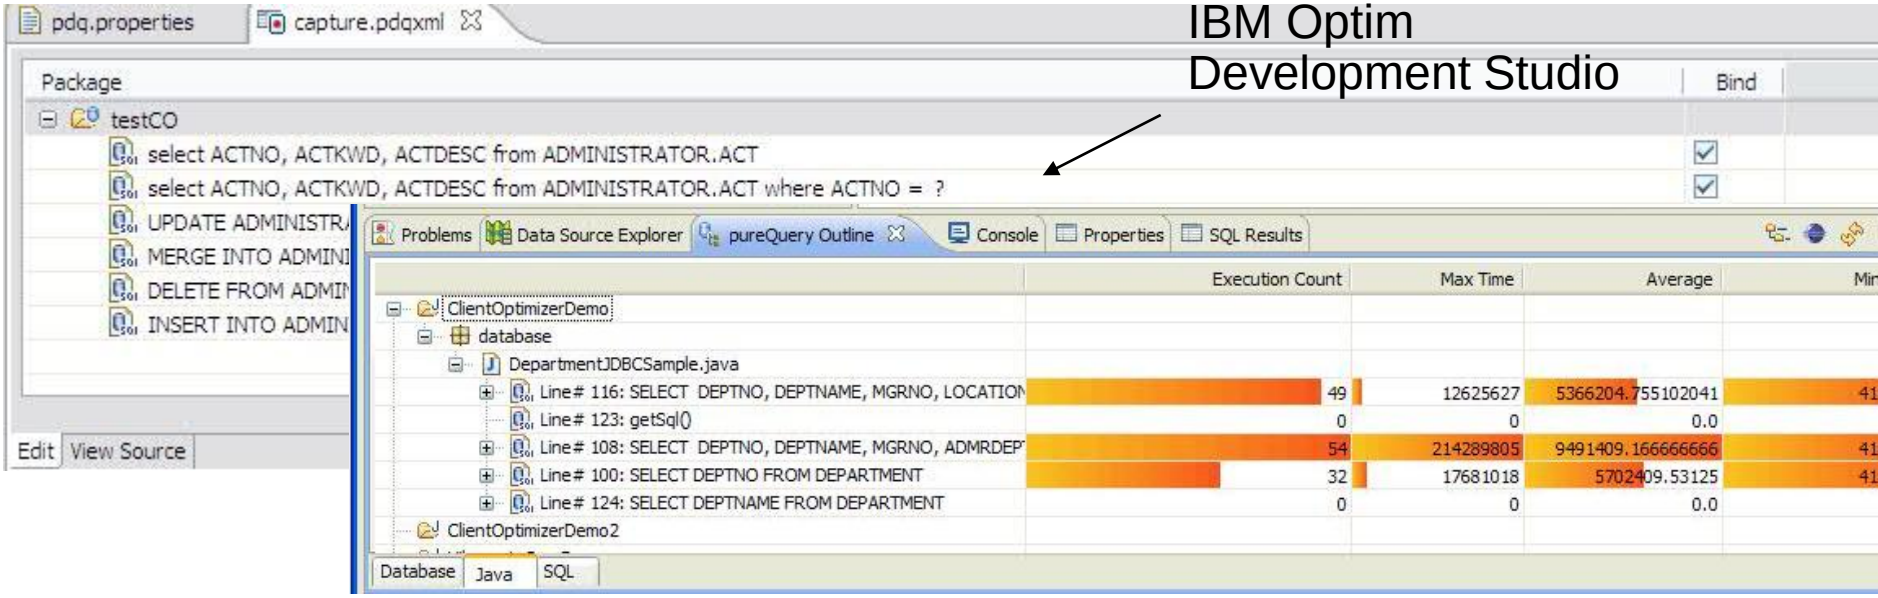

SQL Outline
Speed up problem isolation for developers – even when using frameworks

SHARE
Technology • Connections • Results

- 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

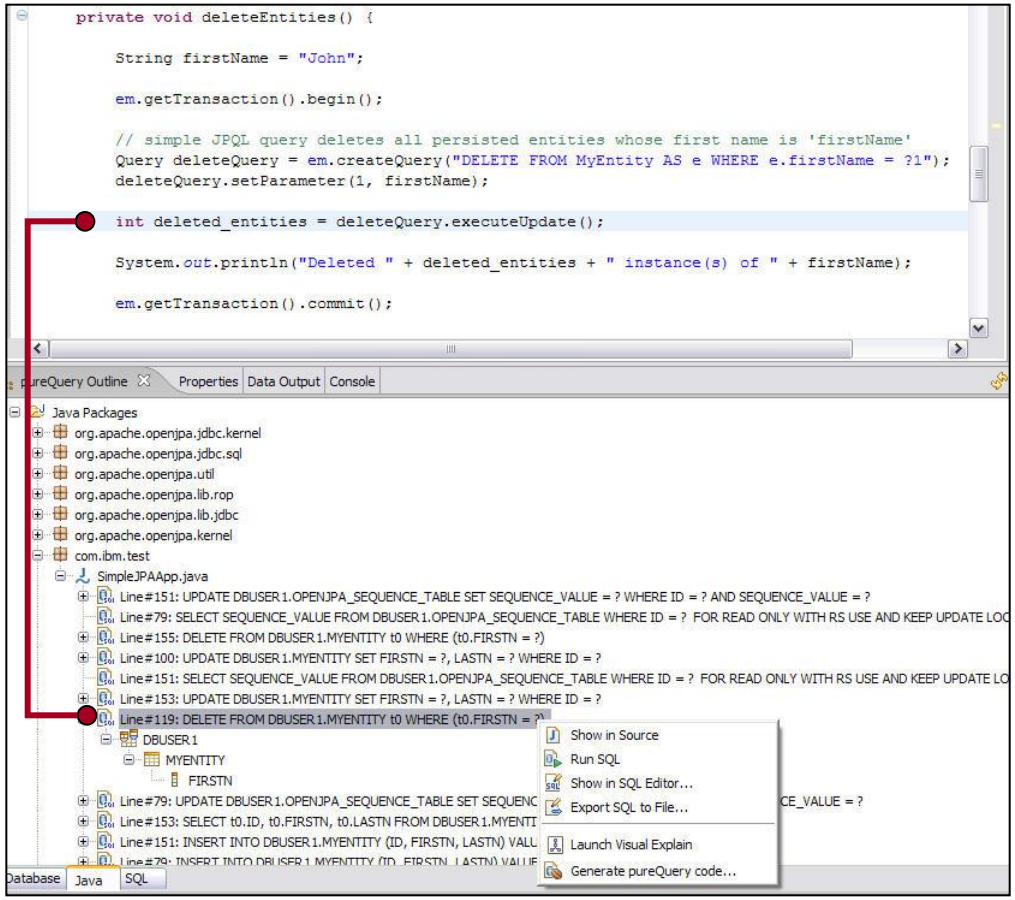

IBM Optim Development Studio

SHARE in Anaheim 2011

Optimize SQL

Capture → Review → **Optimize** → Revise → Restrict

- **Launch Visual Explain**

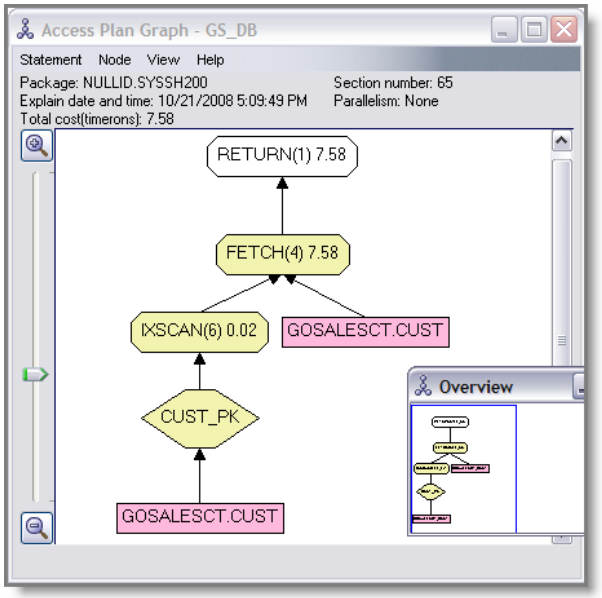

- **Copy SQL to Optim Query Workload Tuner**

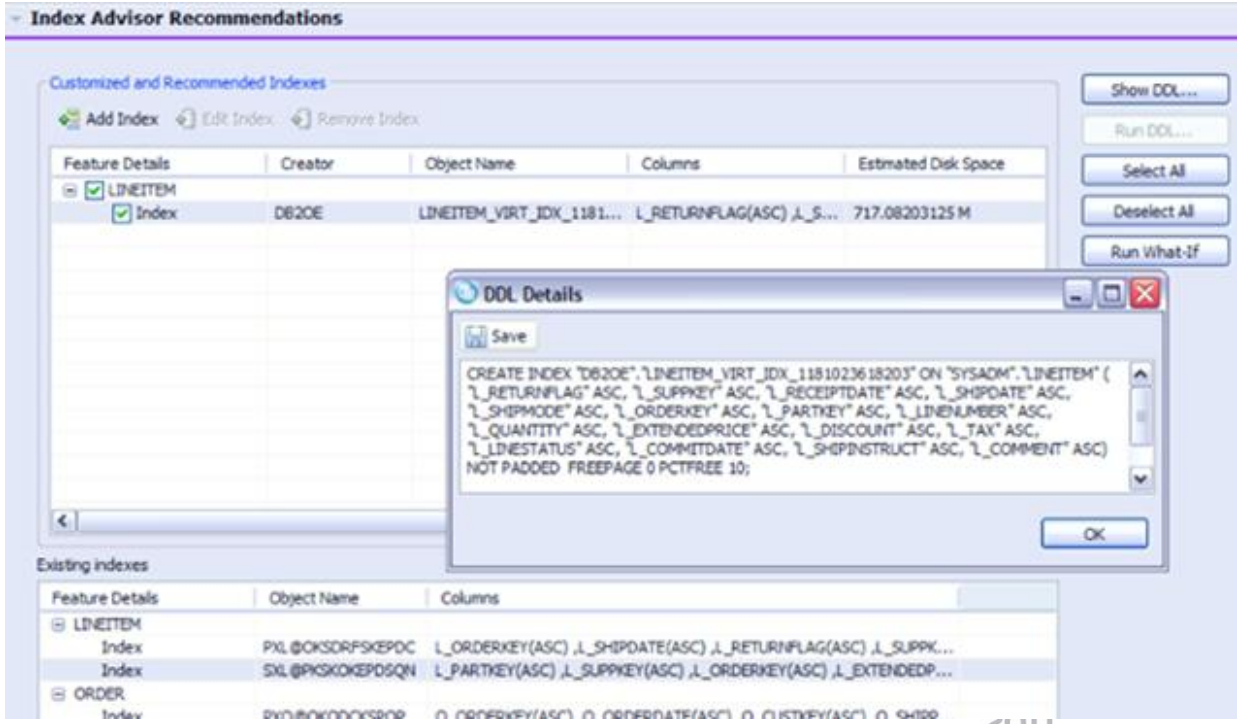

Query Tuning

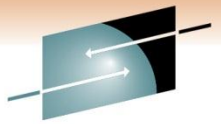

Introduction to query tuning
Query performance causes application outage

- 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

SHARE
Technology · Connections · Results

Define or select workload

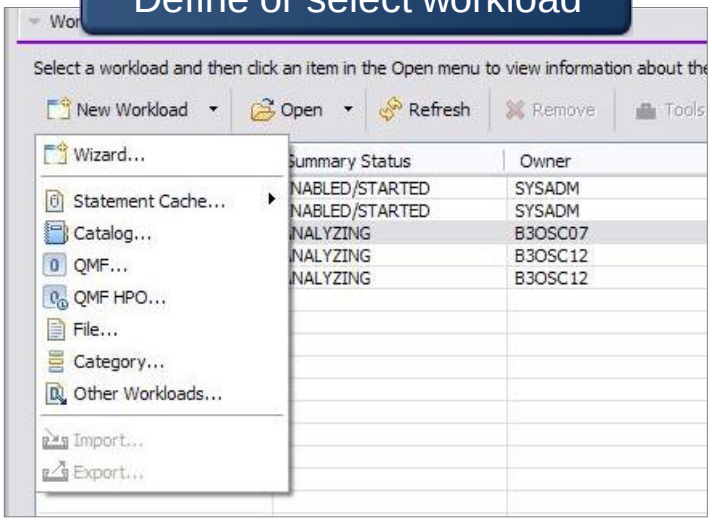

Execute advisors

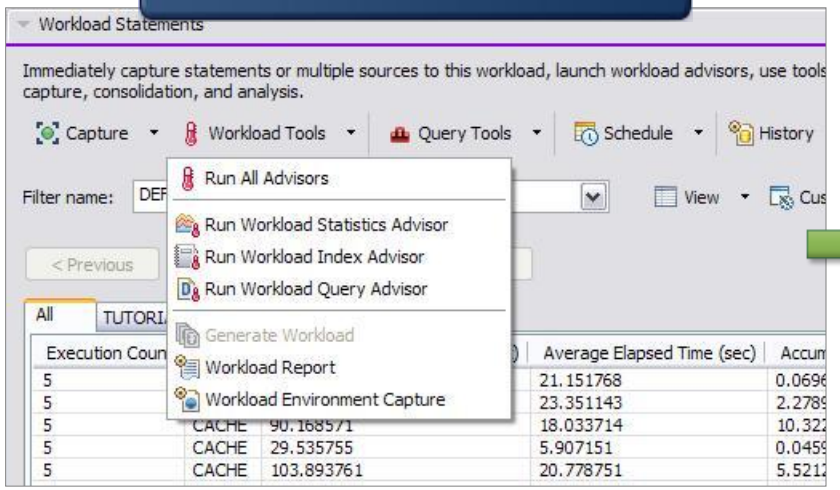

Drill into advice

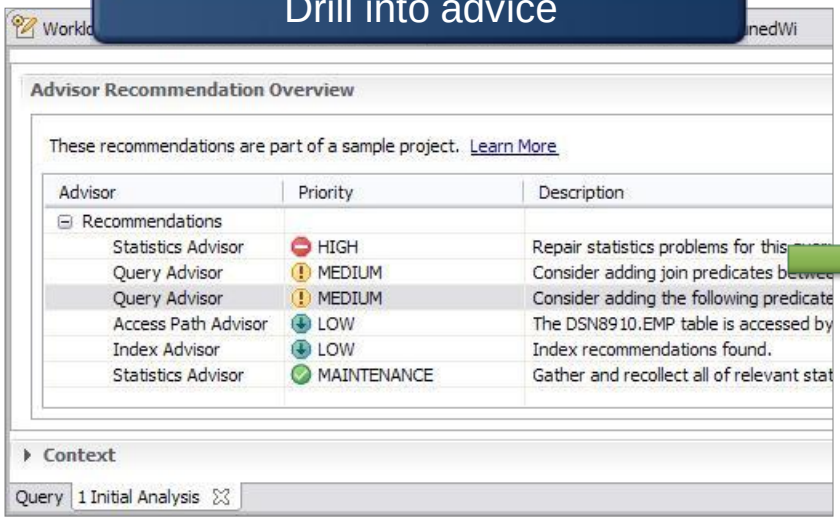

Validate improvement

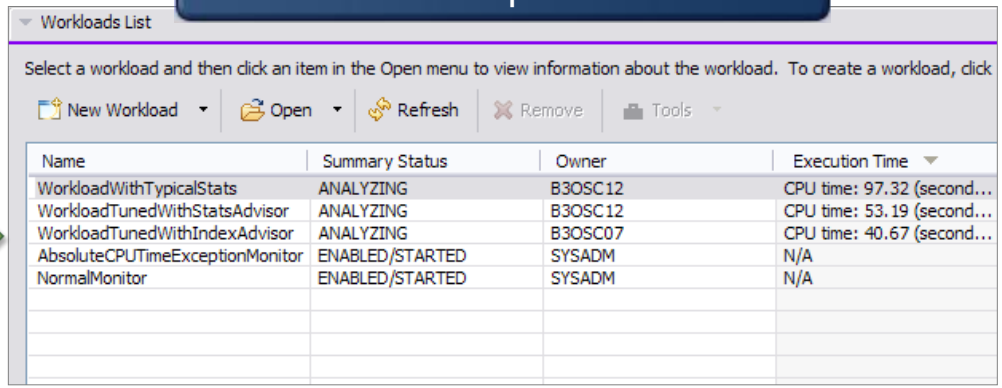

Gather High Cost Queries and Workloads

DB2

- Plan table
- Statement table
- Catalog plan or package
- Statement cache

Optim Development Studio

QMF and QMF HPO

DB2 Query Monitor

OMEGAMON XE for DB2

File, text, or exported workload

User defined category

Define or select workload

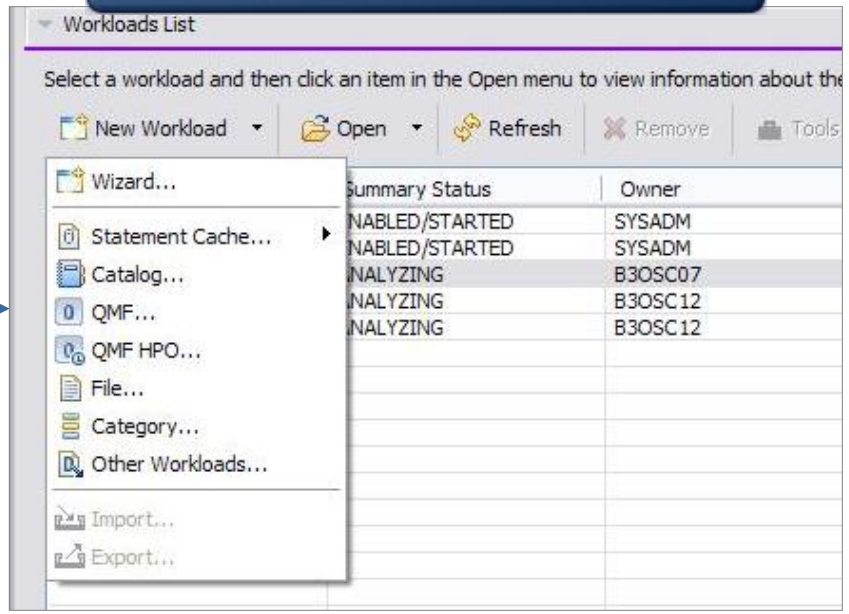

Visualize Queries and Costs to Speed Analysis

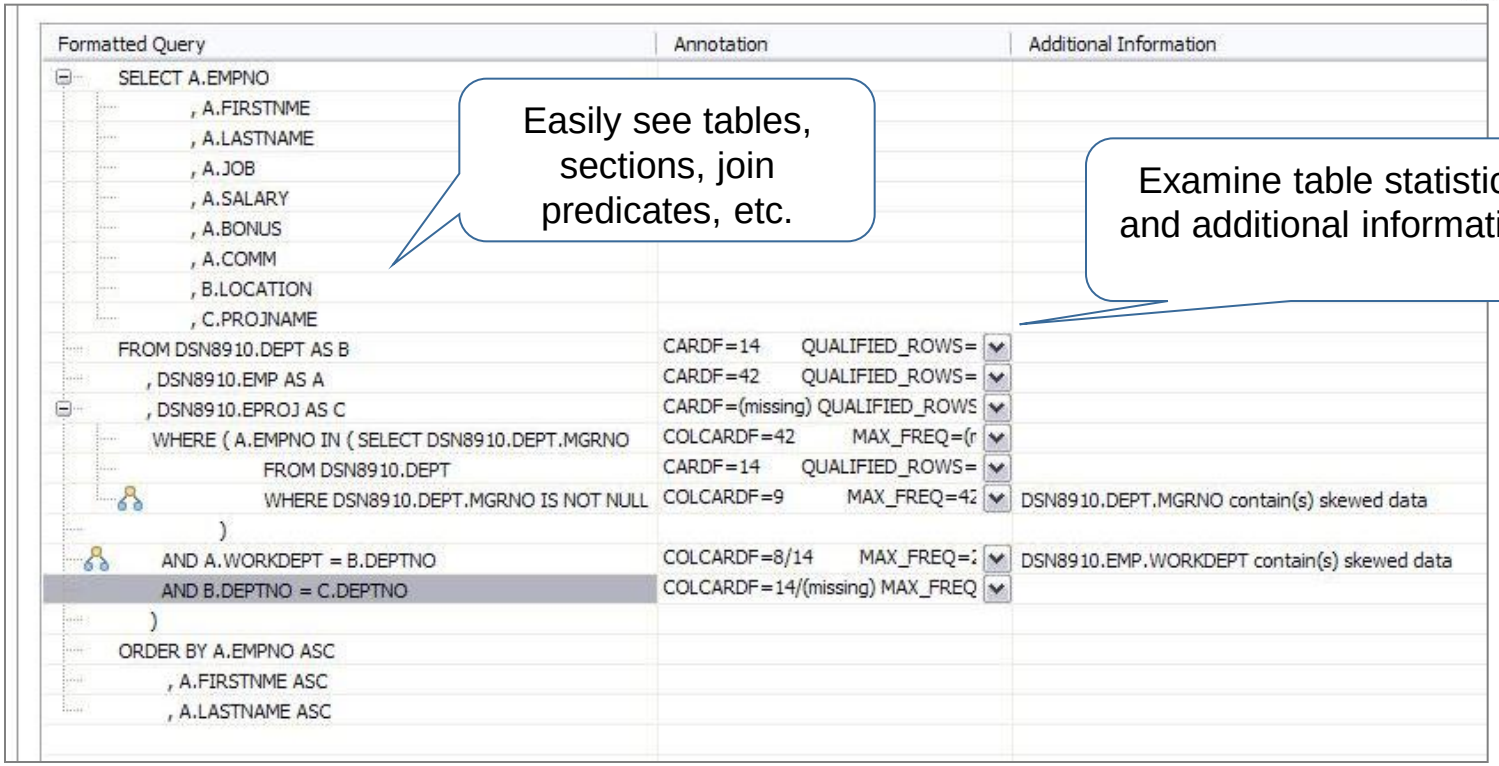

- Accelerate analysis, reduce downtime
  - Spot human errors
  - Identify where filtering should occur

View Optimizer Transformations

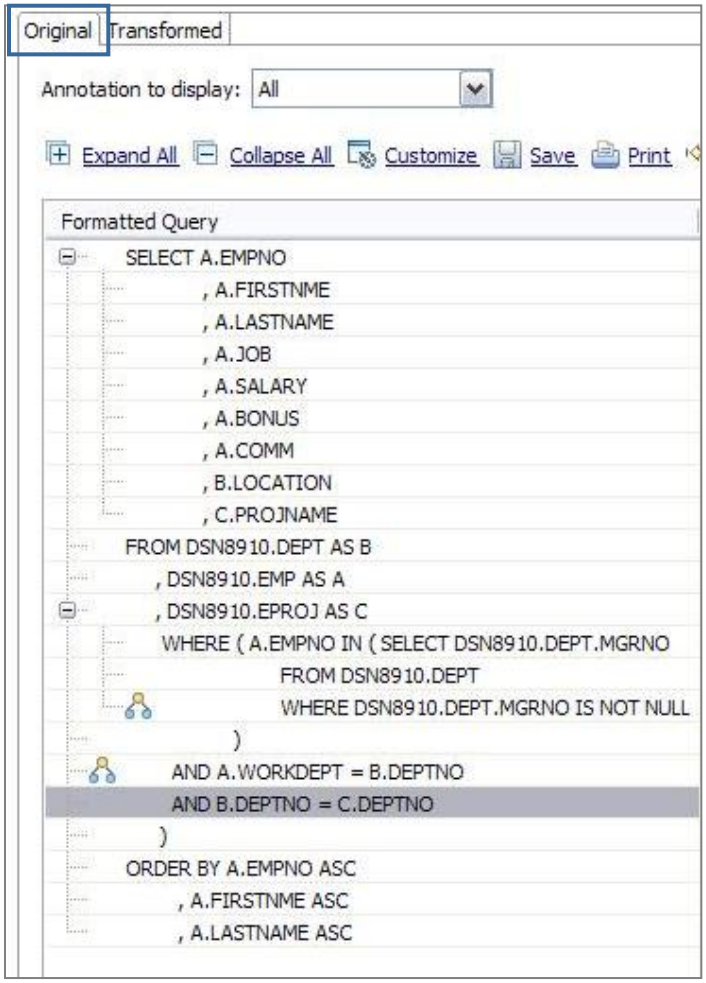

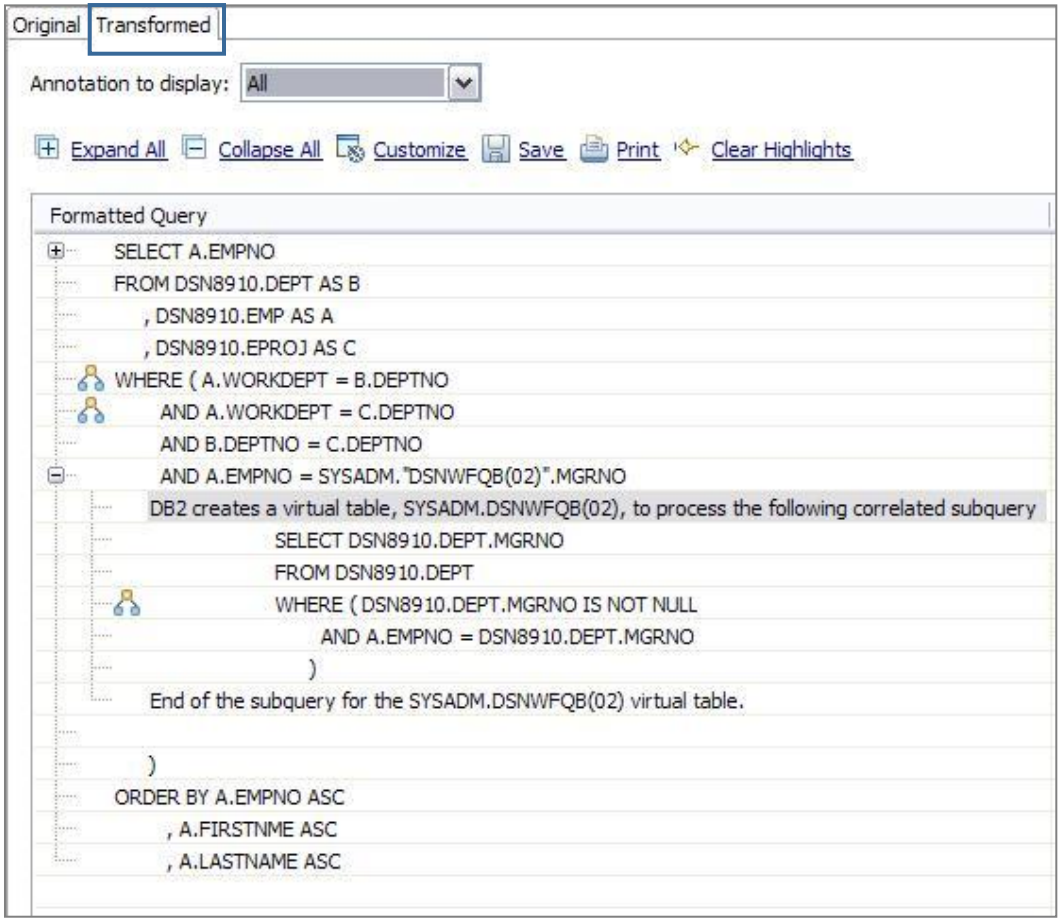

- Accelerate analysis, reduce downtime
  - Spot transformations which will occur

Execute Advisors

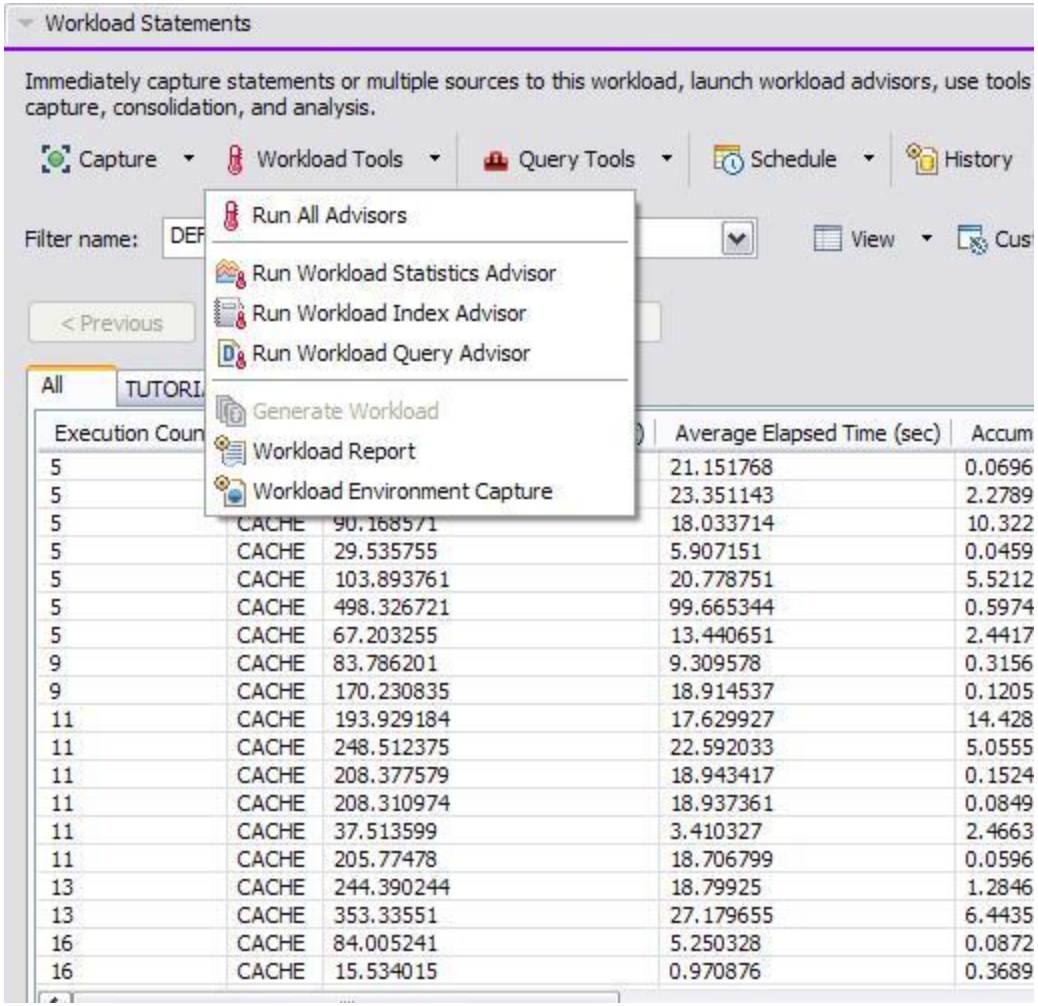

- Statistics
  - Get recommendations on the best statistics to capture to influence access path selection
- Query
  - Get recommendations regarding how to rewrite the query for better efficiency
- Index
  - Get recommendations on additional indexes that can reduce database scans

Analyze Access Plans

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

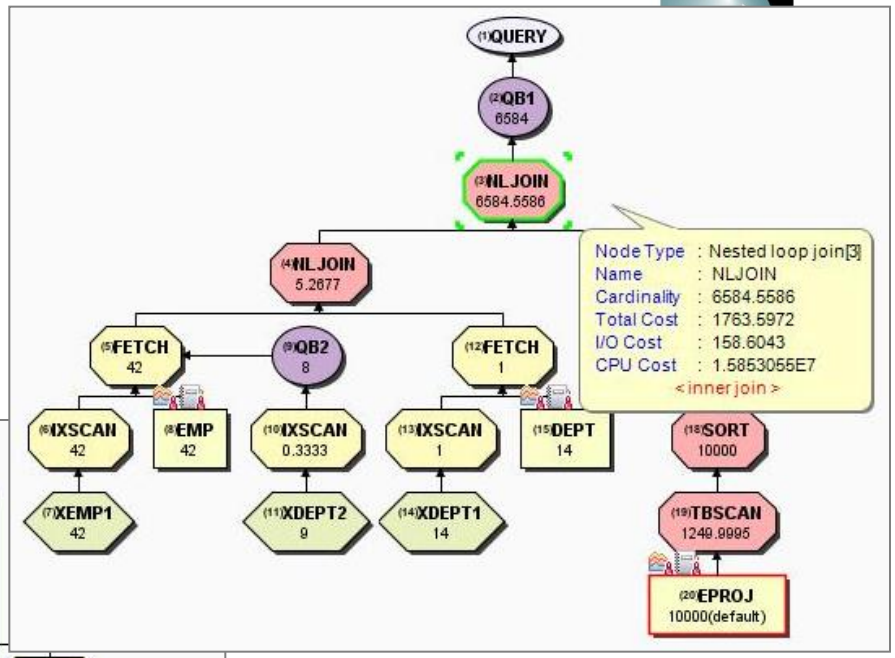

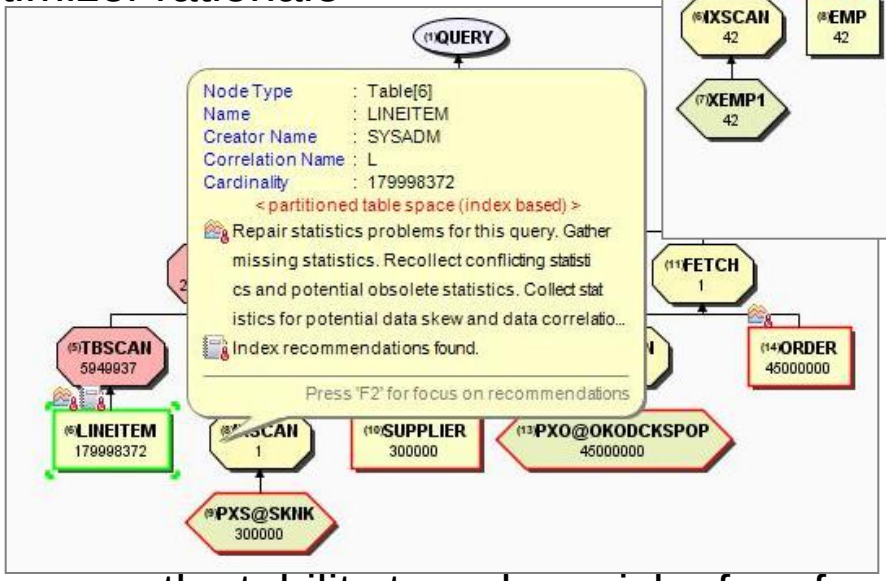

- Assess access path stability to reduce risk of performance regression
  - Is the optimizer able to apply the filtering early?
  - Are there indexes that support an efficient path?
  - Do statistics allow distinction between the choices?

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

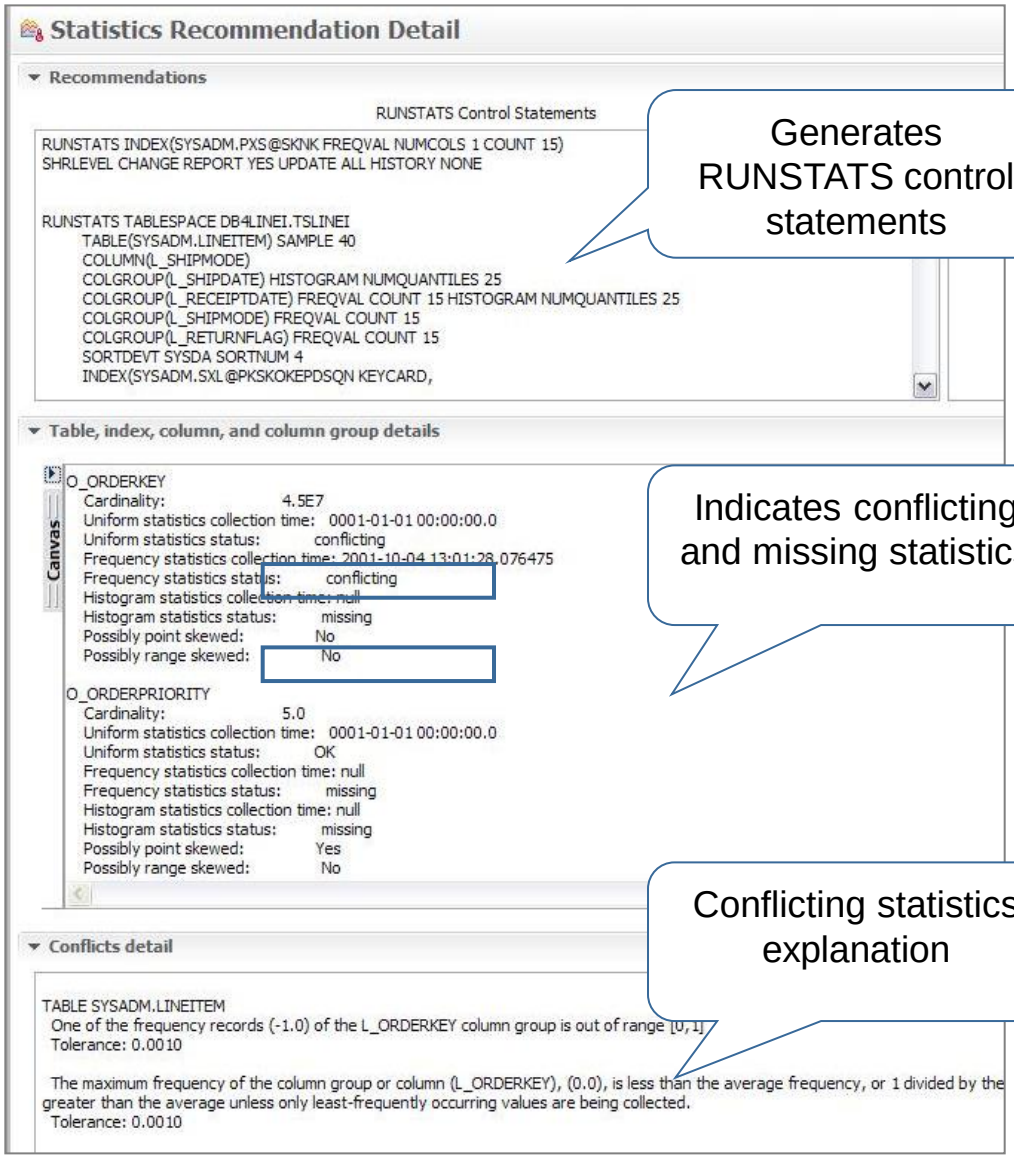

Improve Query Design

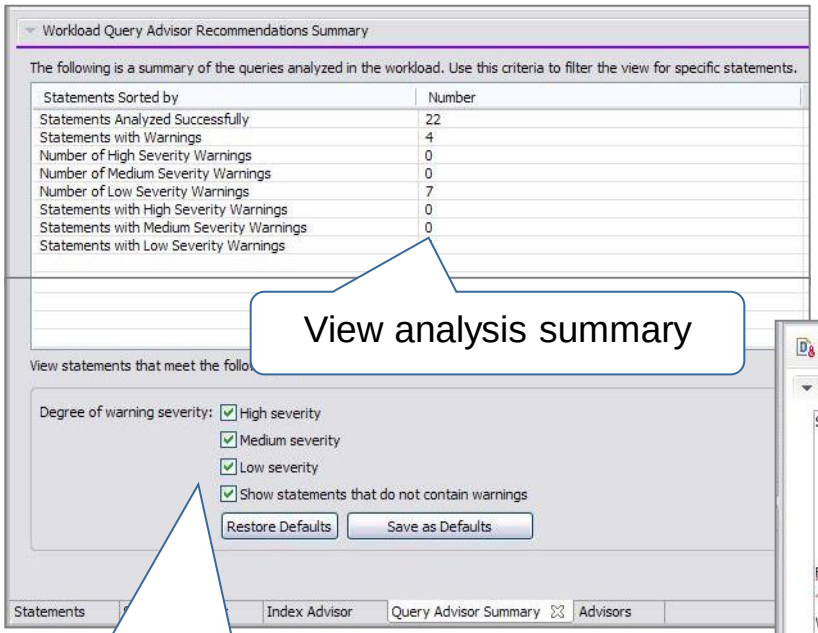

View analysis summary

Filter recommendations by severity

Guard against errors and oversights:

Further constrain query, increase index utilization, and reduce data reads

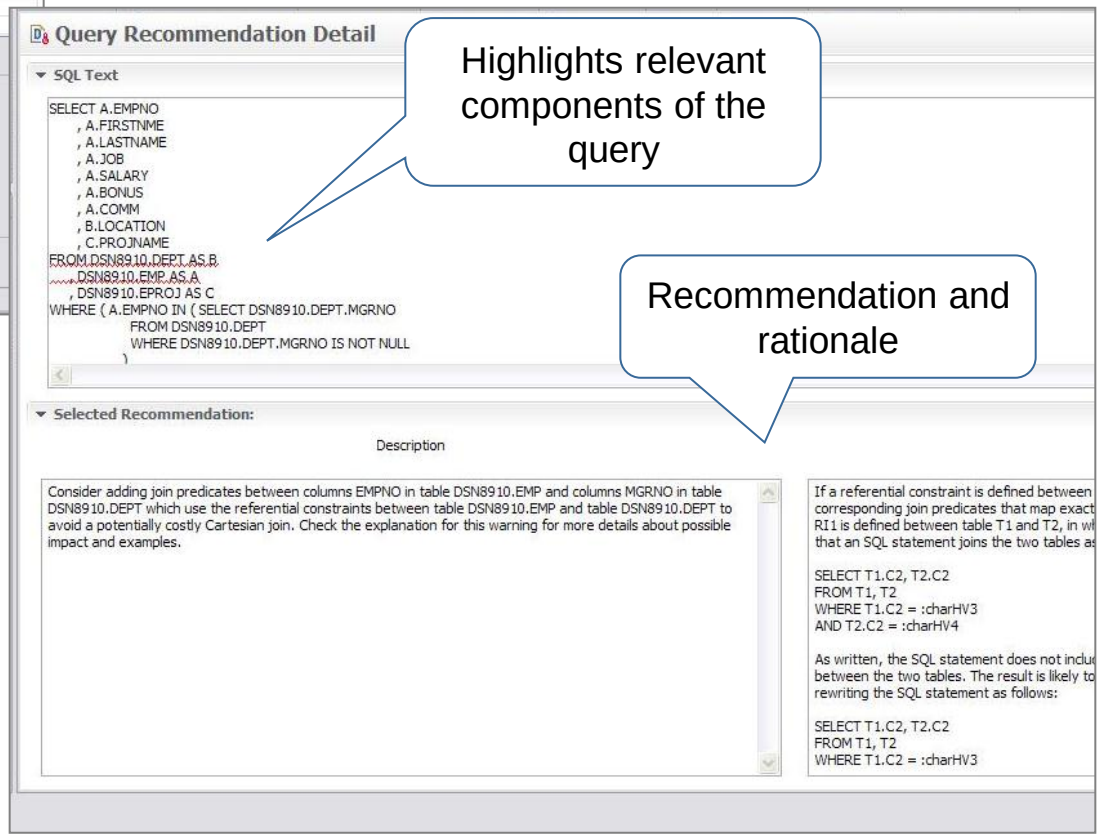

Highlights relevant components of the query

Recommendation and rationale

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

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

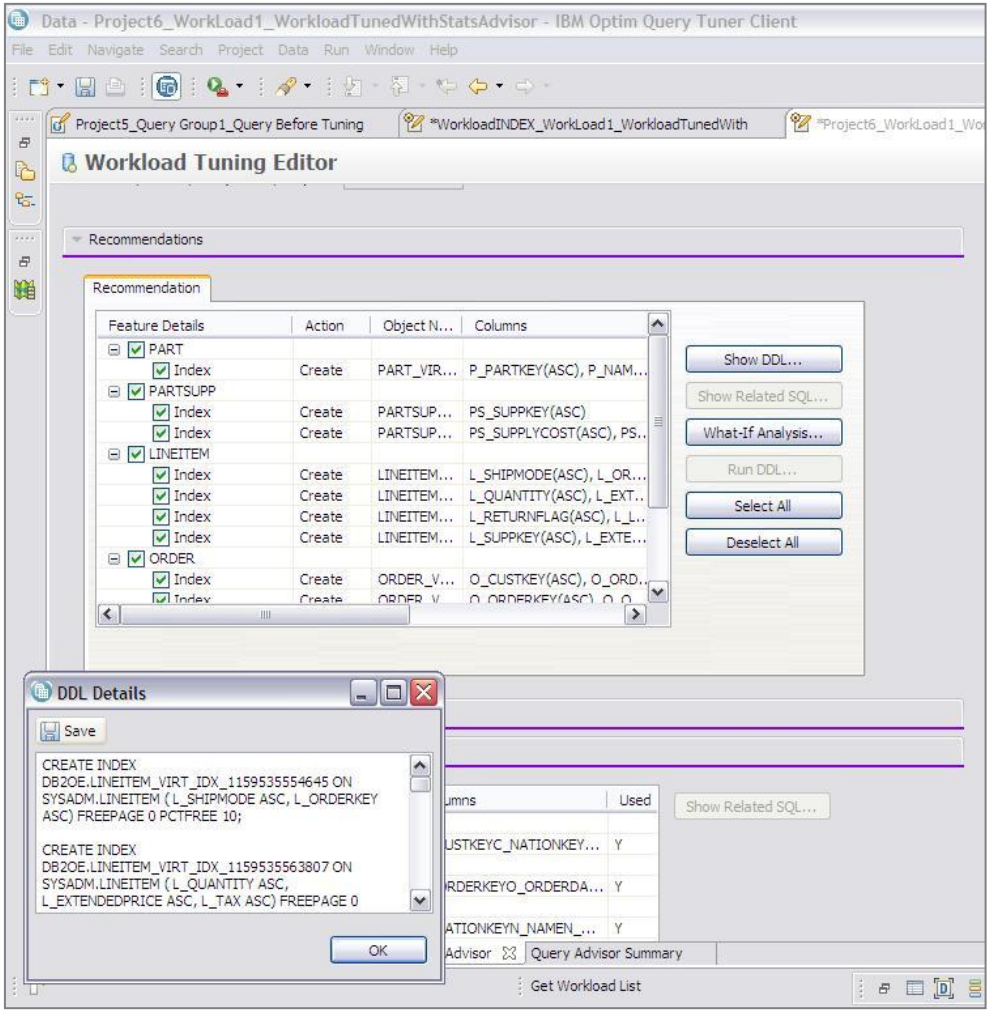

Visual Plan Hints for Experienced DBAs

- 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

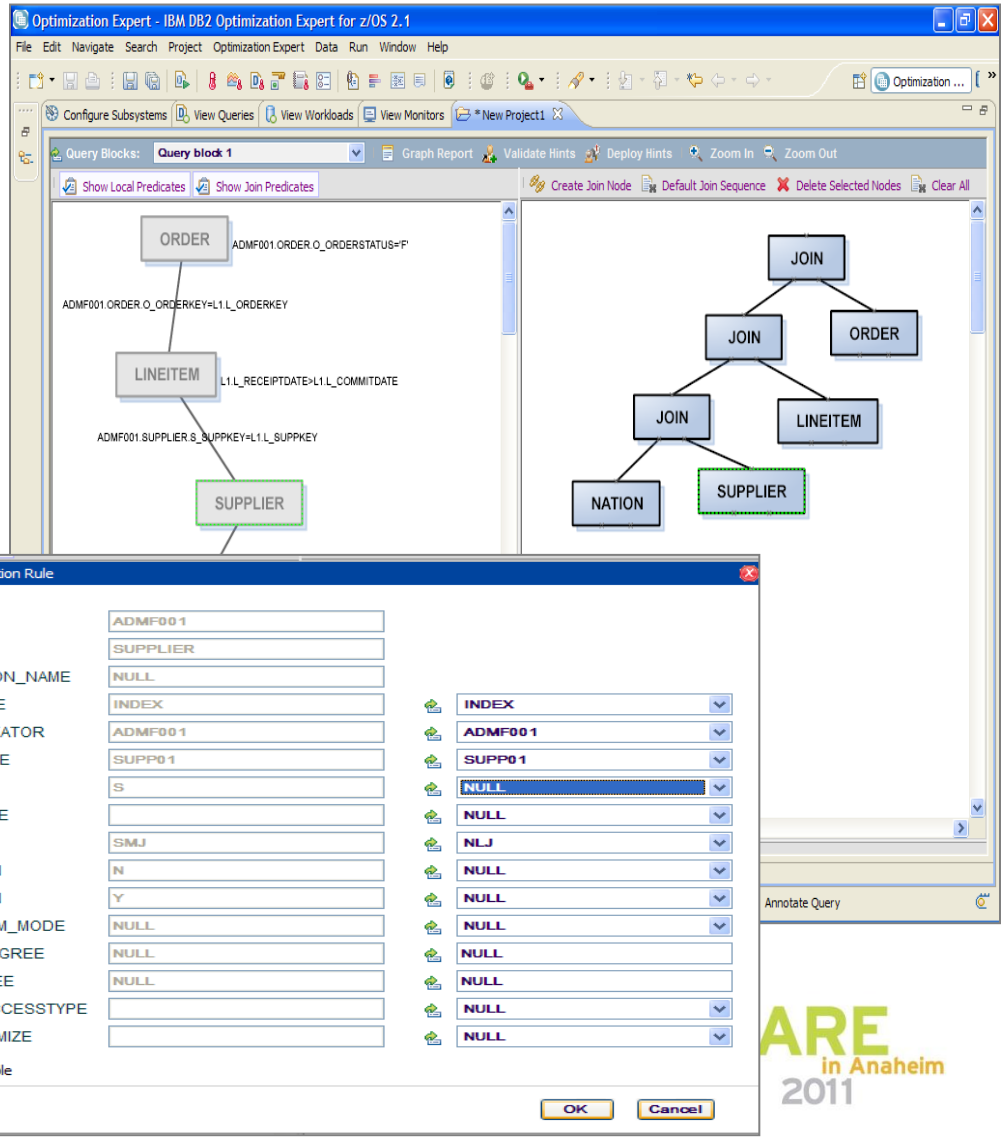

Query tuning tools for z/OS – Unique benefits

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

Capture → Review → Optimize → **Revise** → Restrict

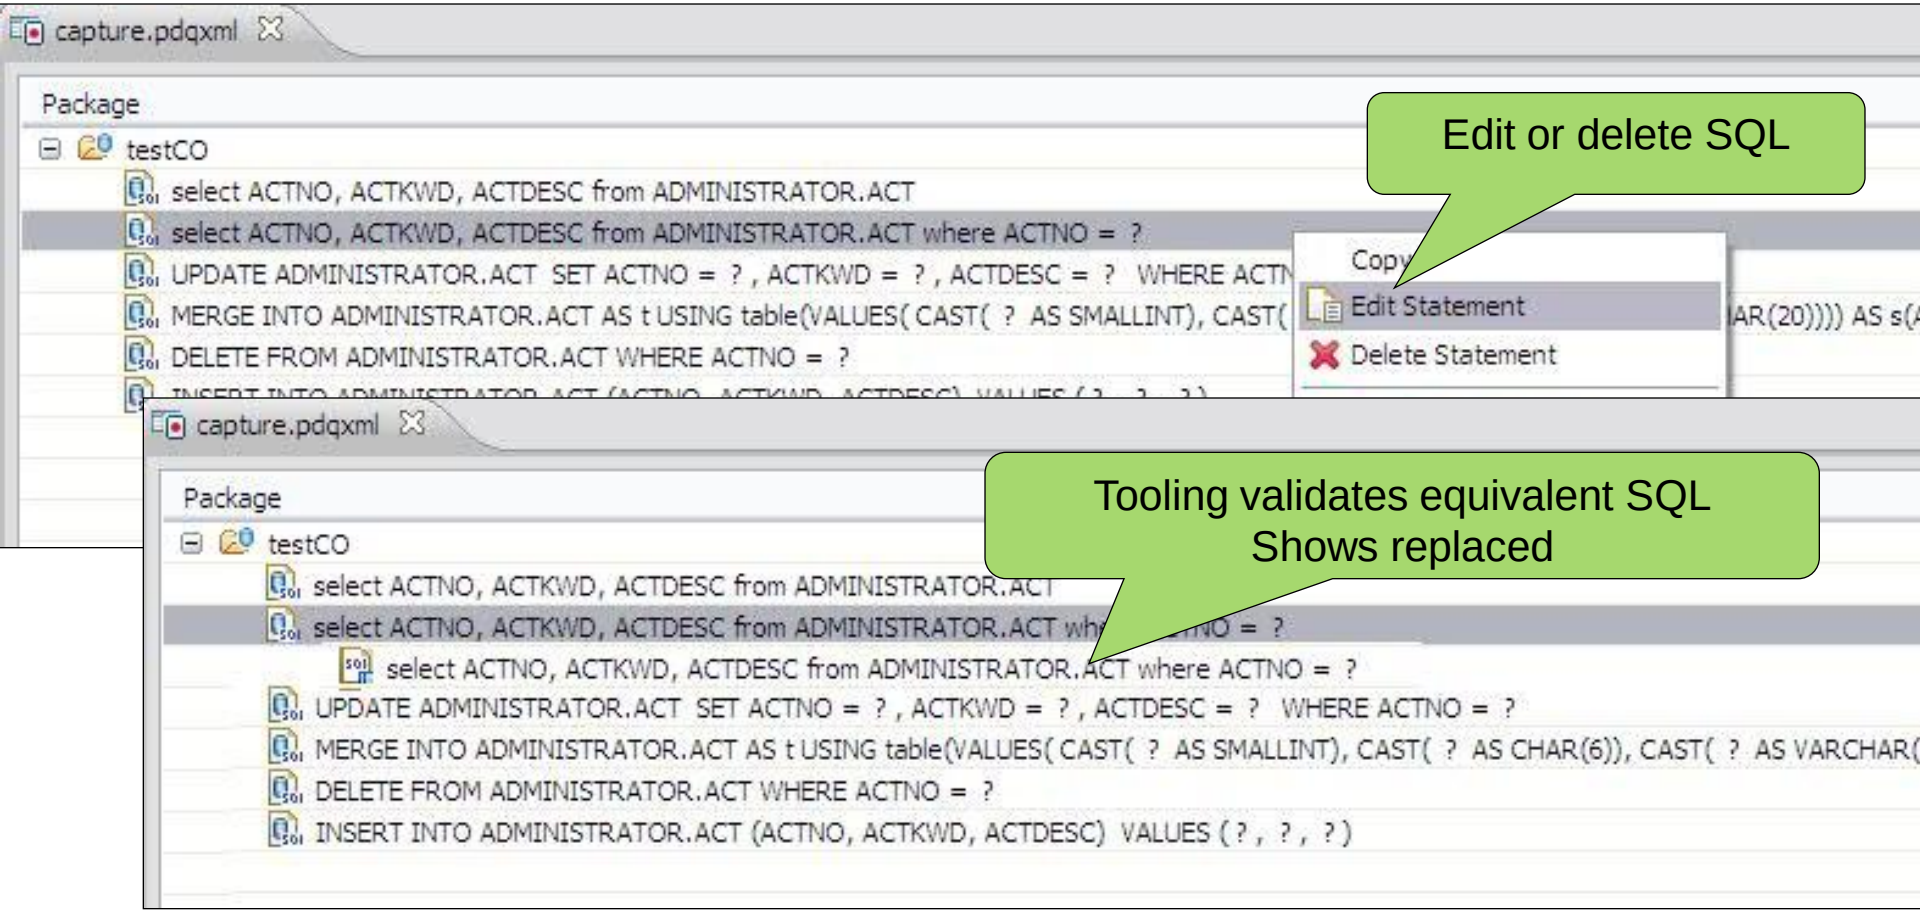

Eliminate SQL Injection

Capture → Review → Optimize → Revise → Restrict

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

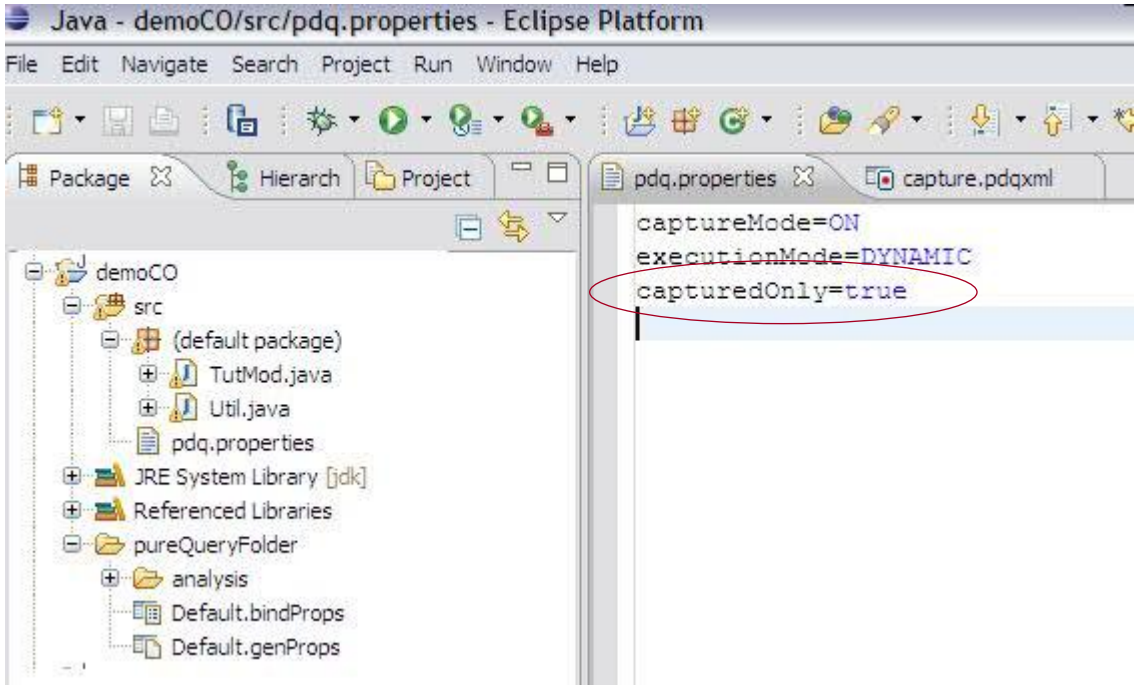

Optimize for WebSphere and DB2 with pureQuery

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

Reduce HW and SW costs

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

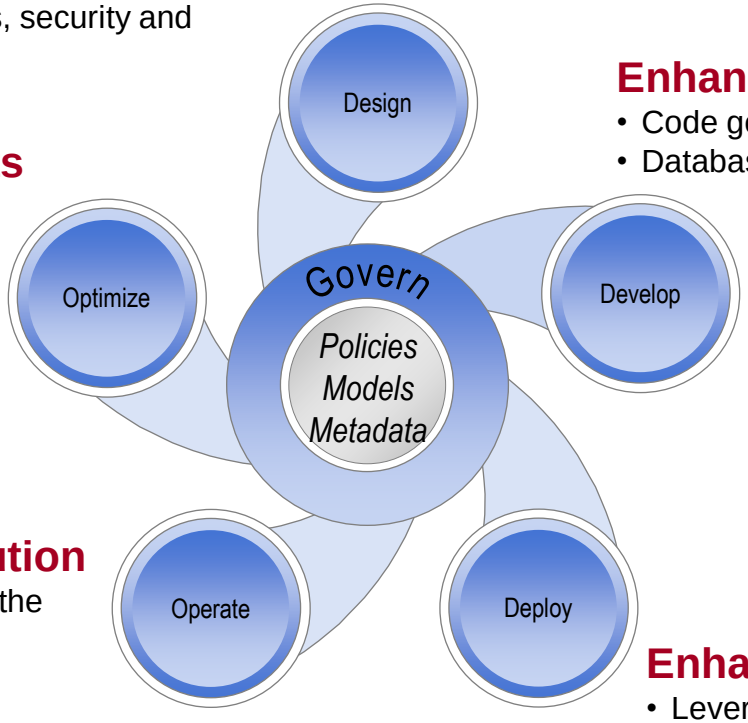

Enhance development productivity

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

Simplify impact analysis

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

Replace SQL without changing the source

- Editor validates equivalency

Focus tuning efforts

- Find and sort by query elapsed time from Java

Speed up problem resolution

- Trace SQL back to line of code in the application

Enhance performance

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

Prevent SQL Injection

- Lock down SQL for dynamic or static execution

Reduce security exposure

- Grant access to queries, not tables

Optim Development Studio and pureQuery Runtime

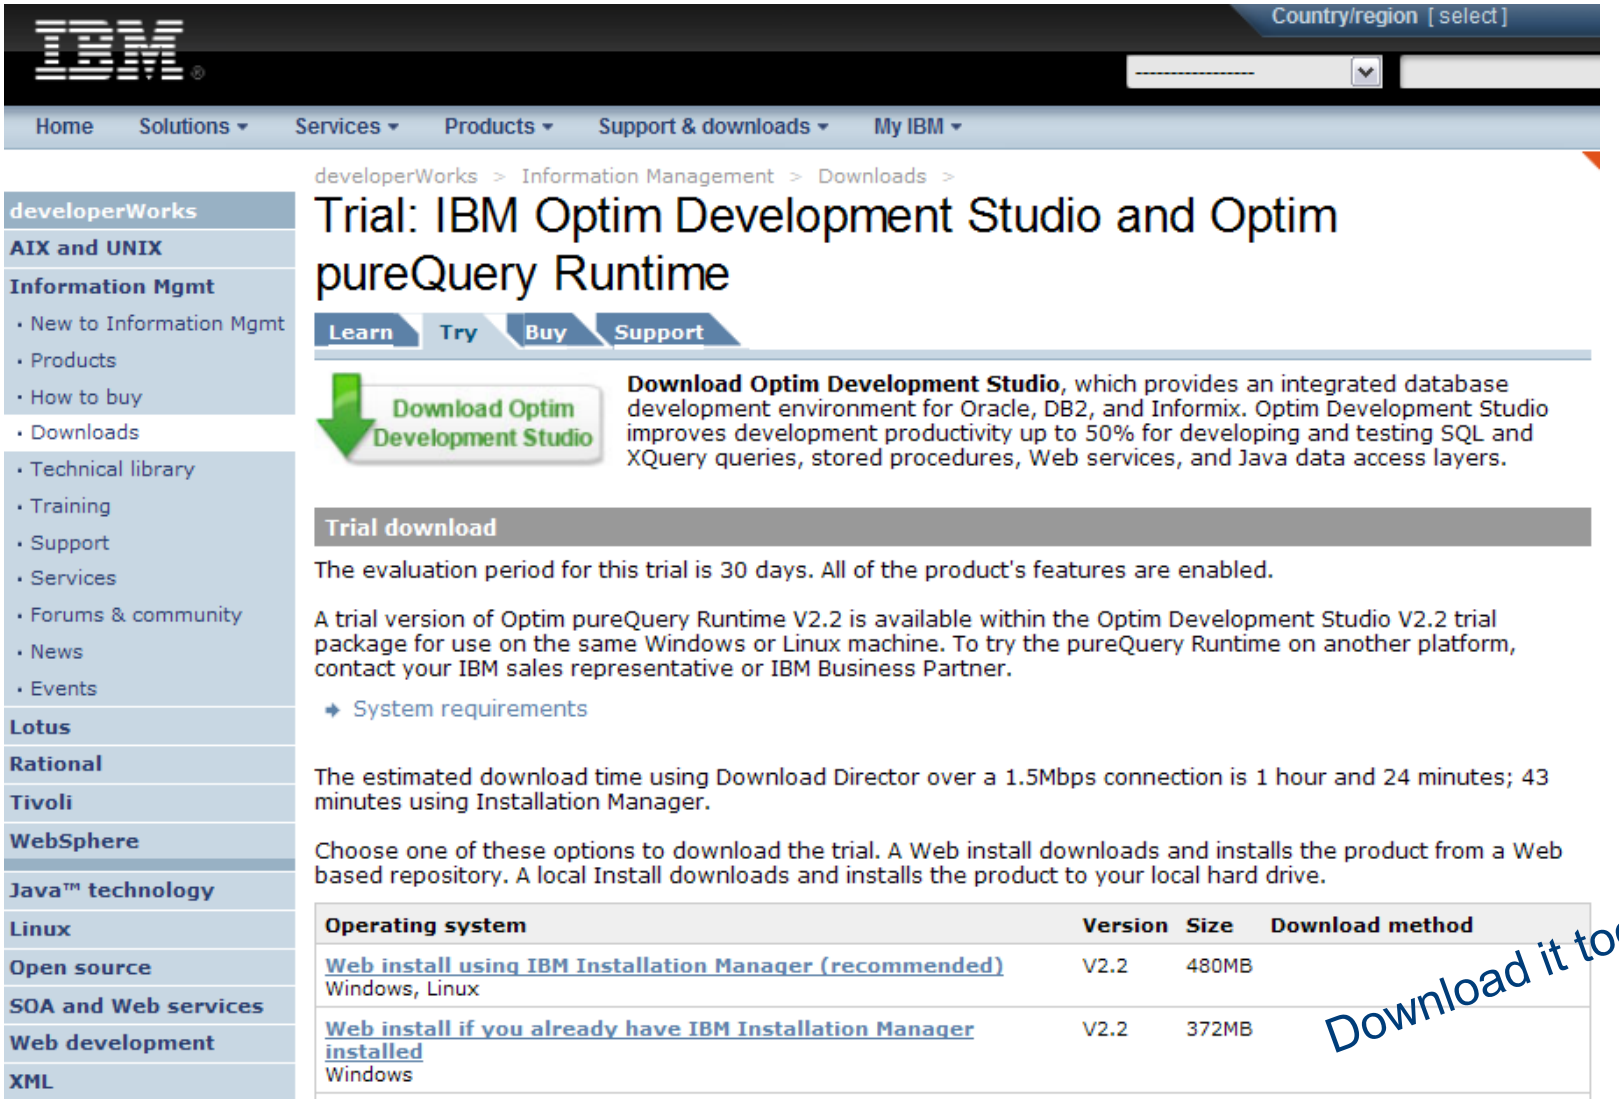

IBM

Country/region [select]

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

developerWorks > Information Management > Downloads >

developerWorks
AIX and UNIX
Information Mgmt
- New to Information Mgmt
- Products
- How to buy
- Downloads
- Technical library
- Training
- Support
- Services
- Forums & community
- News
- Events

Lotus
Rational
Tivoli
WebSphere
Java™ technology
Linux
Open source
SOA and Web services
Web development
XML

Trial: IBM Optim Development Studio and Optim pureQuery Runtime

Learn | Try | Buy | Support

Download Optim Development Studio

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
Web install using IBM Installation Manager (recommended) Windows, Linux	V2.2	480MB	
Web install if you already have IBM Installation Manager installed Windows	V2.2	372MB	

Download it today!

Data movement options

UNLOAD Utility

- **Designed for loading back into DB2 for z/OS**

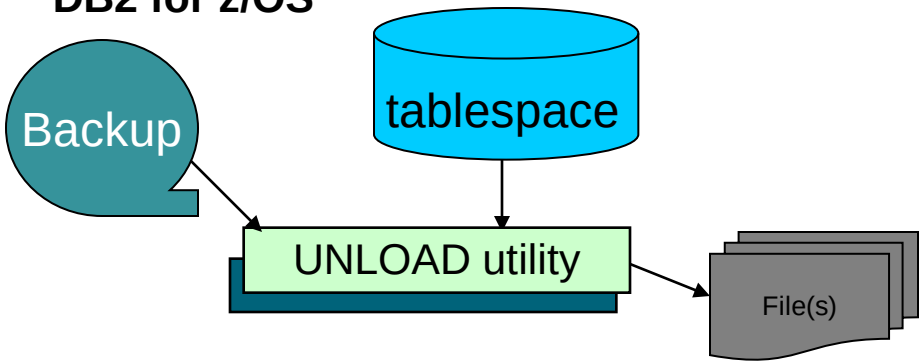

Crossloader (LOAD from INCURSOR DD)

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

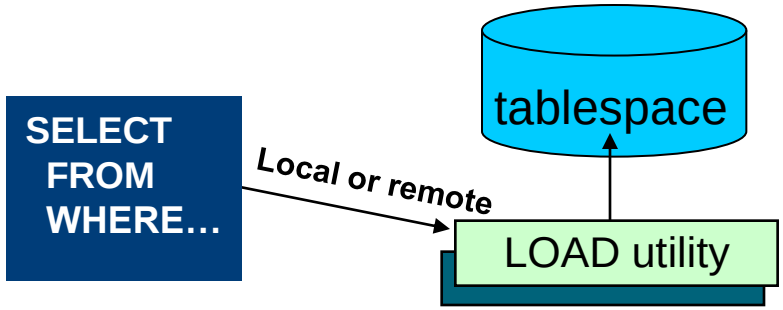

LOAD/UNLOAD via pipes

- Perfect for remote loading/unloading

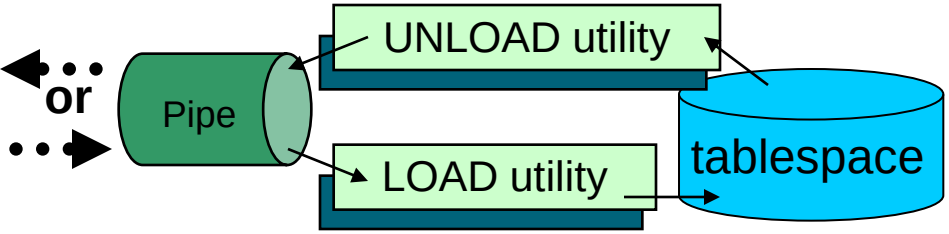

High Performance Unload (z/OS and LUW)

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

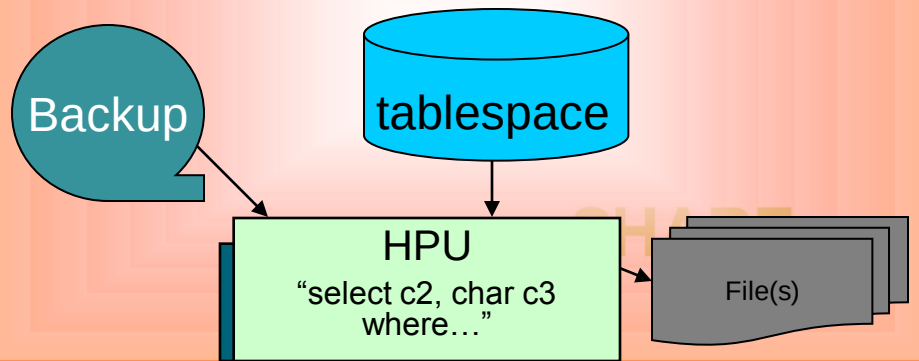

Recovery Tooling

- 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

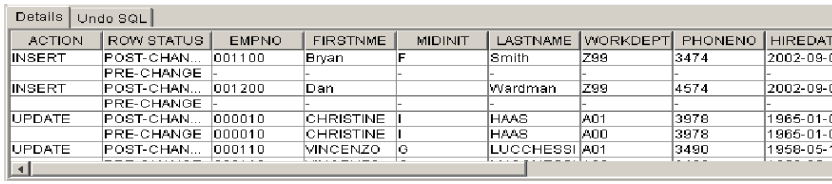

ACTION	ROW STATUS	EMPNO	FIRSTNME	MIDINIT	LASTNAME	WORKDEPT	PHONENO	HIREDAT
INSERT	POST-CHAN...	001100	Bryan	F	Smith	Z99	3474	2002-09-0
	PRE-CHANGE	-	-	-	-	-	-	-
INSERT	POST-CHAN...	001200	Dan		Wardman	Z99	4574	2002-09-0
	PRE-CHANGE	-	-	-	-	-	-	-
UPDATE	POST-CHAN...	000010	CHRISTINE	I	HAAS	A01	3978	1965-01-0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

- 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}

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

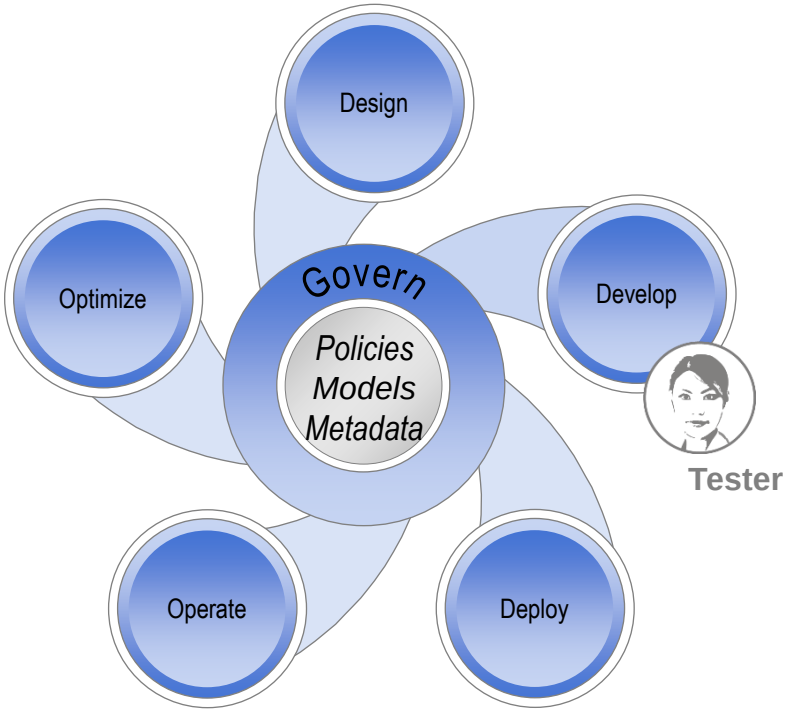

- Increase software quality
  - Use test data that accurately reflects production systems
- Accelerate release schedules
  - Comparison to expected results and fast refresh
- Comply with privacy regulations
  - De-identify personal information with realistic substitutions
- Protect confidential information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

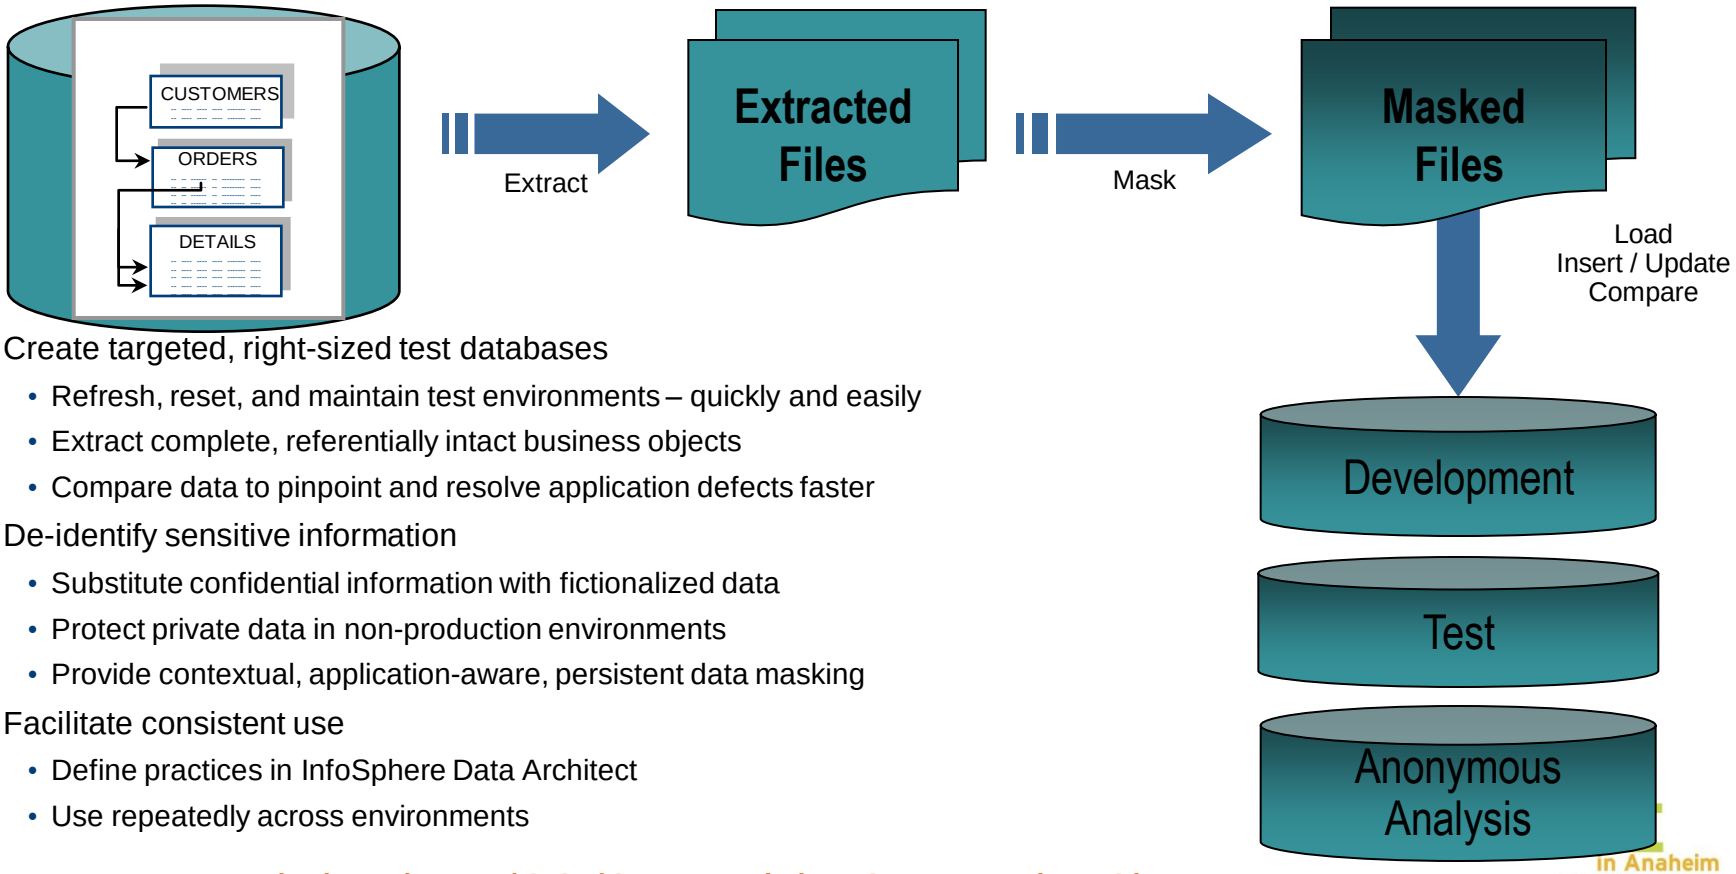

- Create targeted, right-sized test databases
  - Refresh, reset, and maintain test environments – quickly and easily
  - Extract complete, referentially intact business objects
  - Compare data to pinpoint and resolve application defects faster
- De-identify sensitive information
  - Substitute confidential 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

Supports non-relational on z/OS (Sequential, VSAM, and IMS)!

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

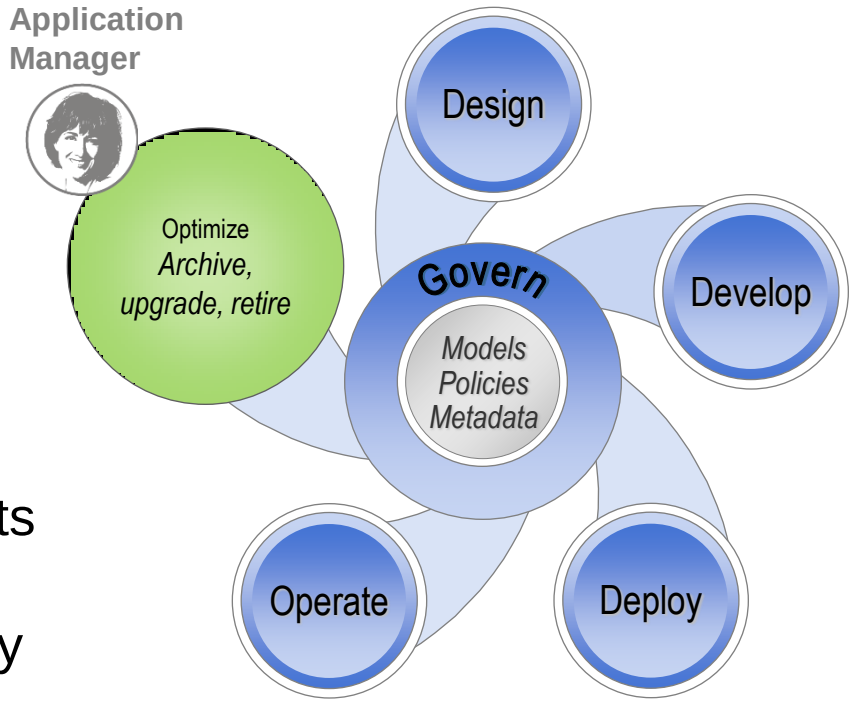

Optim Data Growth Solution

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

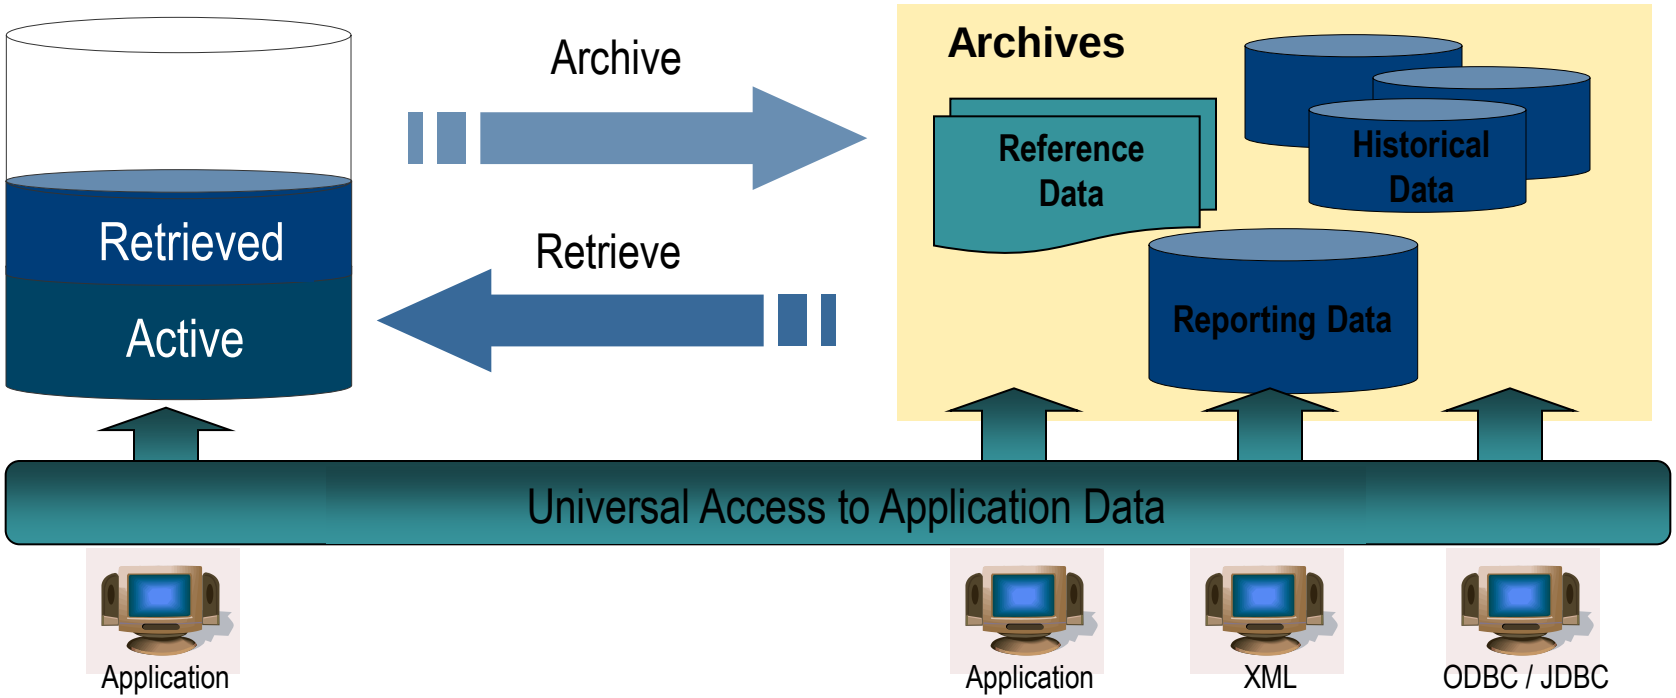

- **Complete business object provides historical reference snapshot of business activity**
- **Storage device independence enables storage cost optimization**
- **Immutable file format enables data retention compliance**

What's cooking back at the lab?

- Performance Management

Where is my DB2 application spending its time?

- **OMEGAMON PE's Extended Insight** is an advanced way to monitor the database workload (SQL) of your applications and solutions
  - Get total response times and response time breakdown (appl, driver, network, data server) per defined workload/cluster (e.g. per system, application, user)
  - Compare workload from various servers / applications
  - Select a time period for analysis
  - Get top SQL statements per defined workload
  - Identify top clients contributing in the workload

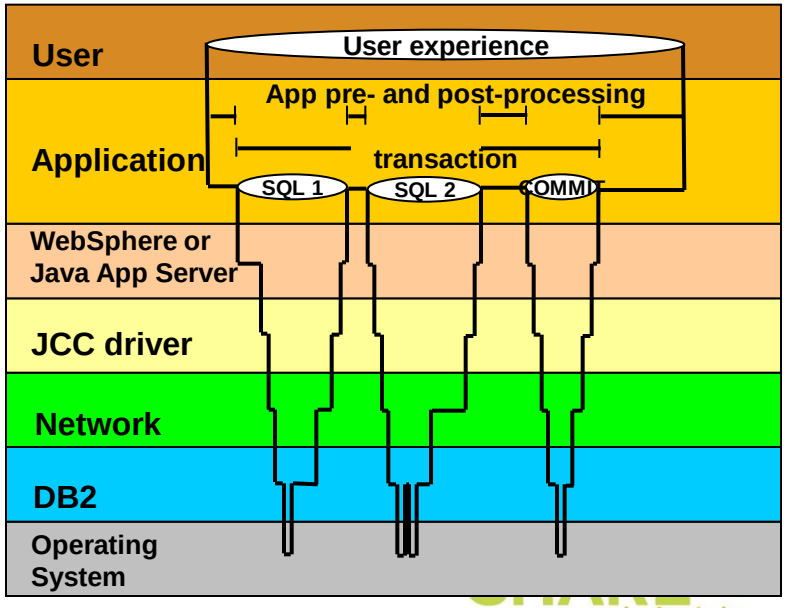

How to start and navigate to the Extended Insight dashboard

- Seamless navigation depending on the usage/problem scenario …

  1. Integration and navigation to Extended Analysis Dashboard as part of new OMEGAMON XE for DB2 PE on z/OS (OMPE) TEP workspace

  2. As a stand-alone web console session

  3. Integrated with ITCAM and navigation to Extended Analysis Dashboard

OMPE on z/OS TEP
Navigation to the Extended Insight Analysis Dashboard

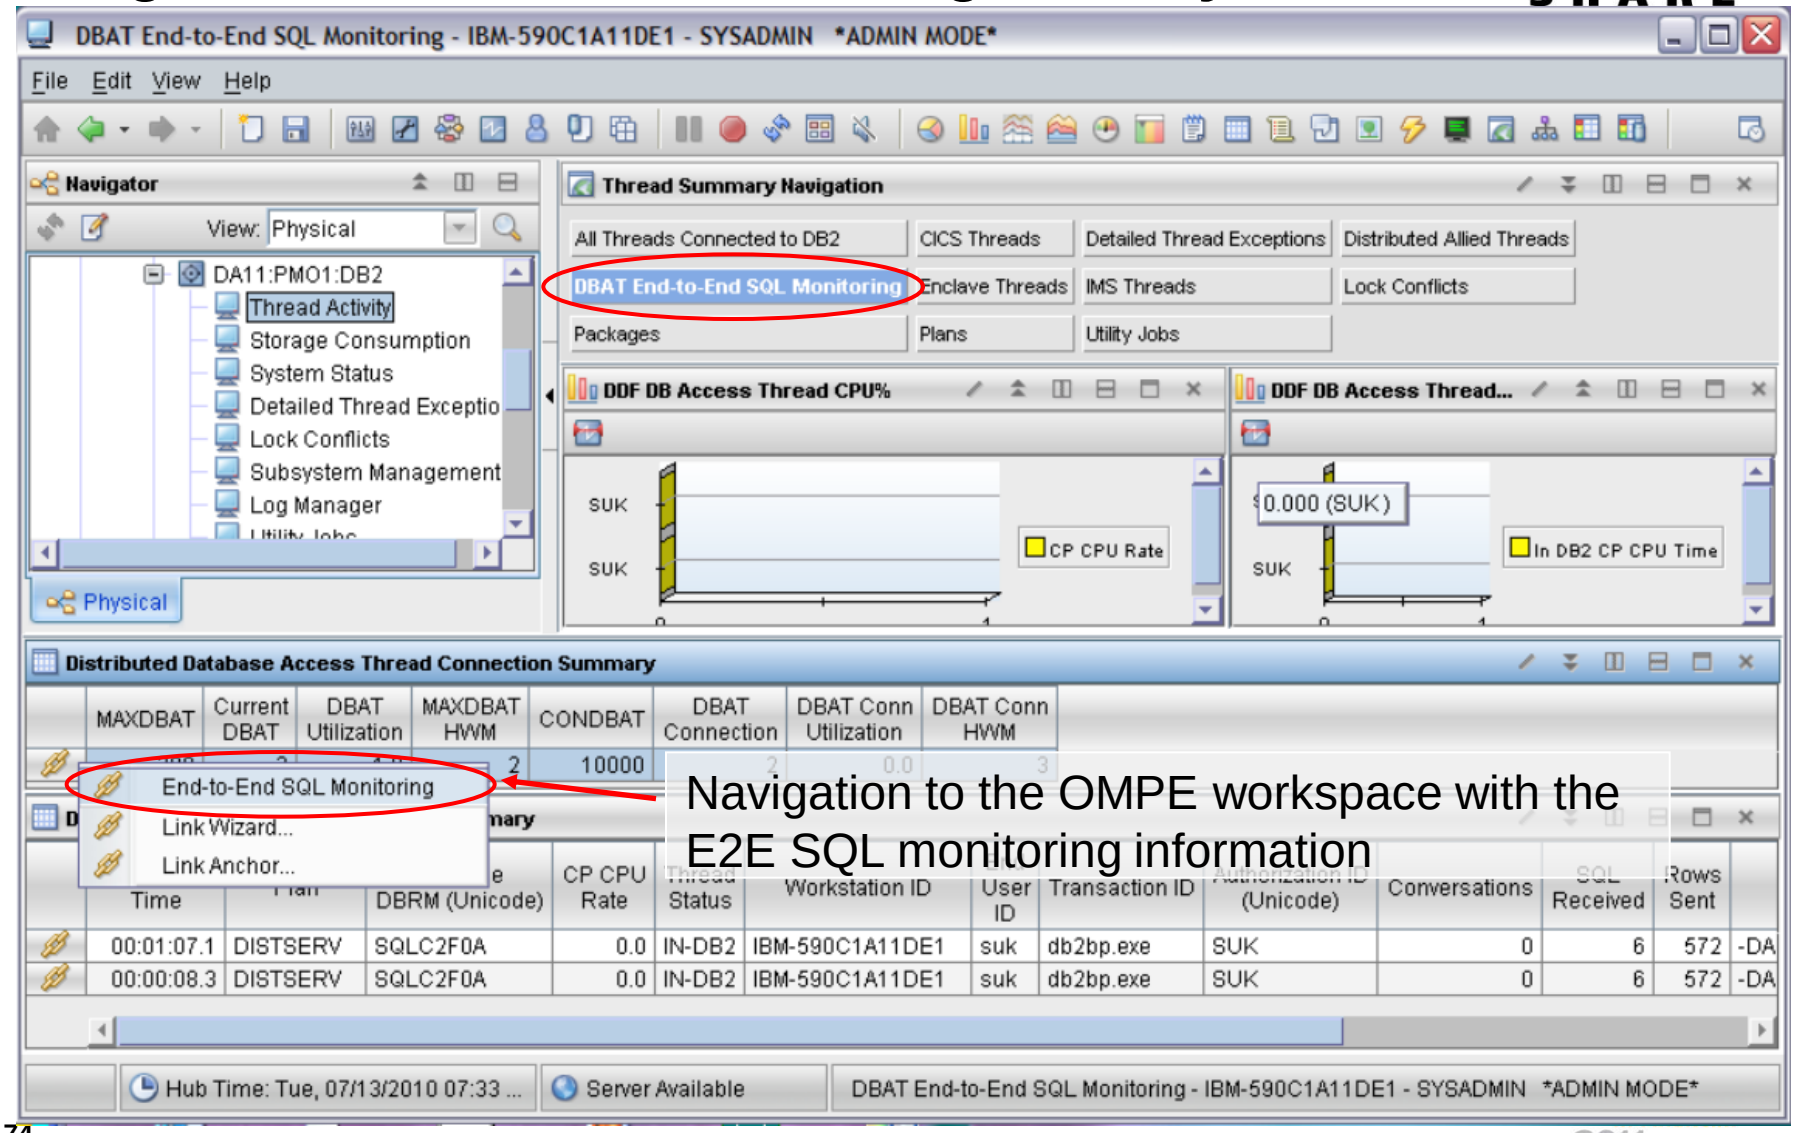

Navigation to the OMPE workspace with the E2E SQL monitoring information

OMPE on z/OS TEP
Navigation to the Extended Insight Analysis Dashboard

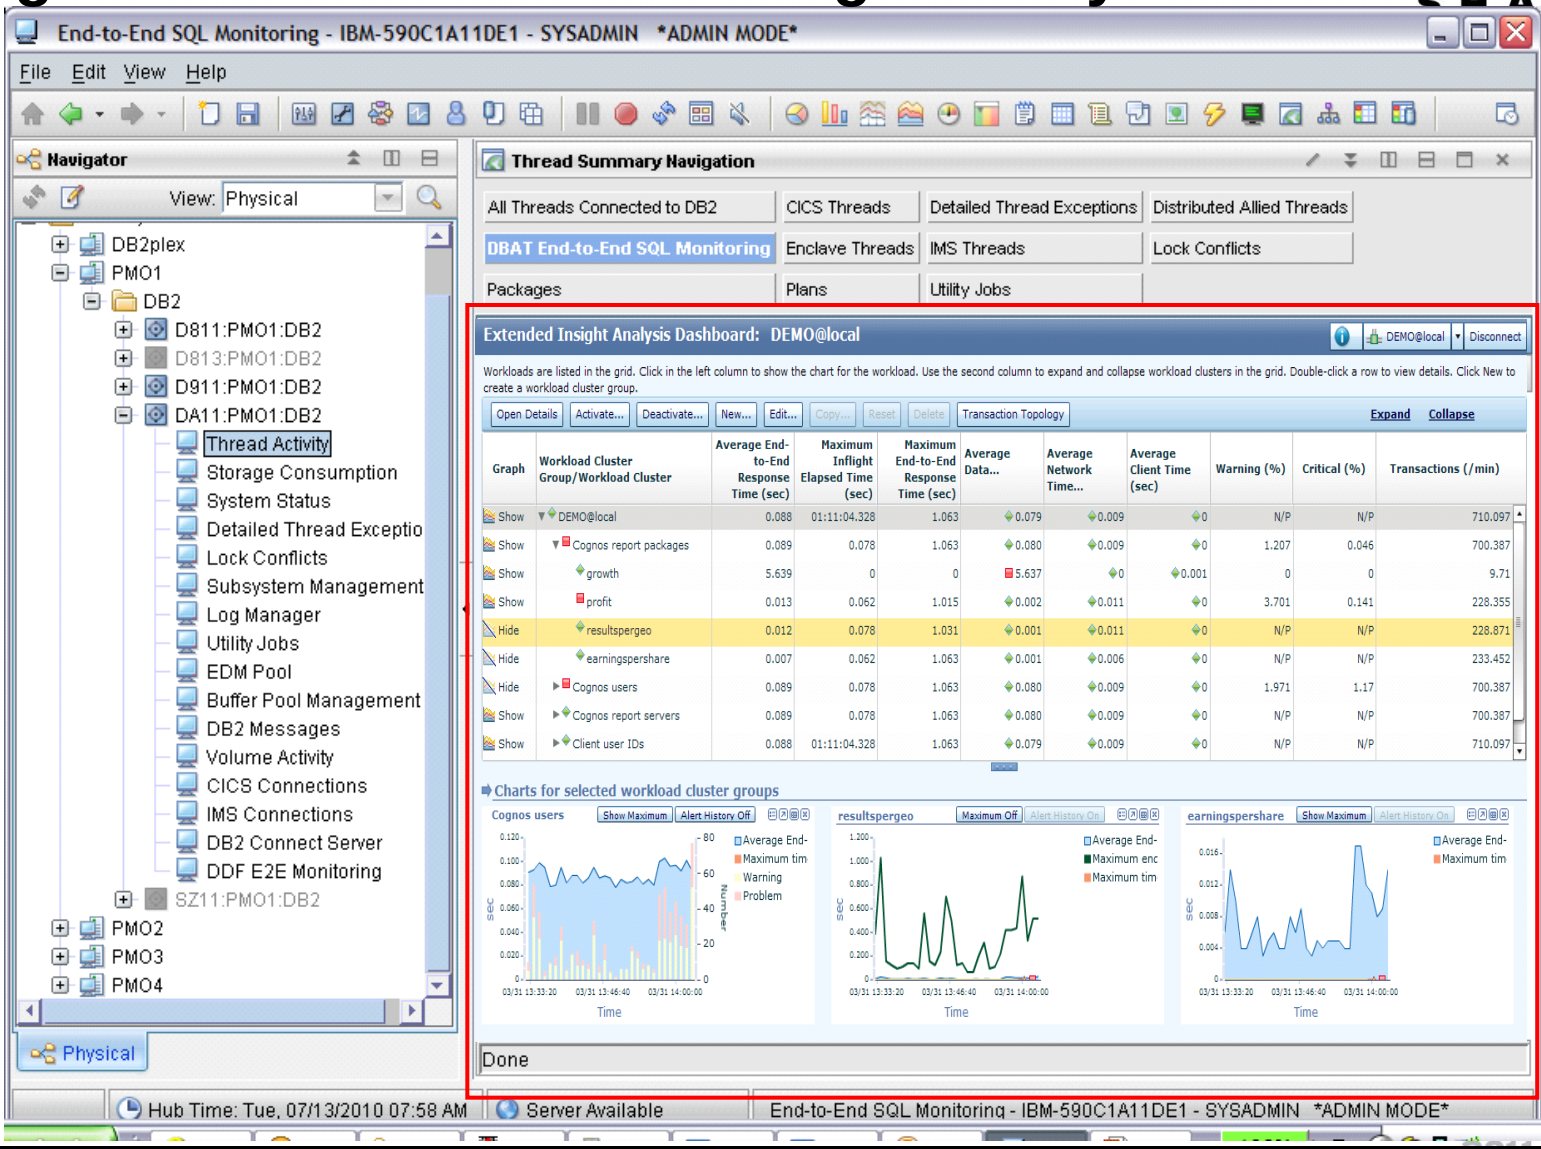

Optim Performance Manager Web Console with Health Summary and navigation to Extended Insight Dashboard

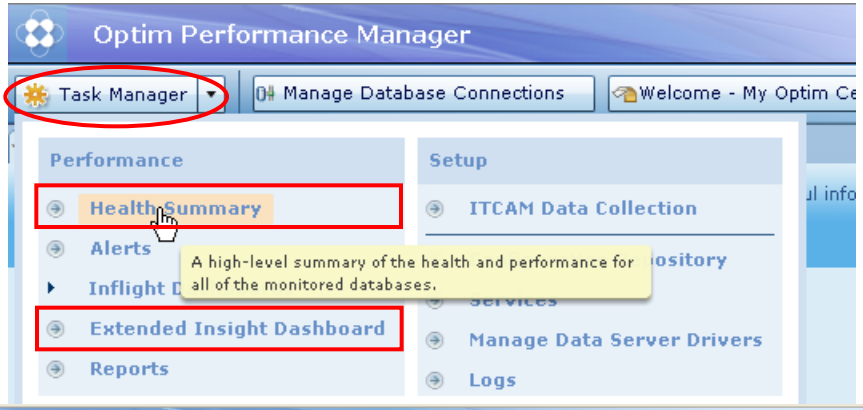

OMEGAMON XE for DB2 PE on z/OS reuses the Optim Performance Manager Extended Insight monitoring capability.. Health Summary and Extended Insight Dashboard are enabled for monitoring of SQL access DB2 for z/OS.

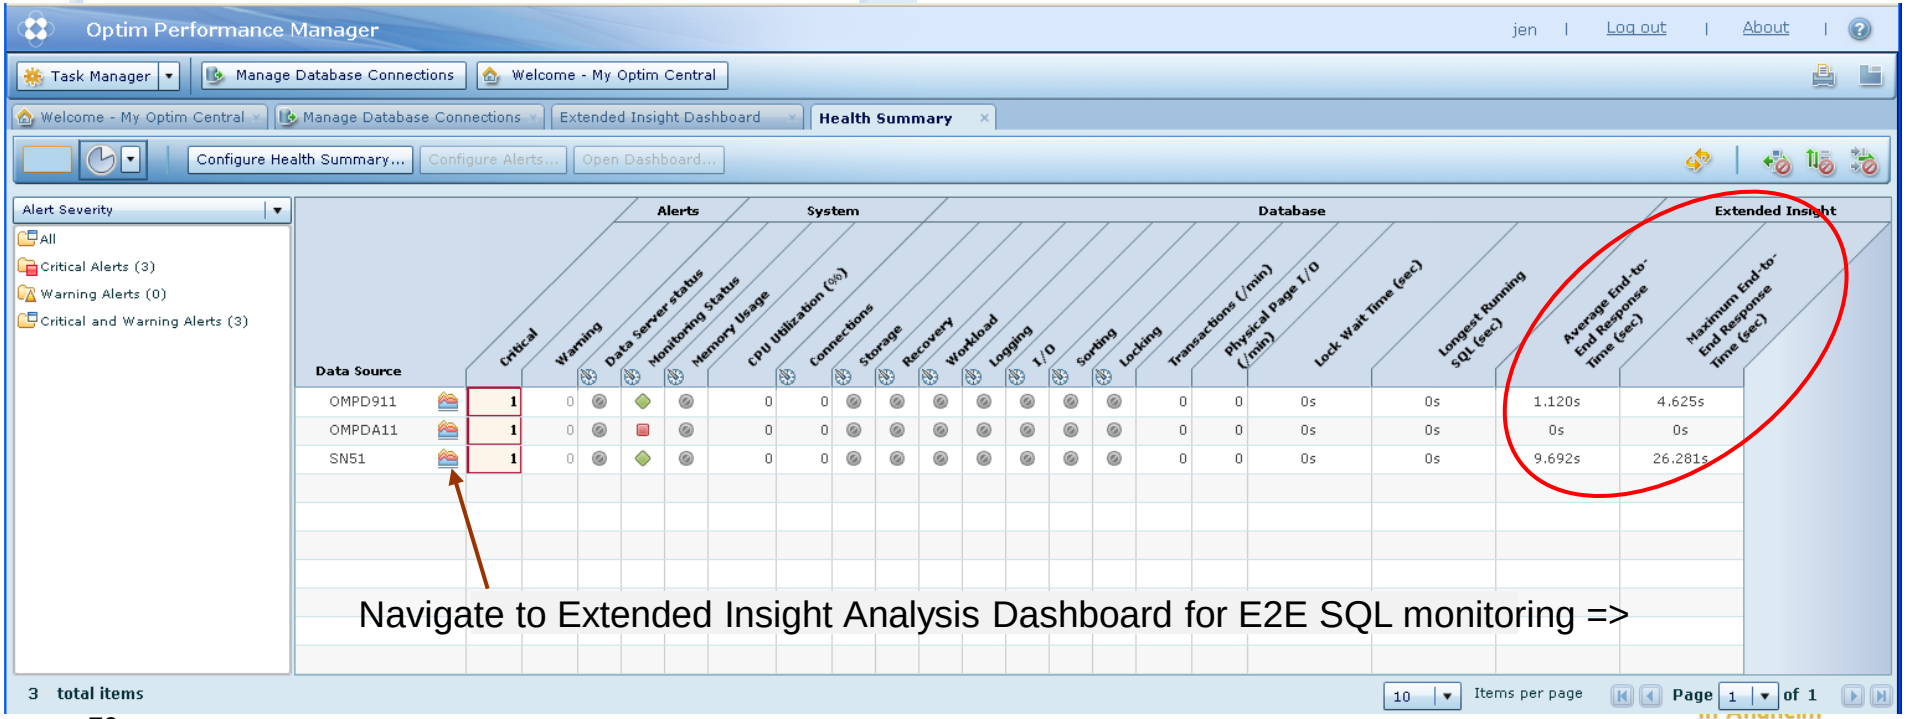

Extended Insight Analysis Dashboard

The slider bar allows selection of the time period to be considered

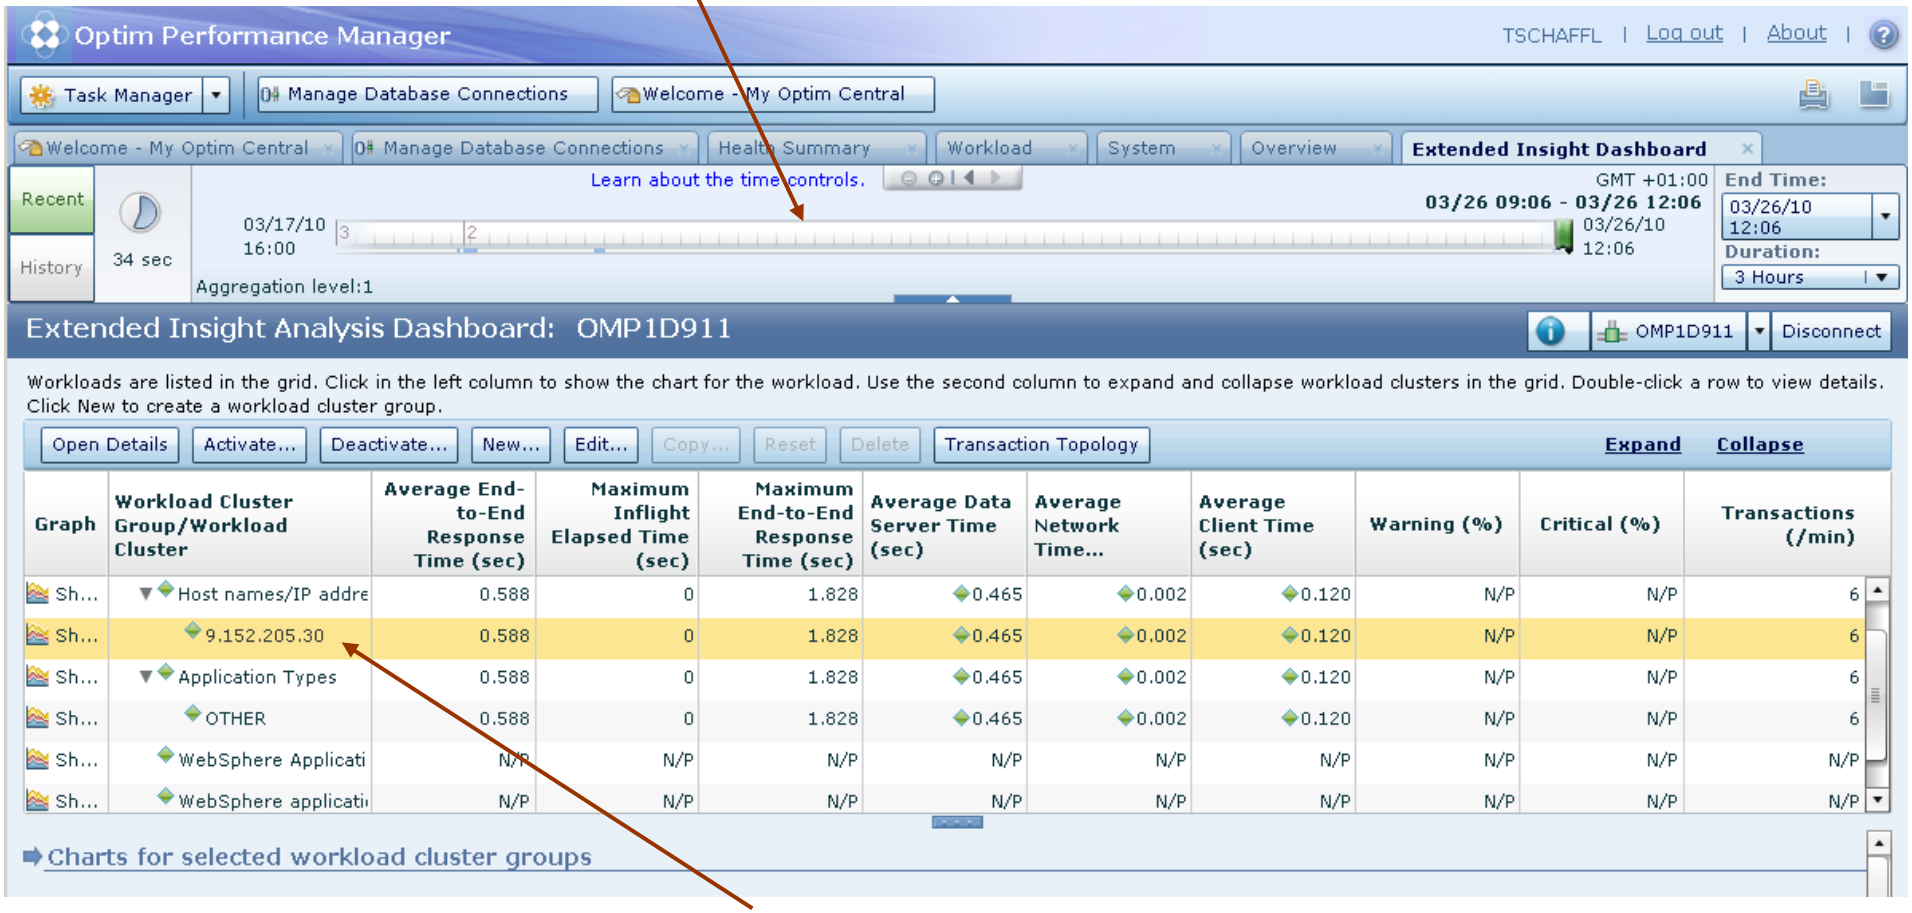

Overview and comparison of “Workload cluster groups” but also on details with the capability to select and further zoom in.

Extended Insight Analysis Dashboard

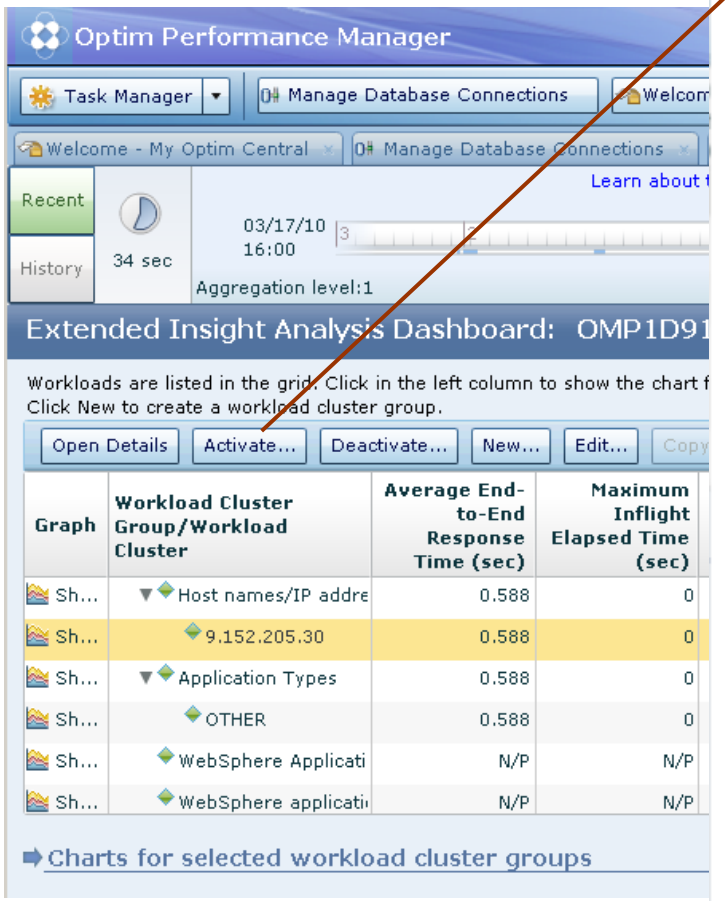

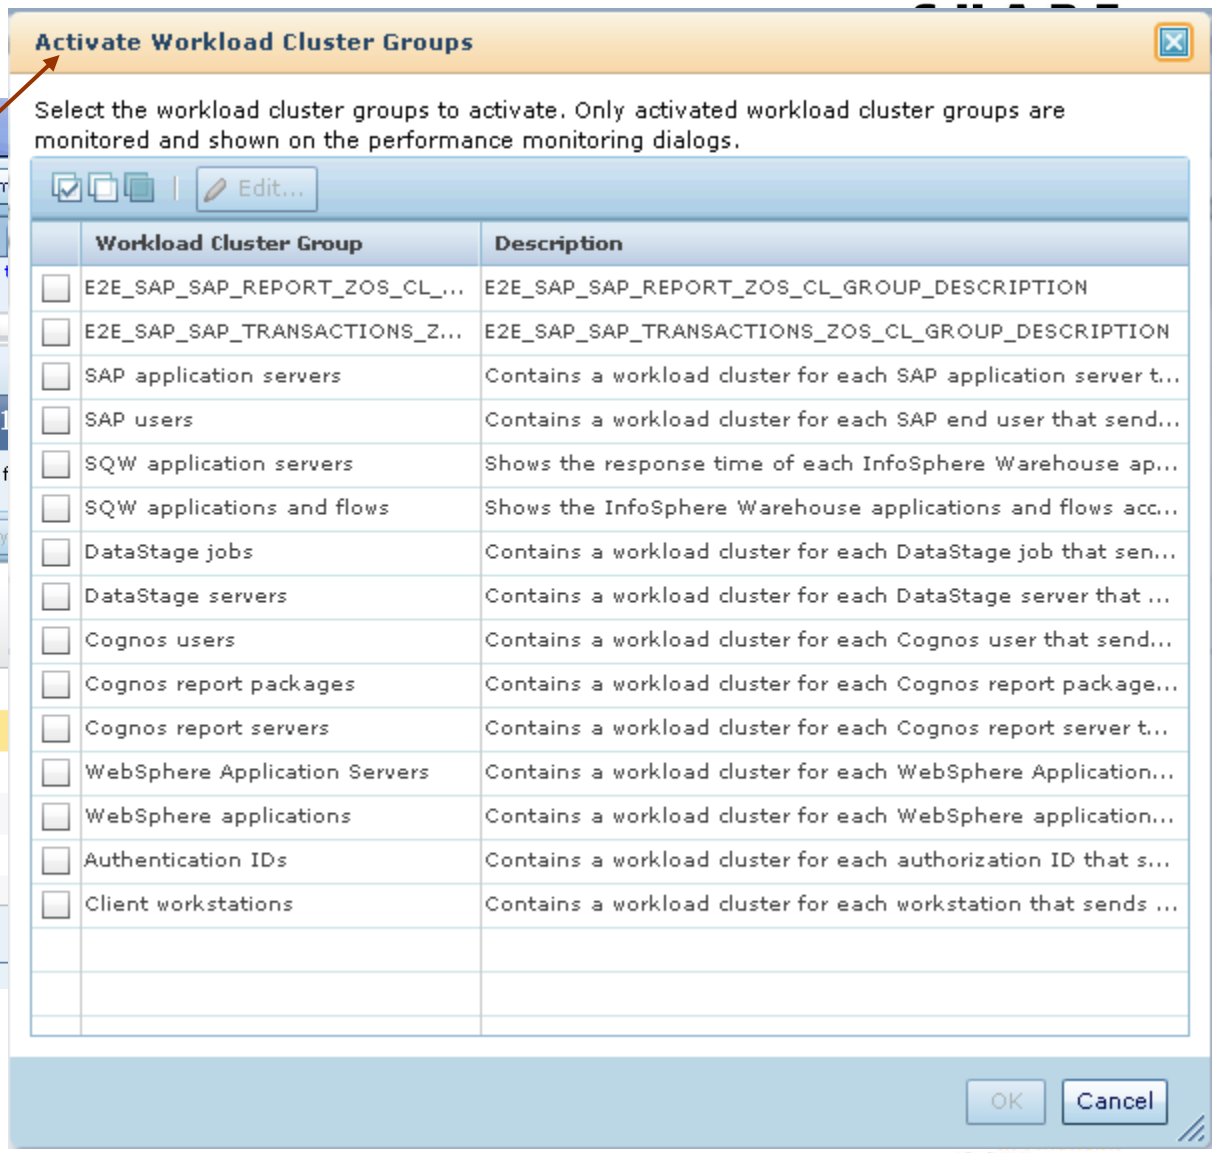

Define your “Workload cluster groups” as you need

Extended Insight Analysis Dashboard

Expand to more details, e.g. expand user data and application data

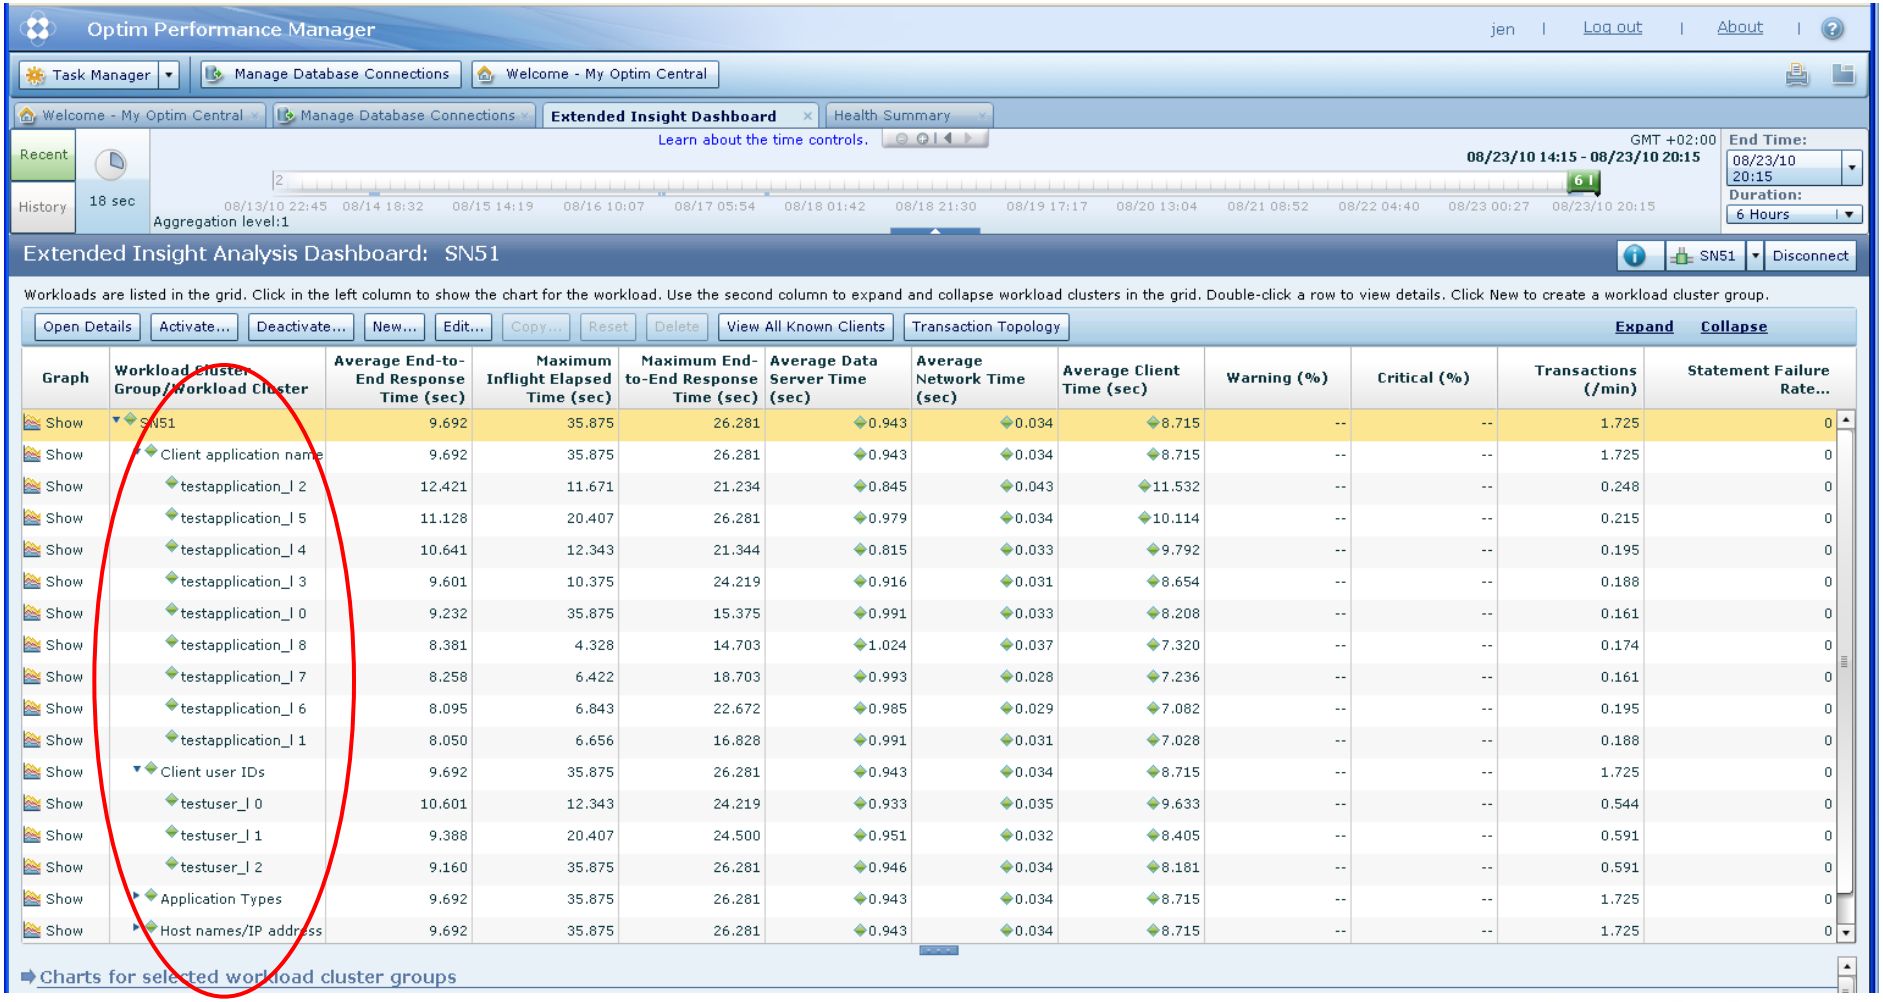

Graph	Workload Cluster Group/Workload Cluster	Average End-to-End Response Time (sec)	Maximum Inflight Elapsed Time (sec)	Maximum End-to-End Response Time (sec)	Average Data Server Time (sec)	Average Network Time (sec)	Average Client Time (sec)	Warning (%)	Critical (%)	Transactions (/min)	Statement Failure Rate...
Show	SN51	9.692	35.875	26.281	0.943	0.034	8.715	--	--	1.725	0
Show	Client application name	9.692	35.875	26.281	0.943	0.034	8.715	--	--	1.725	0
Show	testapplication_I 2	12.421	11.671	21.234	0.845	0.043	11.532	--	--	0.248	0
Show	testapplication_I 5	11.128	20.407	26.281	0.979	0.034	10.114	--	--	0.215	0
Show	testapplication_I 4	10.641	12.343	21.344	0.815	0.033	9.792	--	--	0.195	0
Show	testapplication_I 3	9.601	10.375	24.219	0.916	0.031	8.654	--	--	0.188	0
Show	testapplication_I 0	9.232	35.875	15.375	0.991	0.033	8.208	--	--	0.161	0
Show	testapplication_I 8	8.381	4.328	14.703	1.024	0.037	7.320	--	--	0.174	0
Show	testapplication_I 7	8.258	6.422	18.703	0.993	0.028	7.236	--	--	0.161	0
Show	testapplication_I 6	8.095	6.843	22.672	0.985	0.029	7.082	--	--	0.195	0
Show	testapplication_I 1	8.050	6.656	16.828	0.991	0.031	7.028	--	--	0.188	0
Show	Client user IDs	9.692	35.875	26.281	0.943	0.034	8.715	--	--	1.725	0
Show	testuser_I 0	10.601	12.343	24.219	0.933	0.035	9.633	--	--	0.544	0
Show	testuser_I 1	9.388	20.407	24.500	0.951	0.032	8.405	--	--	0.591	0
Show	testuser_I 2	9.160	35.875	26.281	0.946	0.034	8.181	--	--	0.591	0
Show	Application Types	9.692	35.875	26.281	0.943	0.034	8.715	--	--	1.725	0
Show	Host names/IP address	9.692	35.875	26.281	0.943	0.034	8.715	--	--	1.725	0

Extended Insight Analysis Dashboard

Show additional graphs for selected workload clusters

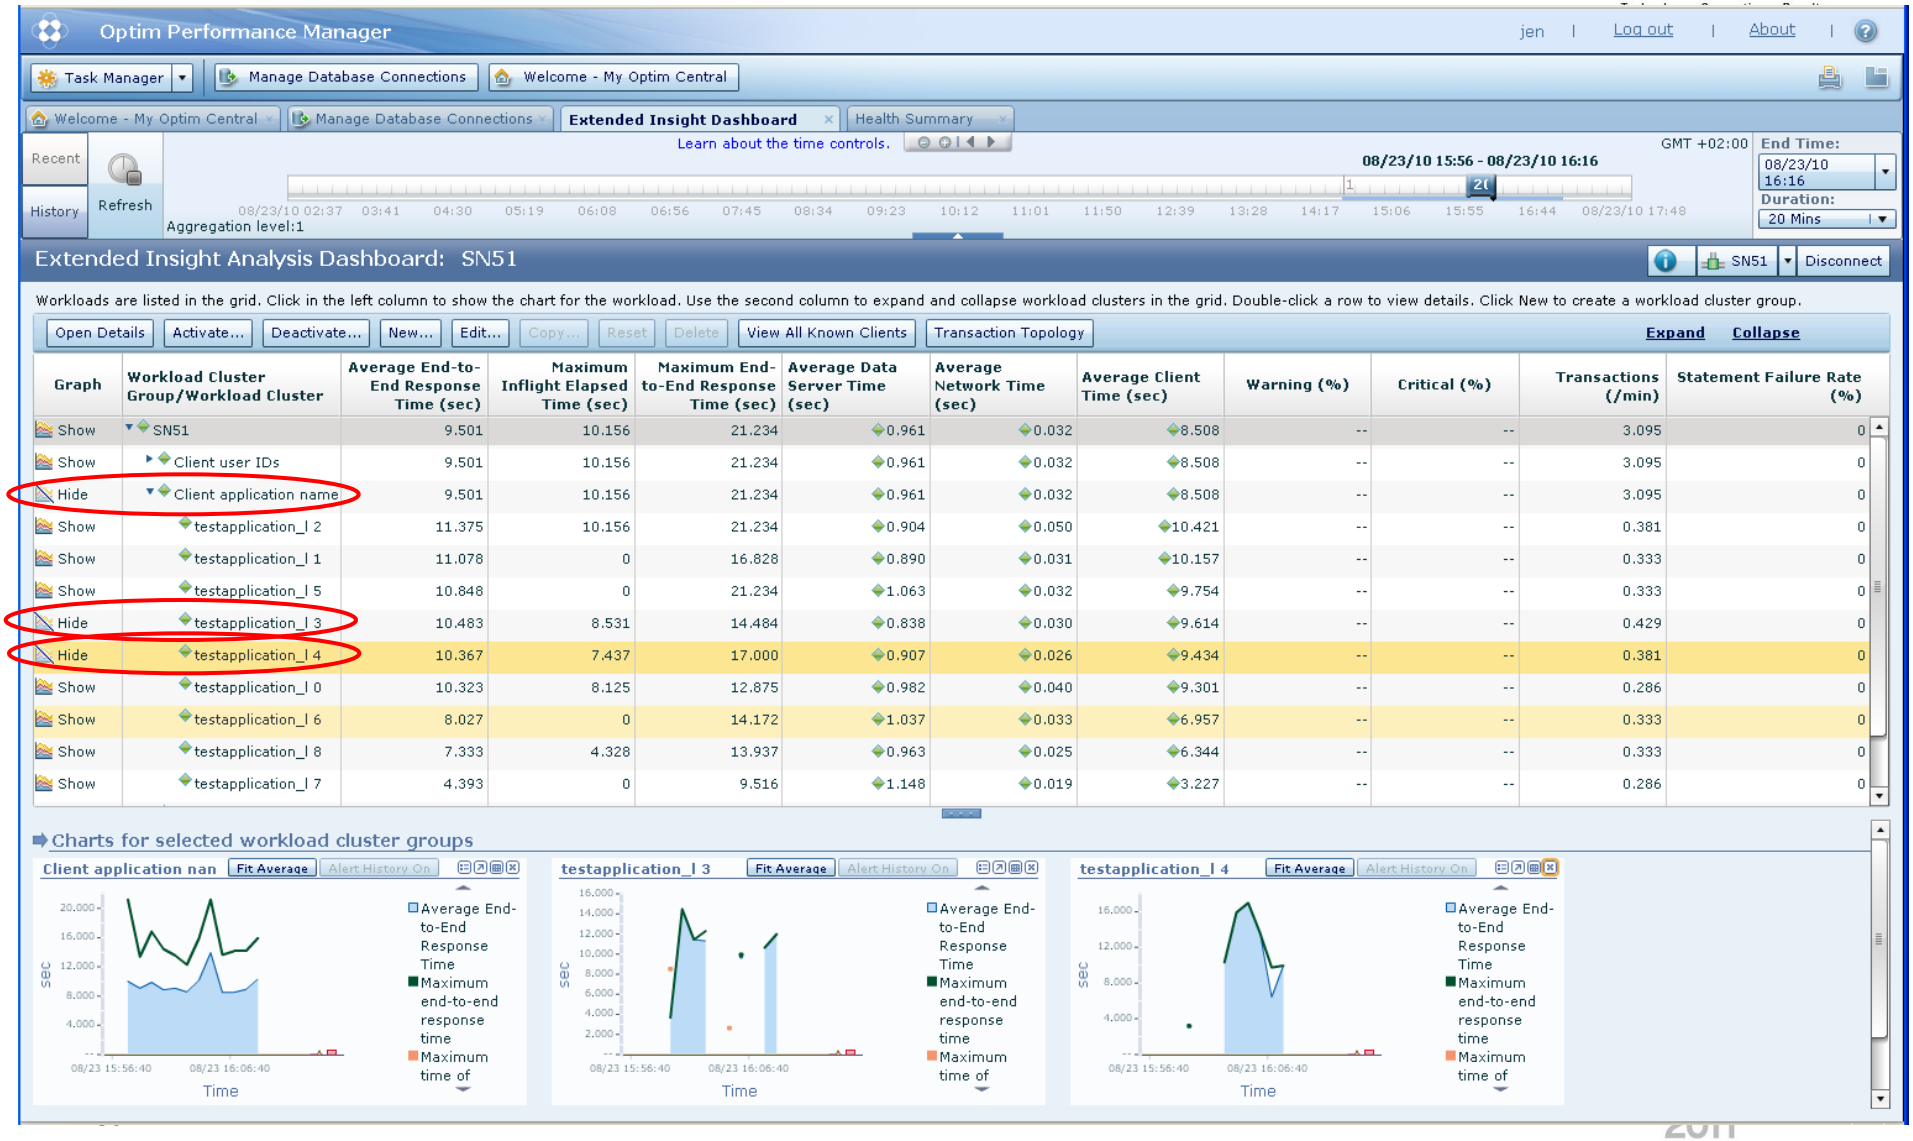

Extended Insight Analysis Dashboard

Show response time histogram for selected workload

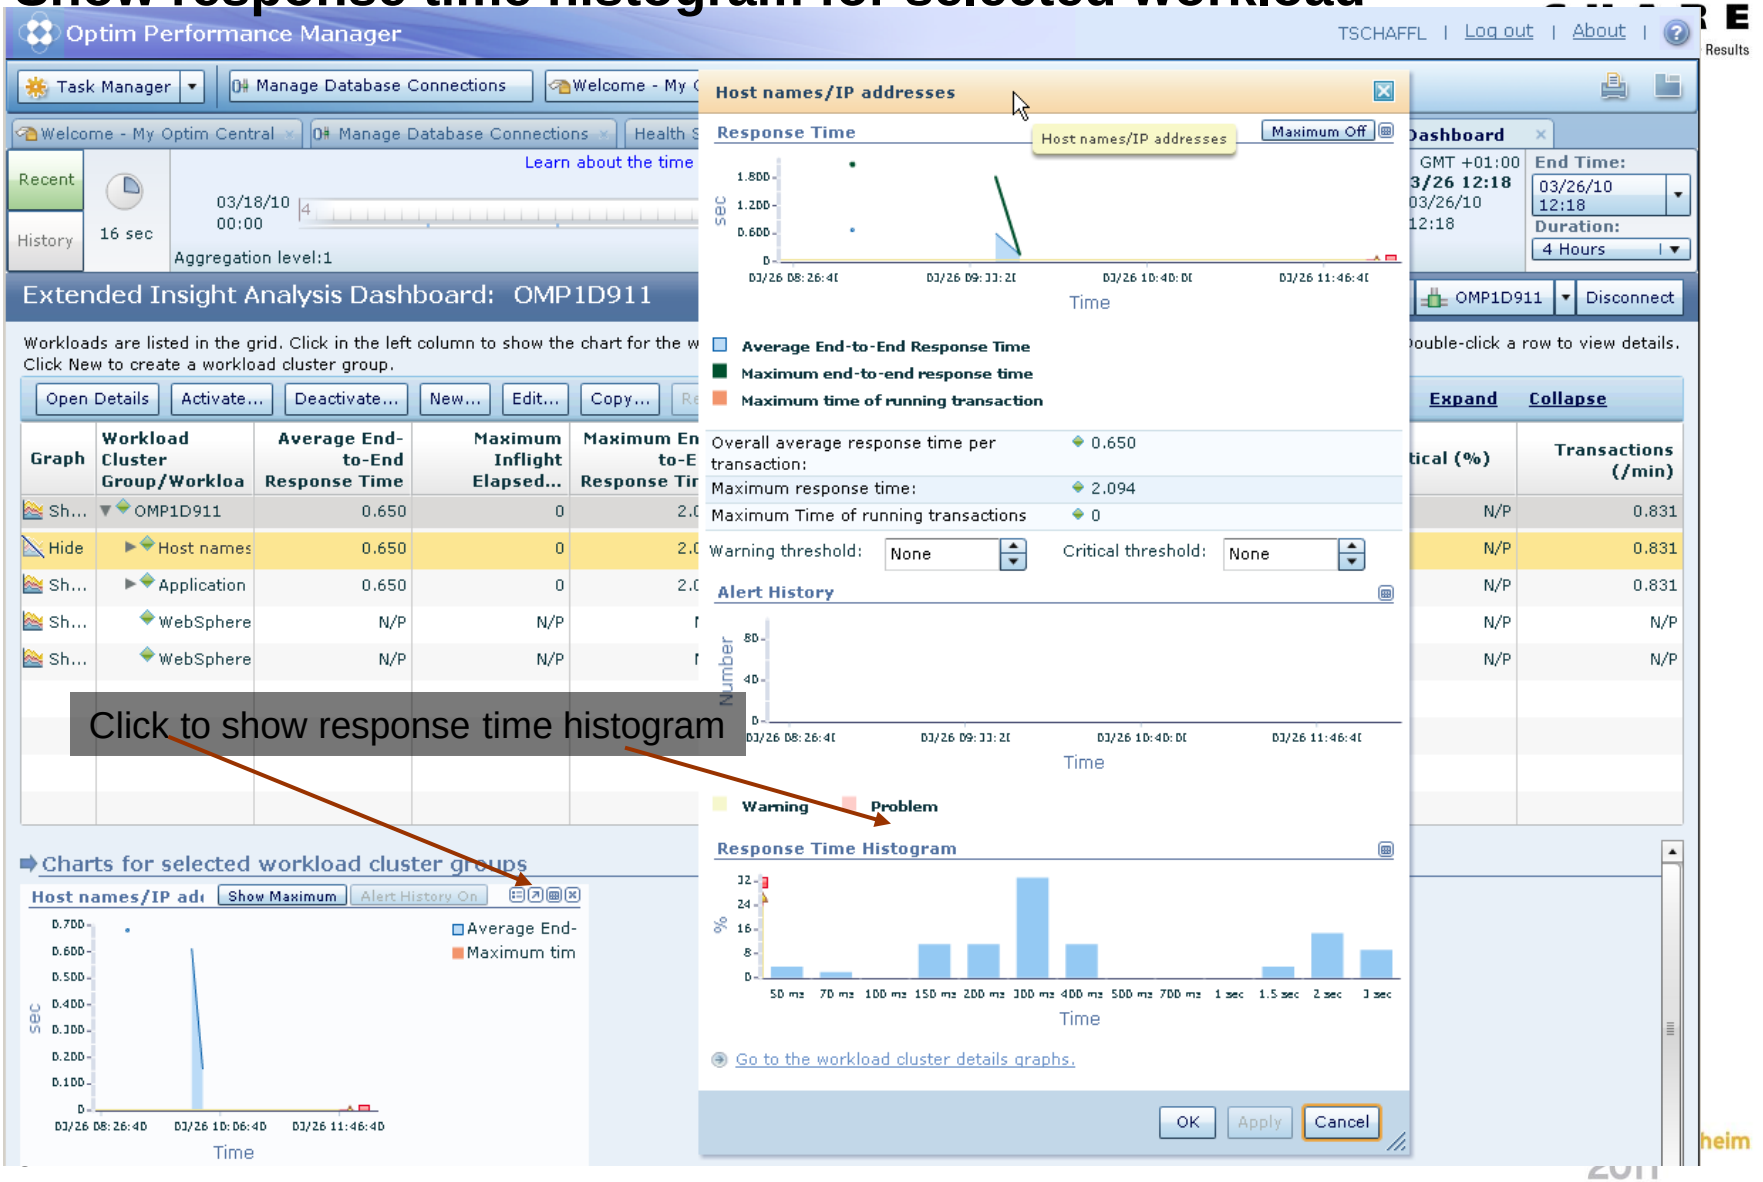

Extended Insight Analysis Dashboard

Zoom into selected workload and see the TOP SQL list

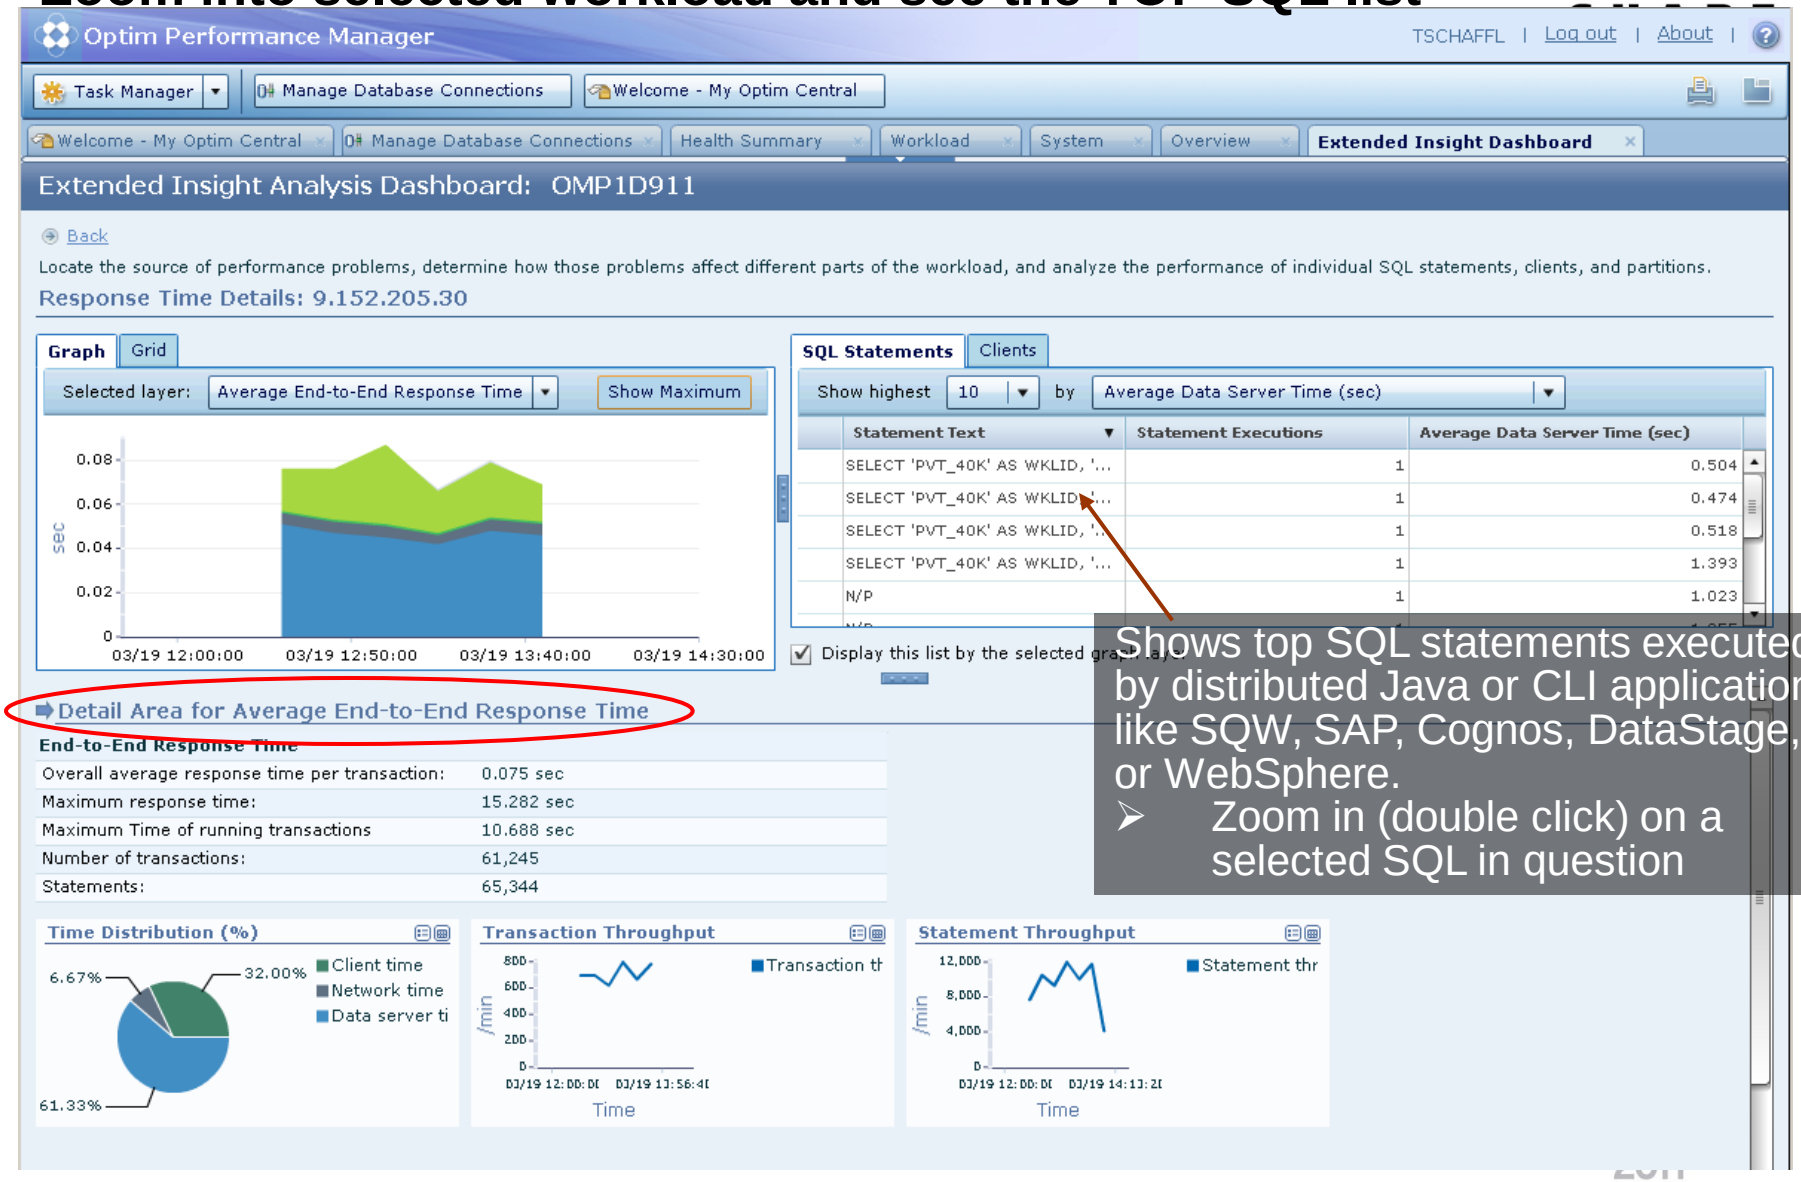

Shows top SQL statements executed by distributed Java or CLI applications like SQW, SAP, Cognos, DataStage, or WebSphere.

- Zoom in (double click) on a selected SQL in question

Extended Insight Analysis Dashboard
Select SQL from list and zoom into SQL level details

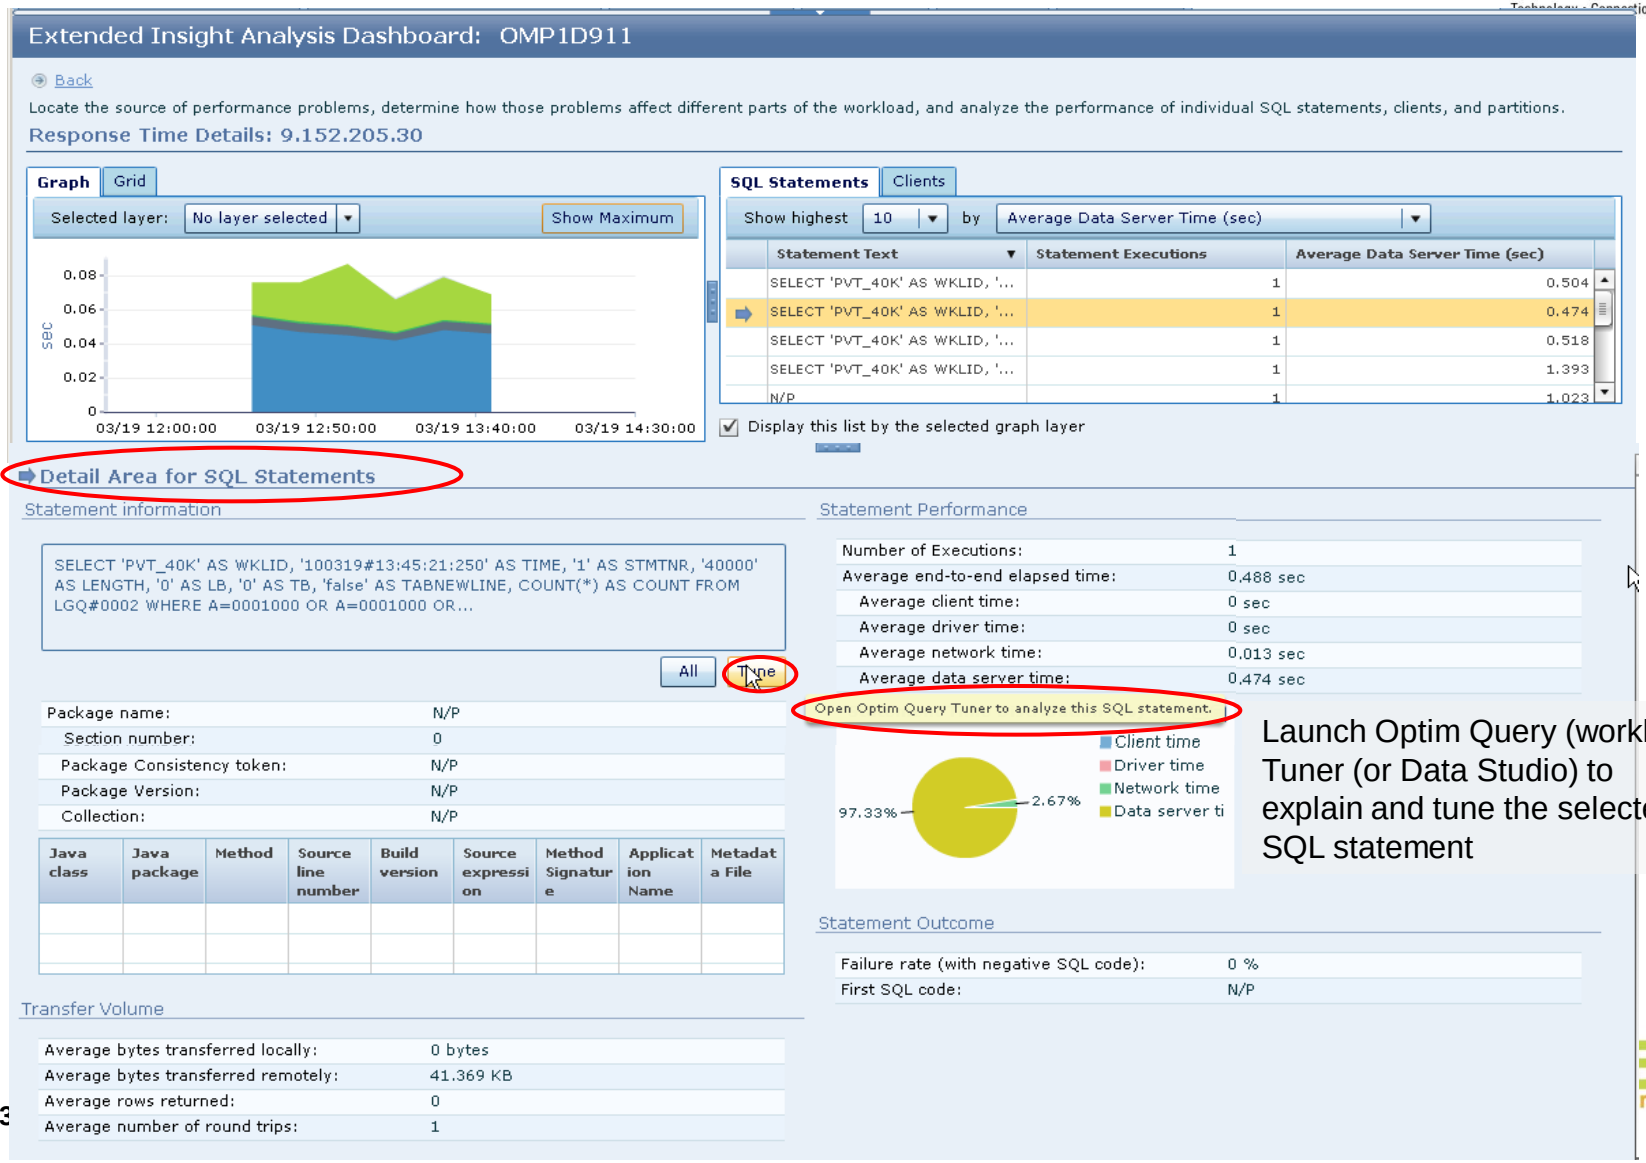

Launch Optim Query (workload) Tuner (or Data Studio) to explain and tune the selected SQL statement

Extended Insight Analysis Dashboard

Page down to review the host Dynamic SQL statement cache metrics

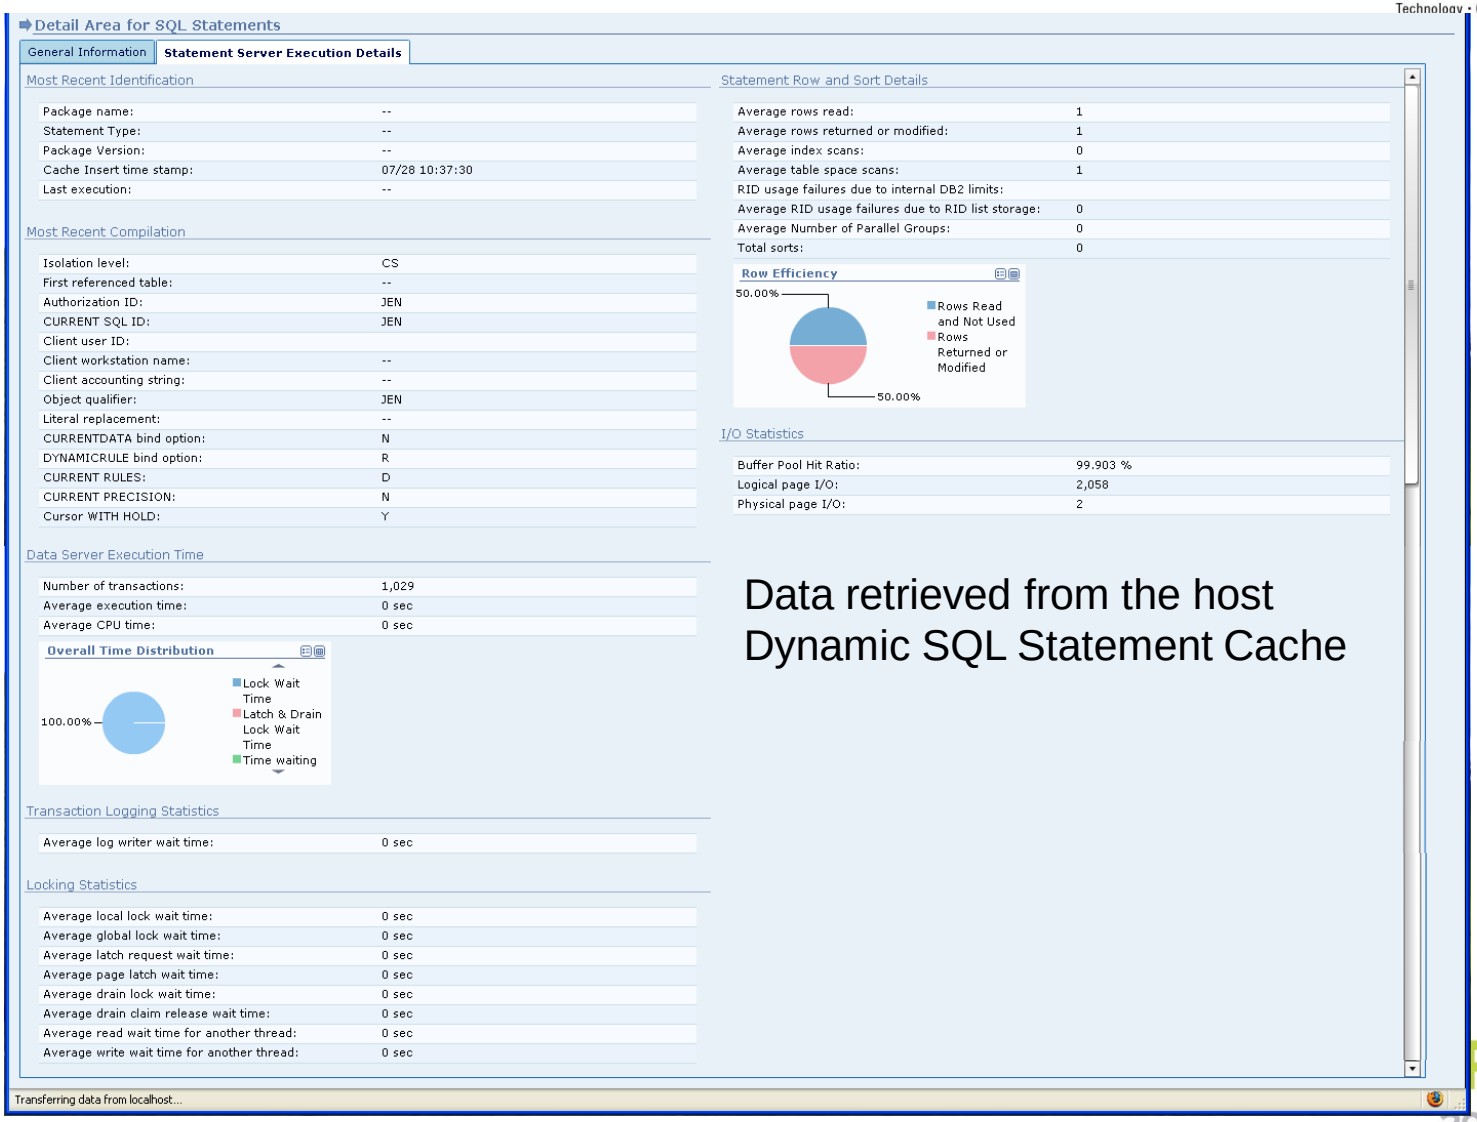

Data retrieved from the host Dynamic SQL Statement Cache

Extended Insight Analysis Dashboard

Select SQL from list and zoom into SQL level details

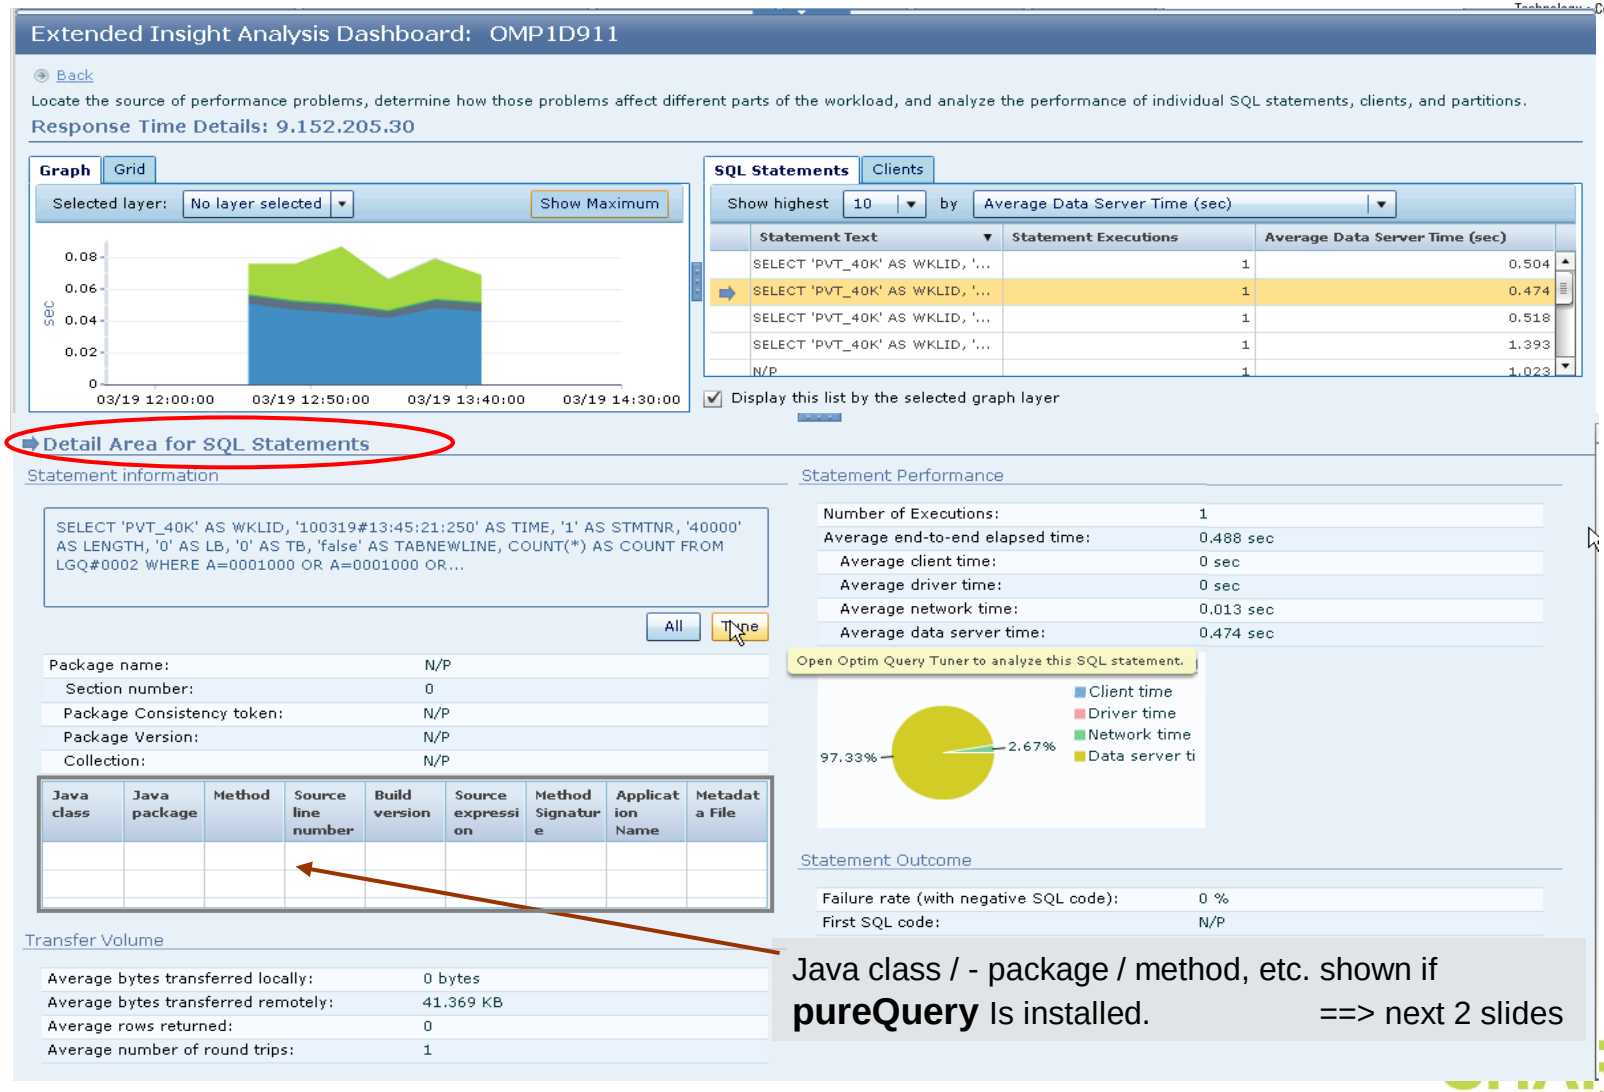

Java class / - package / method, etc. shown if **pureQuery** Is installed. ==> next 2 slides

How pureQuery Runtime can help !

Extended Insight feature can tell you where the query came from …

SHARE
Technology · Connections · Results

Application source

```
public class TestOPM {

    public static void main(String [] args)throws Exception{
        String url ="jdbc:db2://svl-imtestgl2.svl.ibm.com:50000/SAMPLE";
        Connection con = SampleUtil.getConnection(url, "db2admin", "hot6cold");
        ((com.ibm.db2.jcc.DB2Connection)con).setDB2ClientApplicationInformation("blah");

        Statement stmt = con.createStatement();
        for(int i = 0; i<10000; i++){
            stmt.execute("SELECT * FROM DB2ADMIN.INVENTORY");
            Thread.sleep(1000);
            System.out.println(i);
        }
    }
}
```

application name

Capture SQL with pureQuery runtime

```
pdq.captureMode=ON
pdq.executionMode=DYNAMIC
pdq.pureQueryXml=pureQueryFolder/capture.pdqxml
pdq.cmx.controllerURL=9.30.77.61:60000
```

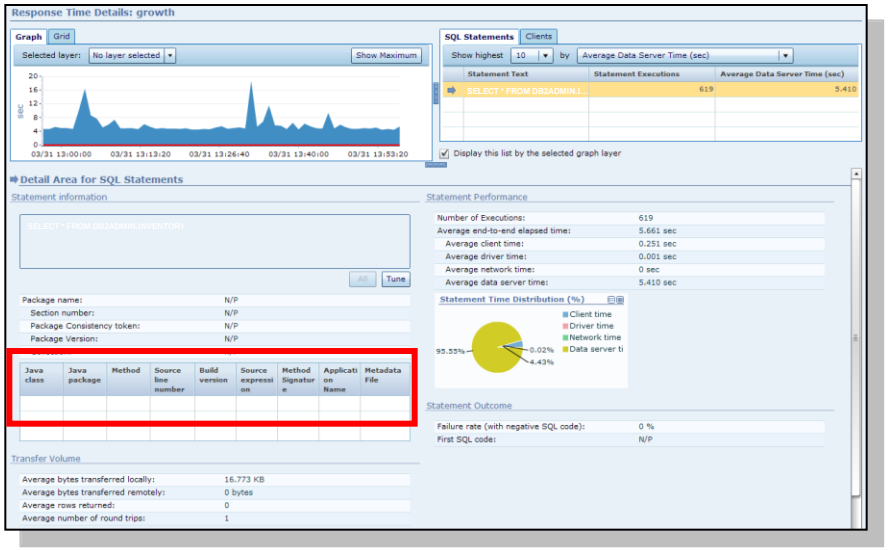

Upload collected metadata into OPM

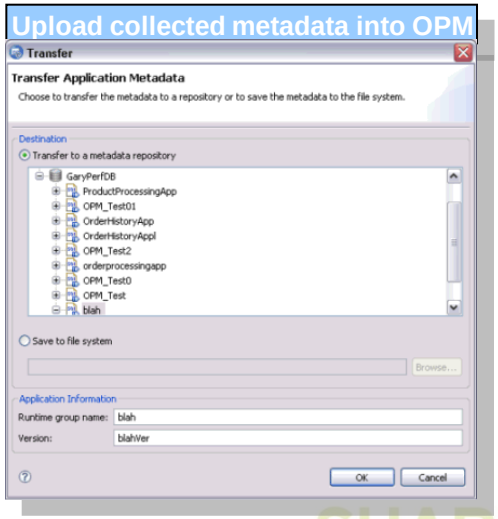

How pureQuery Runtime can help !

Extended Insight feature can tell you where the query came from …

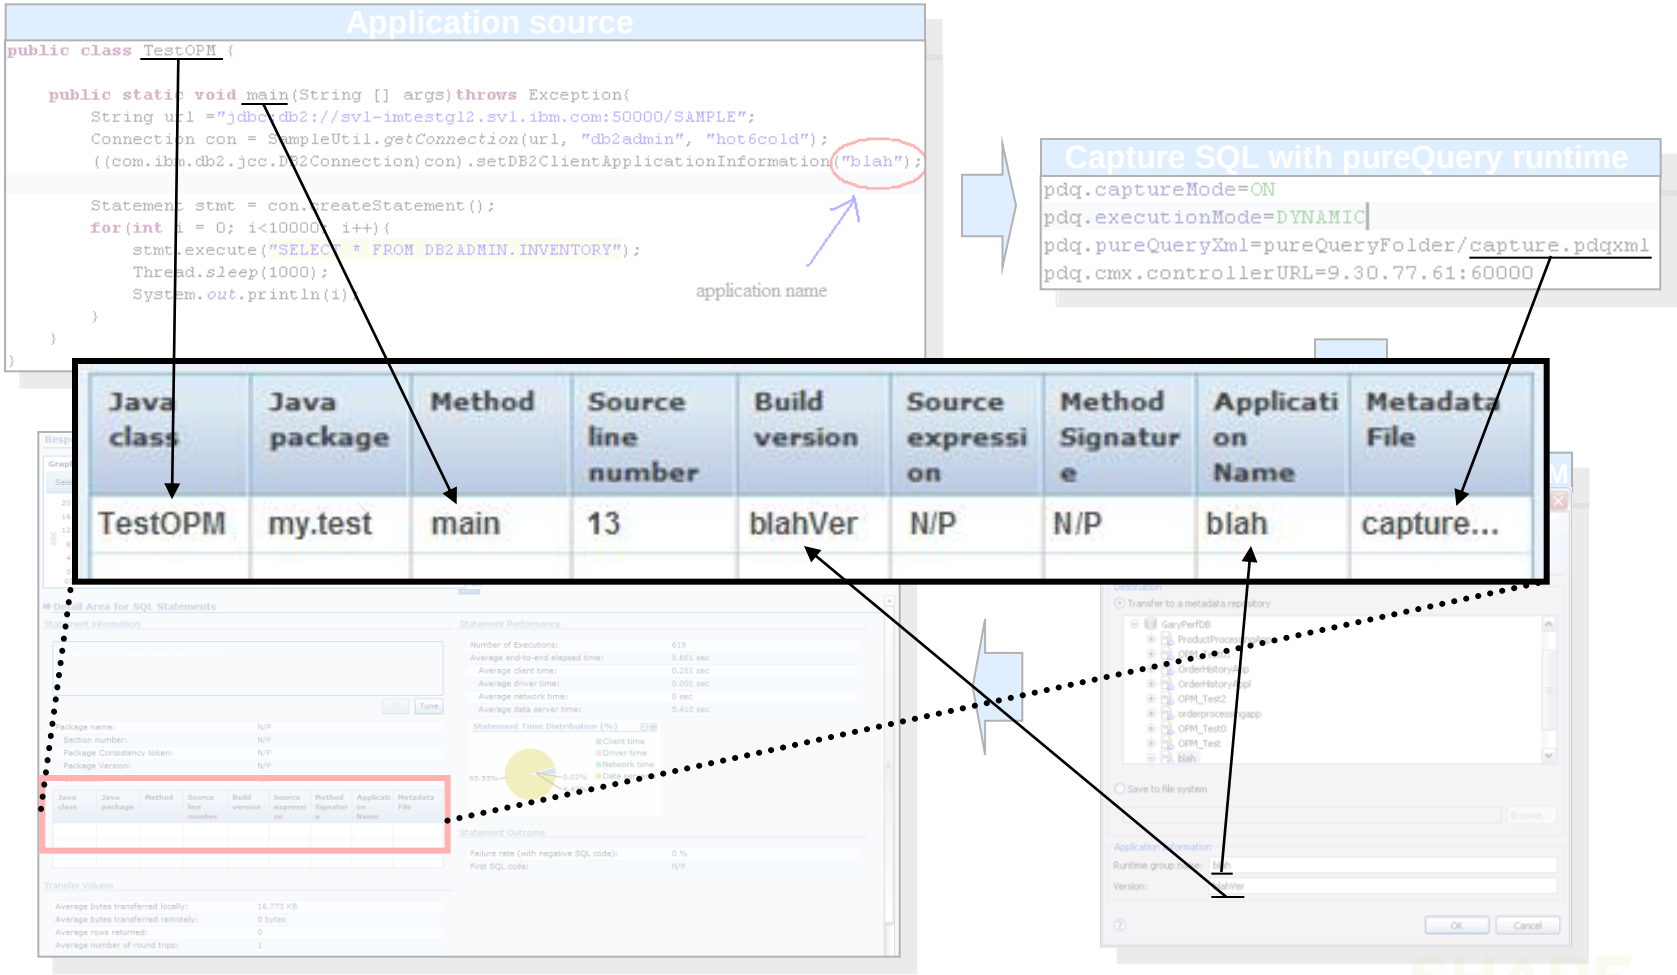

Extended Insight Analysis Dashboard

Select a plot area and see the corresponding performance metrics

Technology • Connections • Results

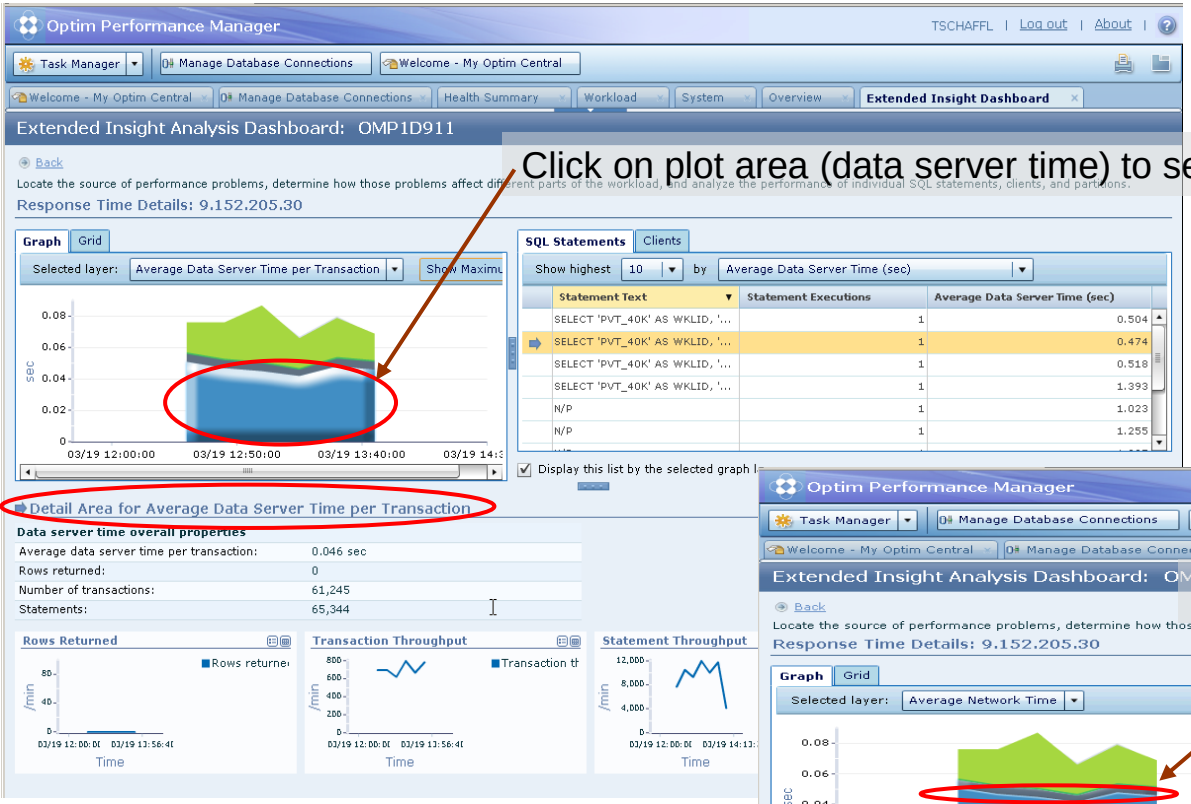

Click on plot area (data server time) to select

… for data server times and transaction throughput

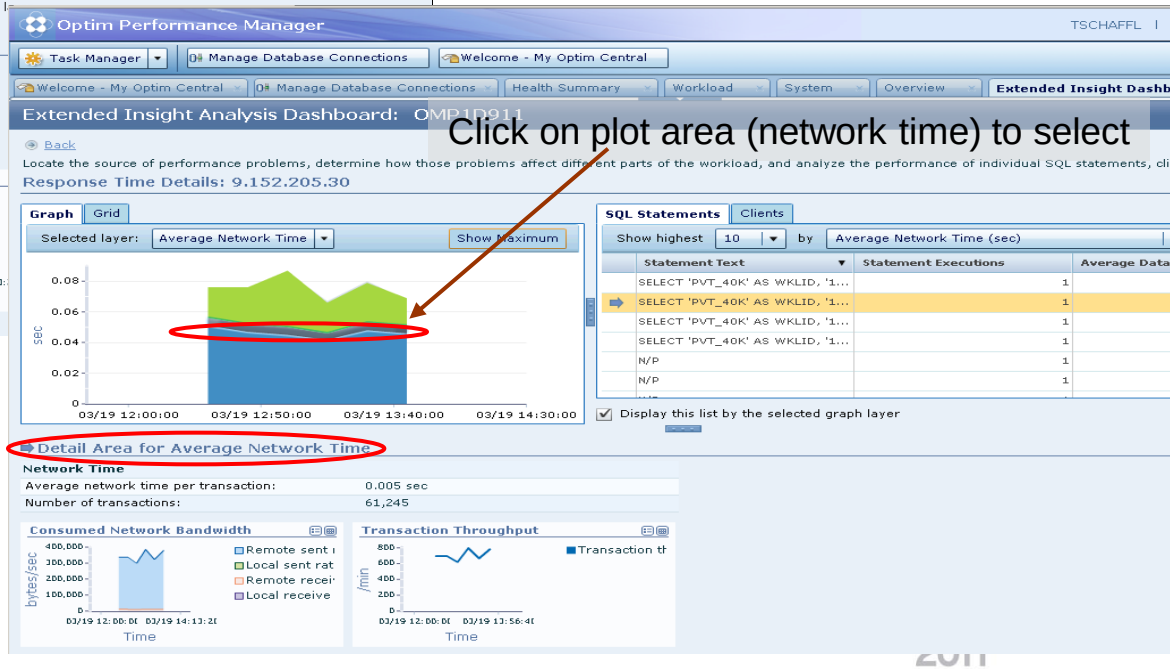

Click on plot area (network time) to select

… for network times and throughput

WebSphere – another area to be monitored in context

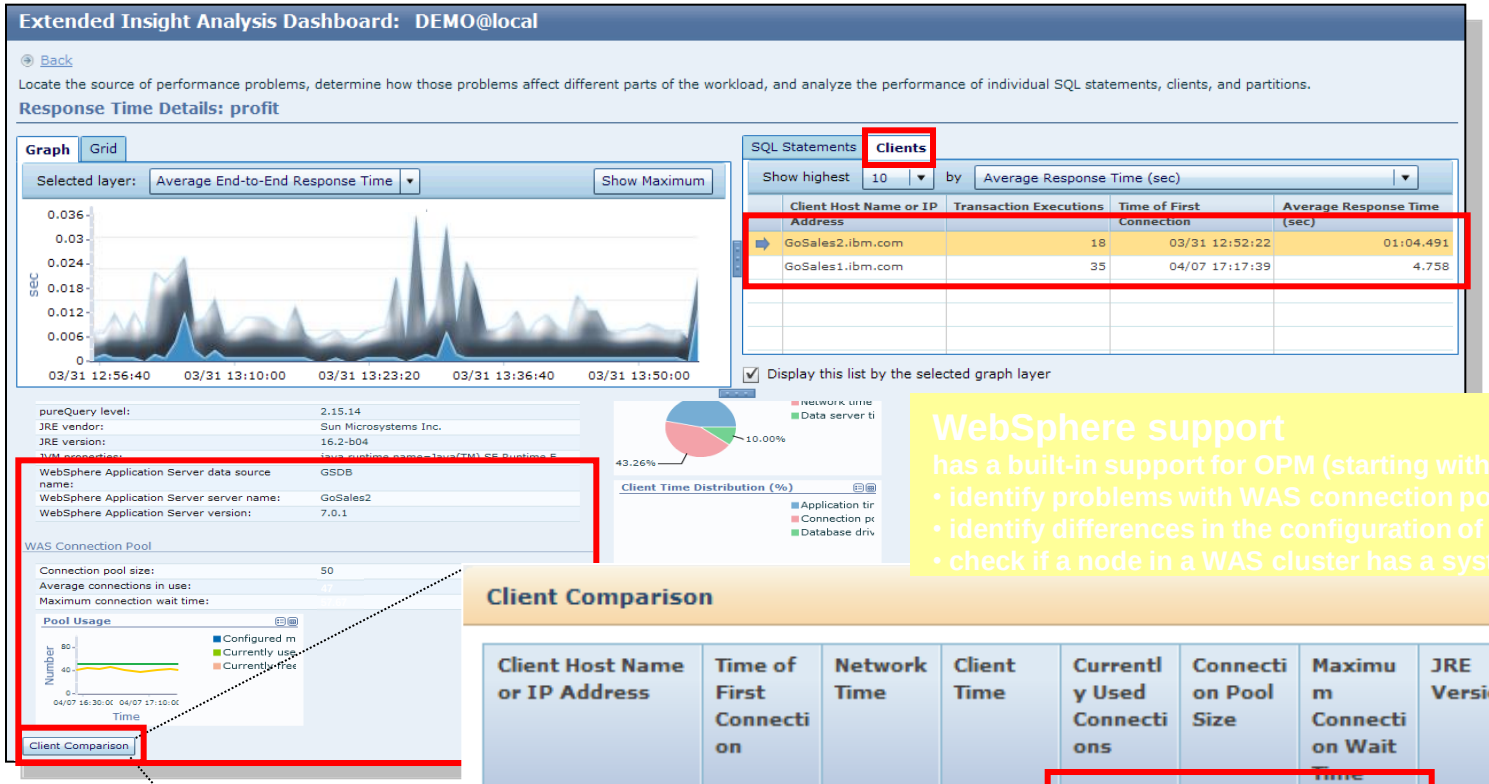

WebSphere support

has a built-in support for OPM (starting with WAS V6.0.21), allowing to …

- **identify problems with WAS connection pool**
- **identify differences in the configuration of nodes in a WAS cluster**
- **check if a node in a WAS cluster has a system or network problem**

Client Comparison

Client Host Name or IP Address	Time of First Connection	Network Time	Client Time	Currently Used Connections	Connection Pool Size	Maximum Connection Wait Time	JRE Version	Operating System	Database Driver Level
GoSales2.ibm.com	03/31 1…	13:50.8…	11:04.0…		50		16.2-b04	Window…	3.58.82
GoSales1.ibm.com	04/07 1…	7.367	25.336		100		16.2-b04	Window…	3.58.82

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

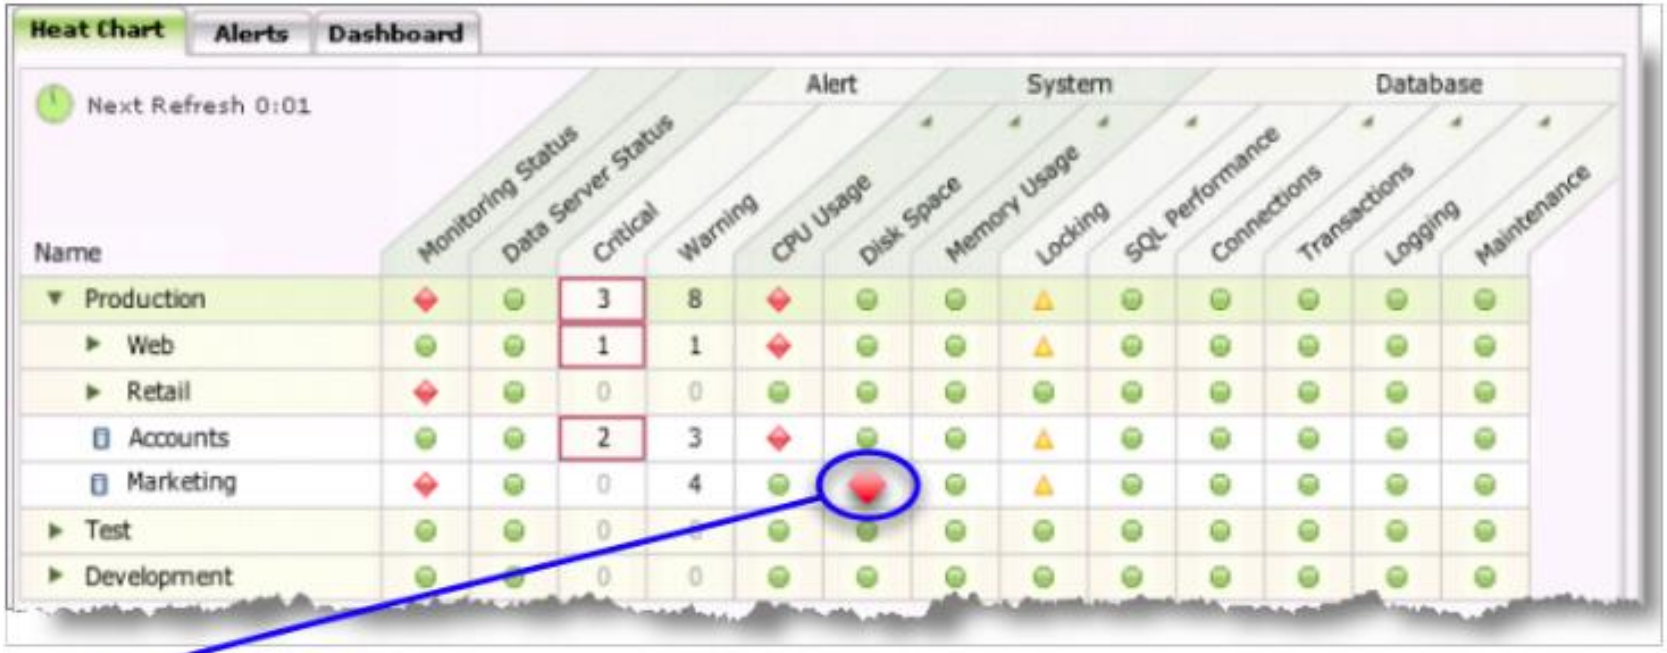

Disk Efficiency Alert

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

fully manual — high — TCO — low — fully automatic

- 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

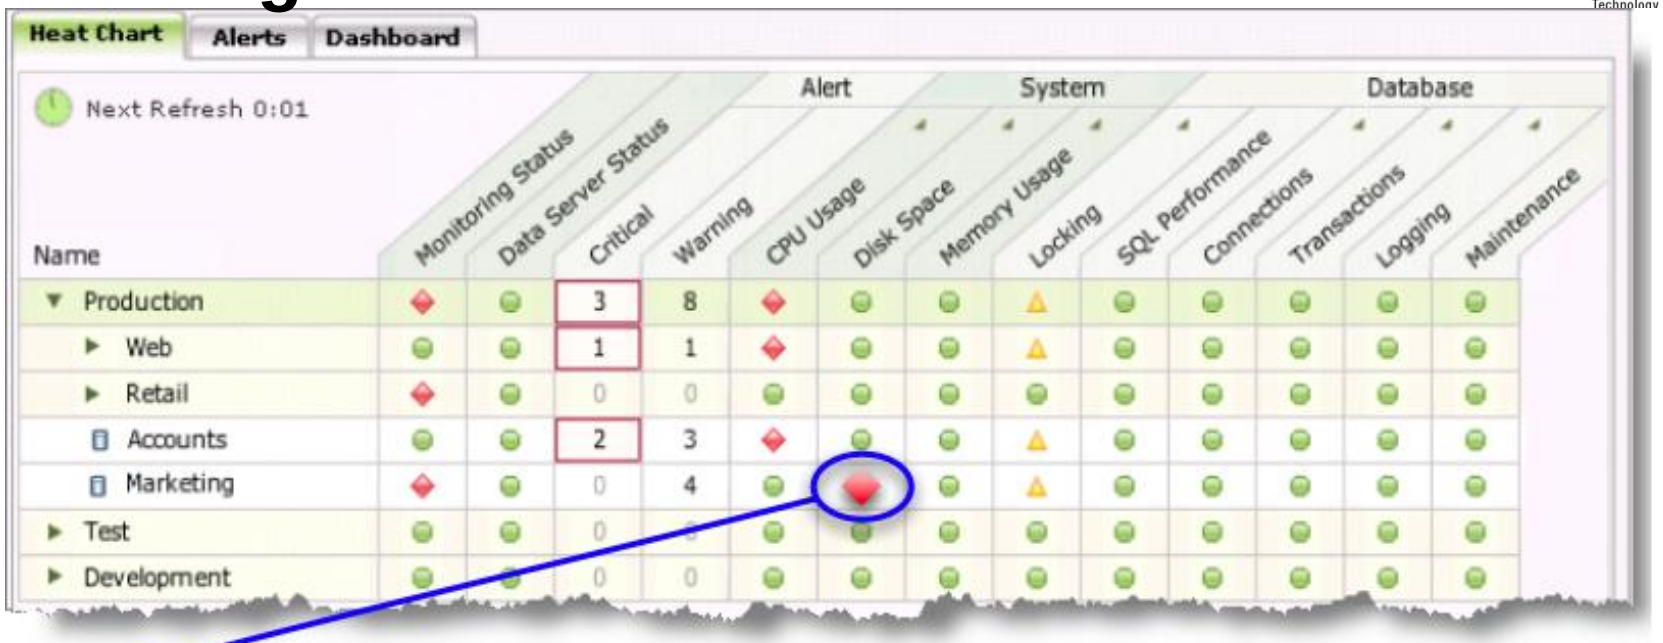

Best Practice Violation Alert: **Log and Data Files Collocated on Single Disk**

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."*

fully manual — TCO high / low — fully automatic

Manage by Exception to Lower TCO

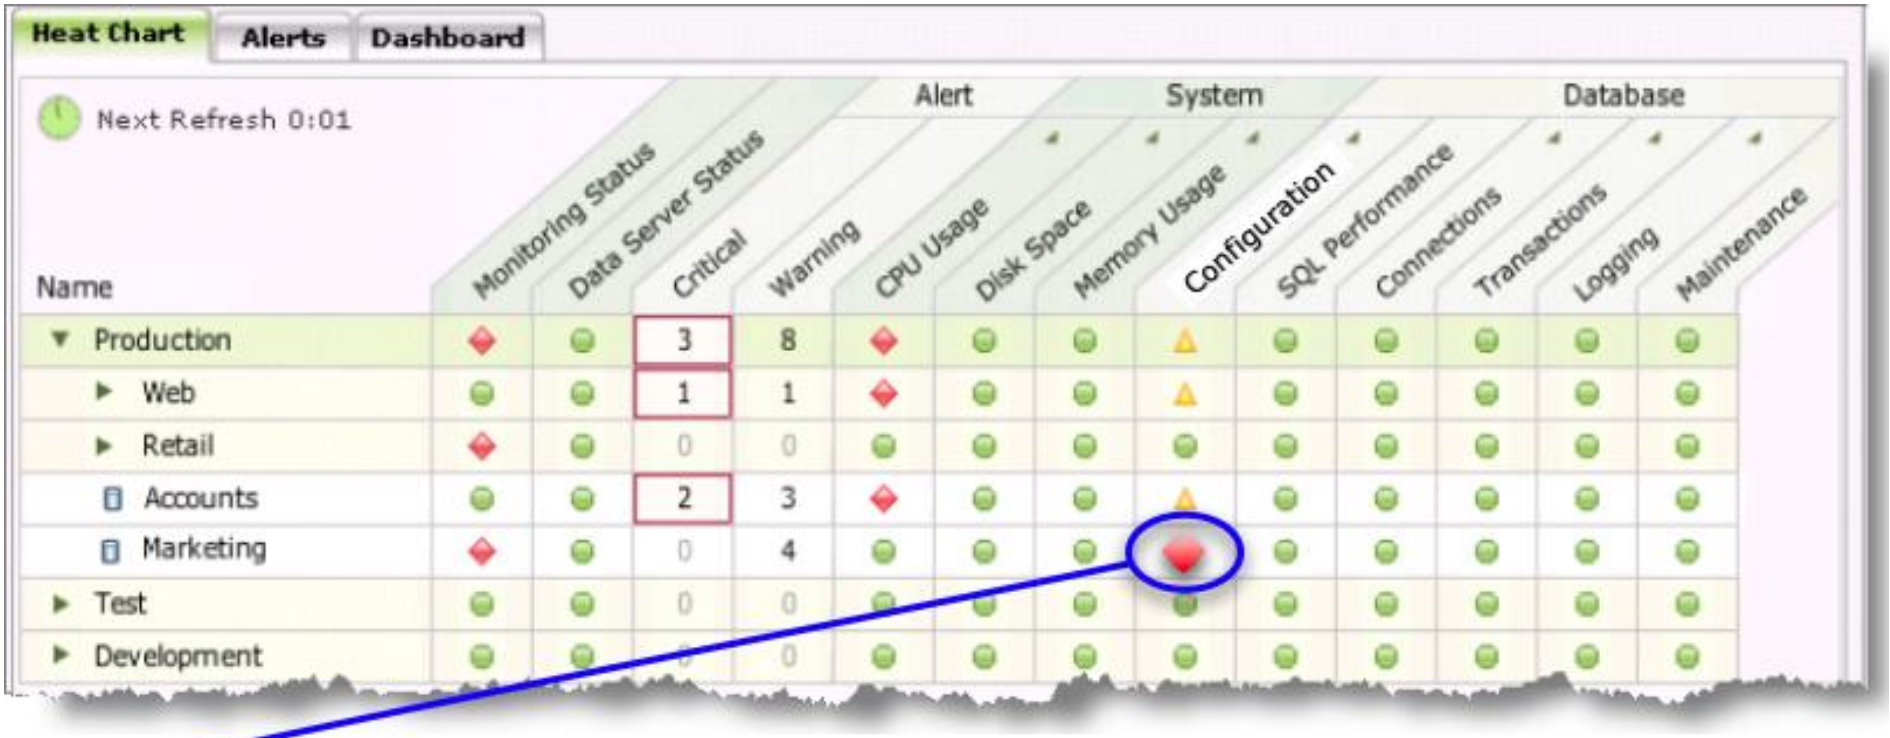

Configuration Change Alerts:

- Table space container 'tblsp25003' is no longer on disk
- Index 'idx00453' has been dropped from table 'mkt.tblA'
- Column 'A' of table 'mkt.tablA' 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

SHARE
Technology • Connections • Results

Data Management Communities for DB2

- IDUG – the worldwide community of DB2 users
  - Membership is FREE – join today! www.idug.org

- Data Management Community – share and interact with peers around the world

IOD 2009
Content Planning

  - 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

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.