



# Managing the Performance of Your IBM DB2 z/OS Applications by Using IBM InfoSphere Optim Solutions

**Peter Schouboe, IBM SWG Nordic**  
**Norbert Jenninger, IBM**



# Disclaimer

© Copyright IBM Corporation 2010. All rights reserved.

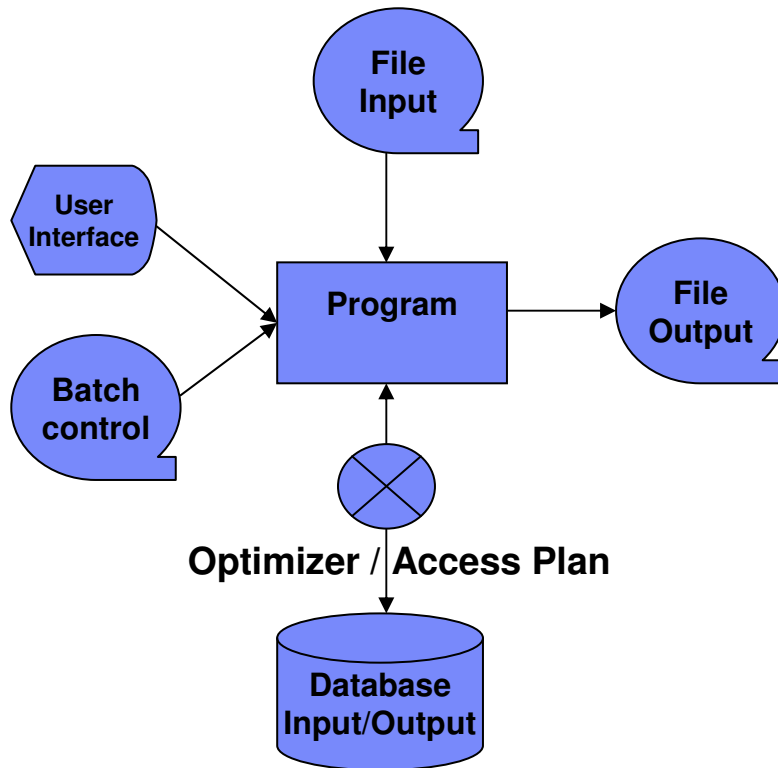
*U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.*

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

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

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

# Performance Issues



## ■ Performance issues – main areas

- Program
  - Number of file / database trips
  - DataChecks - Commitpoints
  - Internal data manipulation
  - Complex calculations
- Transaction Managers
  - Number of threads / active users vs. Region sizes
  - Transaction mgmt protocol
- Database
  - Optimized database design
  - Index use vs. Direct data access
  - Advanced vs simple SQL
  - Database Optimizer help
- Communication
  - Number of trips to database
  - Amount of data shipped
  - Number of threads

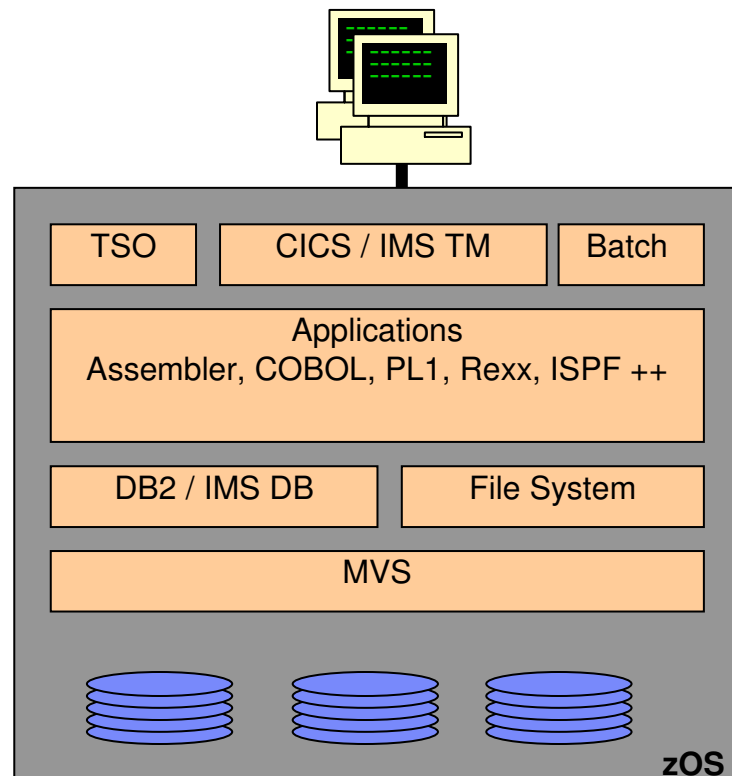
### Main components to watch for performance reasons:

Operating Systems – Transaction Managers – Database System – Program language compilers – Physical devices / box, connector, memory etc...

## Application development on System z – old school ☺

### ▪ Traditional Mainframe development

zOS Application Server



### ▪ Developing for performance relies on 2 min factors:

- Monitors for each main component
  - Transactions, Batch jobs, Databases, Operating system
    - AddressSpace Monitors
  - SQL monitors
- Advise for optimization
  - App. Analyzers
  - Explain (DB2 Optimizer Guess)

Complicated enough, but all inside one known box, under one ruleset – the zOS / MVS .

Complications mainly from,

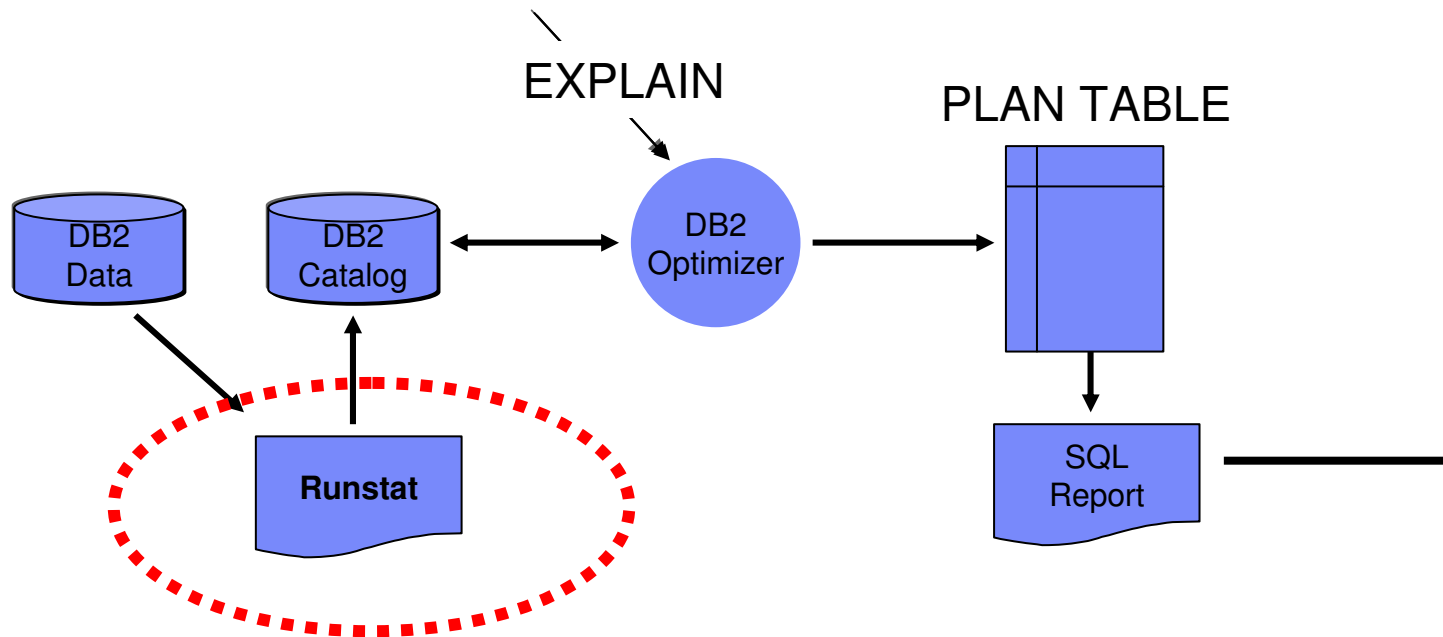
- system & core sw. upgrades
- database changes
- physical improvements



# Optimizing SQL – The process

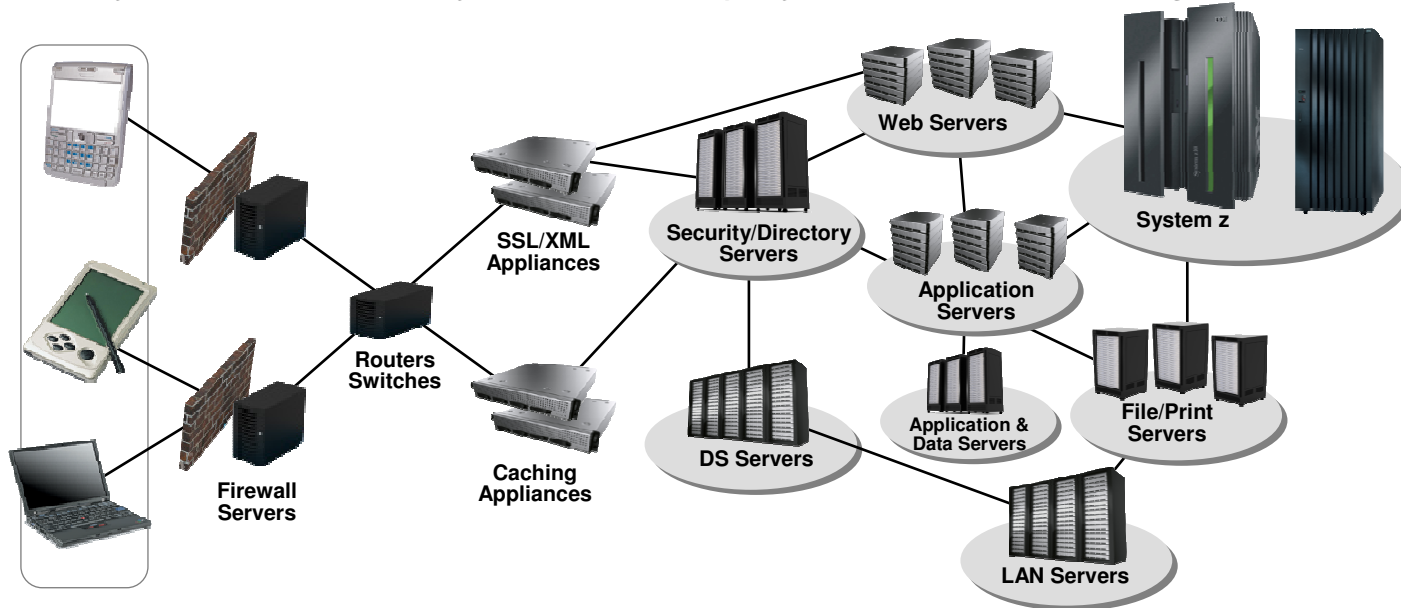
```

SELECT GOSALESC.T.CUST_ORDER_DETAIL.CUST_ORDER_DETAIL_CODE,
GOSALESC.T.CUST_ORDER_DETAIL.CUST_ORDER_NUMBER,
GOSALESC.T.CUST_ORDER_DETAIL.CUST_SHIP_DATE,
GOSALESC.T.CUST_ORDER_DETAIL.CUST_UNIT_PRICE
FROM
GOSALESC.T.CUST_ORDER_DETAIL JOIN GOSALESC.T.CUST_ORDER_HEADER ON
GOSALESC.T.CUST_ORDER_DETAIL.CUST_ORDER_DETAIL_CODE =
GOSALESC.T.CUST_ORDER_HEADER.CUST_ORDER_NUMBER
WHERE GOSALESC.T.CUST_ORDER_HEADER.CUST_CODE =
GOSALESC.T.CUST_ORDER_HEADER.CUST_CODE
    
```




Today, The IT model looks different ...

Real-time event-driven workloads, richer content, and modular technologies alter the composition of *systems and how systems are deployed as well as managed*

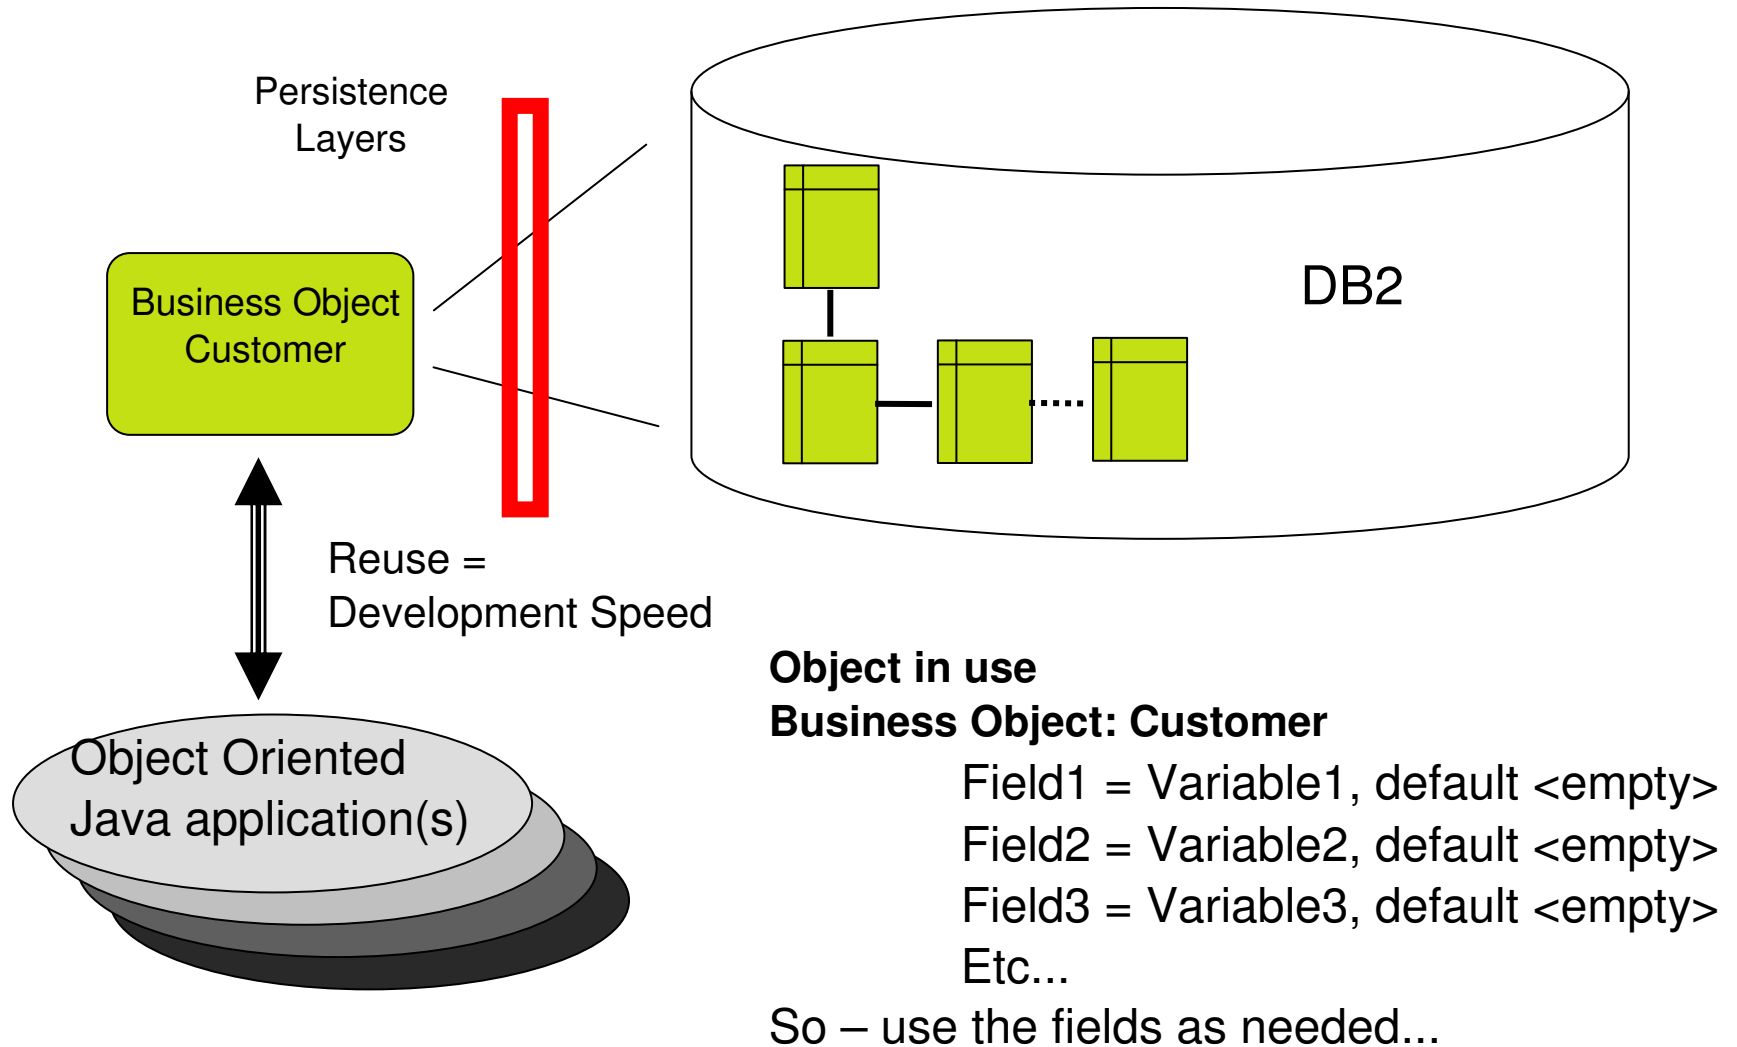


- Application software and middleware is being deployed on several platforms
- Much more complicated to monitor and optimize.
- New, and different methods need to be considered when developing, monitoring and optimizing applications.



# Object oriented versus Relational

## The built-in original conflict



# Issues in the Optimization process

- **Explain requires that you have access to /can find the code**
  - It requires that you have access to the SQL
  - Many times it's very hard to predict dynamic SQL
- **Explain will not advise on how to change program / SQL**
- **Explain bases itself on known Index structures (Fix ?Virtual Index, 8, 9 & 10)**
- **Explain does not cover trips across the Net**
- **Explain does not help minimize data shipped**
- **Explain needs access to production statistics....(Fixed in DB2 10 ☺ !!)**
- **It's easy to type errors, and do you know your data ?**
- **Testing corrections is difficult**
- **The problem of introducing rewrites**
  - Price of corrections grows the later you have to change....
- **It's becoming still harder to figure out if there is an issue, and where you need to optimize.**

## Other Issues – Dynamic vs. Static SQL

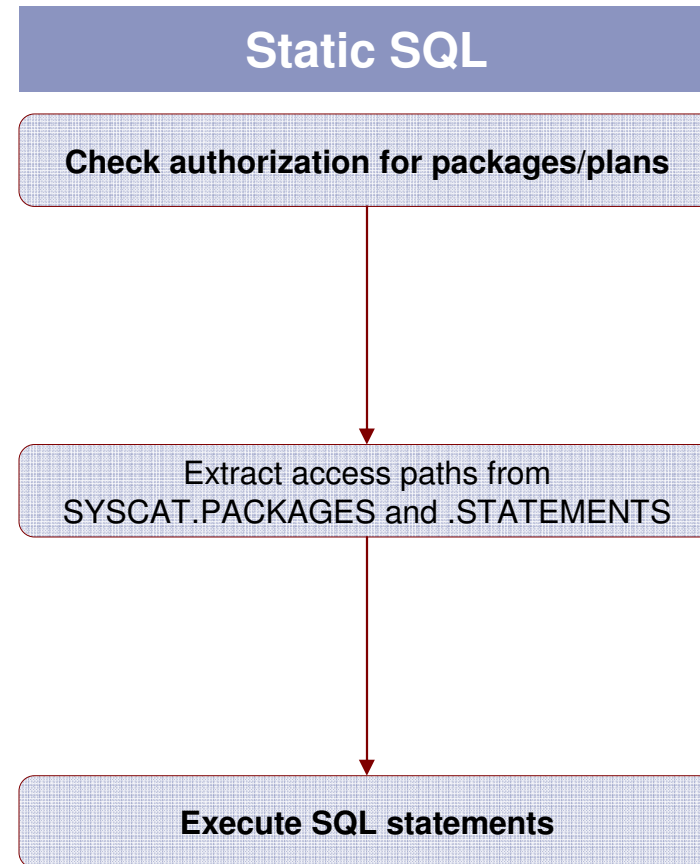
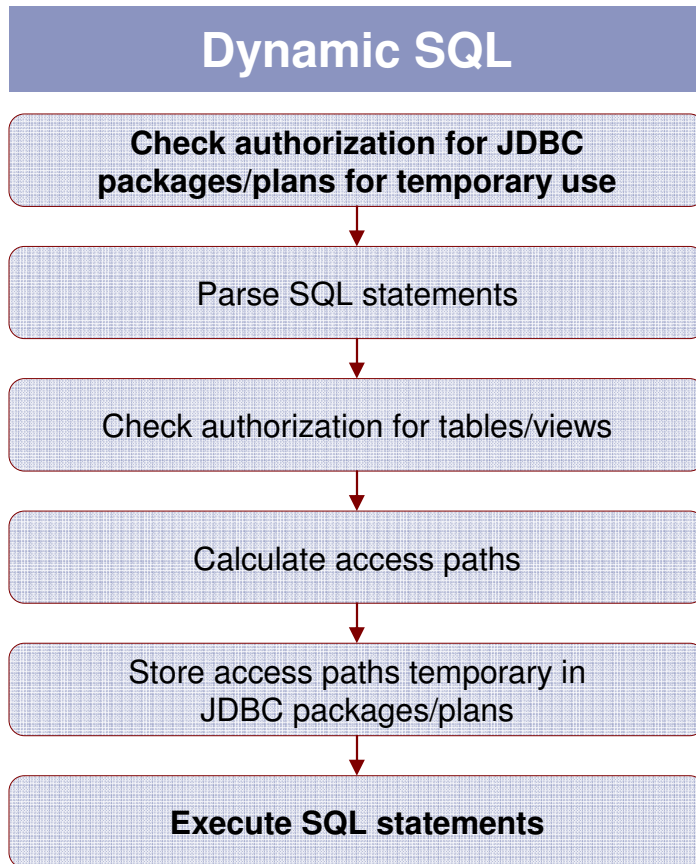
- **Dynamic SQL is to development what Chokolade is to a child....**
  - It's quick...
  - It solves your problem...
  - It comes in many forms and shapes...
- **But ....**
- **It's not going to help meet diet goals**
- **It will 'cheat' the system burn-rates**
- **And – if left alone.....**



www.shutterstock.com · 31572991



# Difference between “healthy food” and “chokolate”....



DB2 will reuse SQL from dynamic statement cache if possible. **New feature “&” in DB2 10 will help.** Alternatively using “?” parameter markers rather than Literals will enable reuse.

**Still – sheer number of threads might flush statement from the dynamic cache.....**

SQL is bound into a **package** beforehand. All SQL parsing and initial optimization **work is done ahead of time** and only once. Statements are fully optimized during execution.

## Management Issues with dynamic SQL...

- Every statement potentially fires – Prepare, Bind, Execute
- All uses same Plan/User/Authority
- Access needs to be grant on DB2 Tables/Views
- Literal vs Parameter markers (better in DB2 10)
- Hard to find specific SQL

```

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

```



## IBM Tools to support the optimization process

- **z/OS / ISPF:**           **SQL performance Analyzer**
  
- **z/OS / GUI:**           **Data Studio (free)**
  - Visual Explain
  - SQL Editor
  - Statistics Advisor
  
- **z/OS / GUI:**           **Infosphere Optim Development Studio**  
**Infosphere Optim pureQuey**  
**Infosphere Optim Query(Workload)Tuner**
  
- **z/OS / GUI:**           **Omegamon for DB2 PE Extended Edition**

# Editing, Testing and Optimizing SQL with Development Studio

The screenshot displays the IBM Development Studio interface. The top window is the SQL sensitive Editor, showing a Java file with SQL code. The bottom window is the Performance Data Set, displaying a table of execution metrics. The SQL Outline window is also visible, showing a list of SQL queries.

**SQL sensitive Editor**

```

resultSet.close();
statement.close();

statement = connection.createStatement();
if(statement.execute("SELECT CUST_CODE, STDDEV(ORD_STOT_COST) AS TOTAL, CUST_CODE FROM GOSALESC")
resultSet = statement.getResultSet();
while(resultSet.next()) {
    System.out.print("customer id: " + resultSet.getString(1) + ", ");
    System.out.print("net sales:" + resultSet.getDouble(2) + "\n");
}

```

**Performance Data Set**

Schemas	Number of Times Run	Total Time	Max Time	Average Time	Min Time
GOSALESC					
CUST					
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	1	90.92	90.92	90.92	90.92
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	2	297.06	273.61	148.53	23.45
SELECT CUST_CODE, CUST_CITY FROM GOSALESC.CUST ORDE	3	20.09	7.99	6.70	4.97
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	1	2.70	2.70	2.70	2.70
SELECT count(CUST_CODE) FROM GOSALESC.CUST	14	61.04	40.82	4.36	1.20
SELECT count(CUST_CODE) FROM GOSALESC.CUST where C	14	61.04	40.82	4.36	1.20
ReportGenerator	-	-	-	-	-
CUST_CODE	-	-	-	-	-
CUST_ORD					
SELECT CUST_CODE, STDDEV	4	7.11	2.18	1.78	1.35
SELECT ORD_NBR AS ORDER	5853	3335.63	46.98	0.57	0.22
SELECT CUST_CODE, SUM(ORD	3	5.80	2.38	1.93	1.63
SELECT CUST_CODE, AVG(ORD	4	5.17	1.40	1.29	1.14
SELECT CUST_CODE, CORRELA	3	3.24	1.20	1.08	1.00

**SQL Outline**

- Find in Source
- Run SQL
- Show in SQL Editor
- Export SQL to File...
- Compare
- Retrieve EXPLAIN Data
- Open Visual Explain
- Tune SQL
- Generate pureQuery Code...
- Find in pureQueryXML

Compare performance data of various versions of your application

Visualize application SQL

Visualize execution metrics

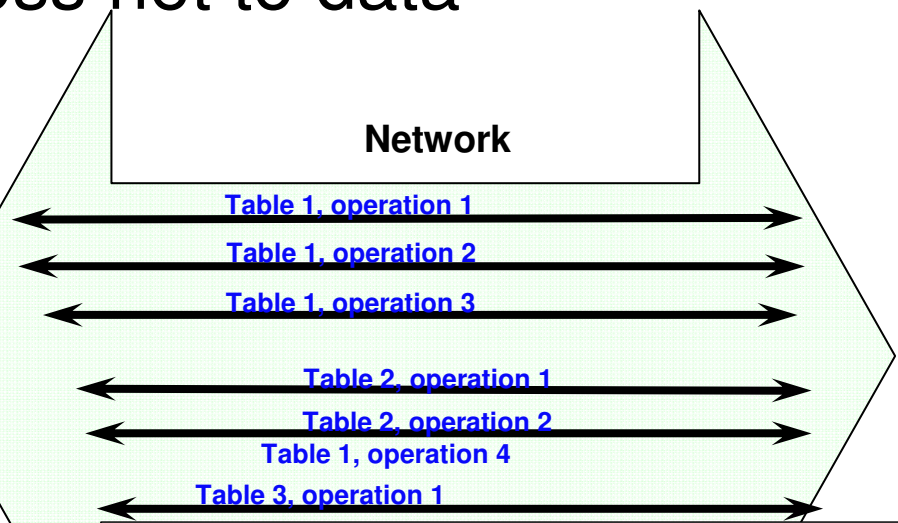

Correlate with data source attributes

Execute, tune, share, trace, explore SQL

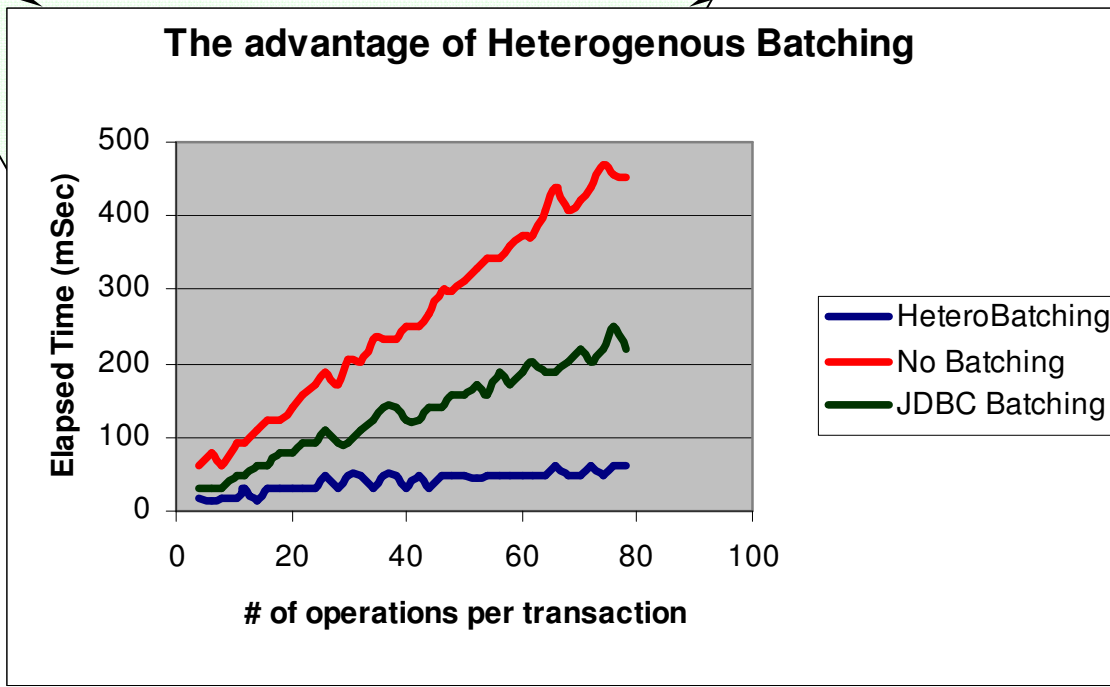


# Reduce trips across net to data

- Heterogenous Batching



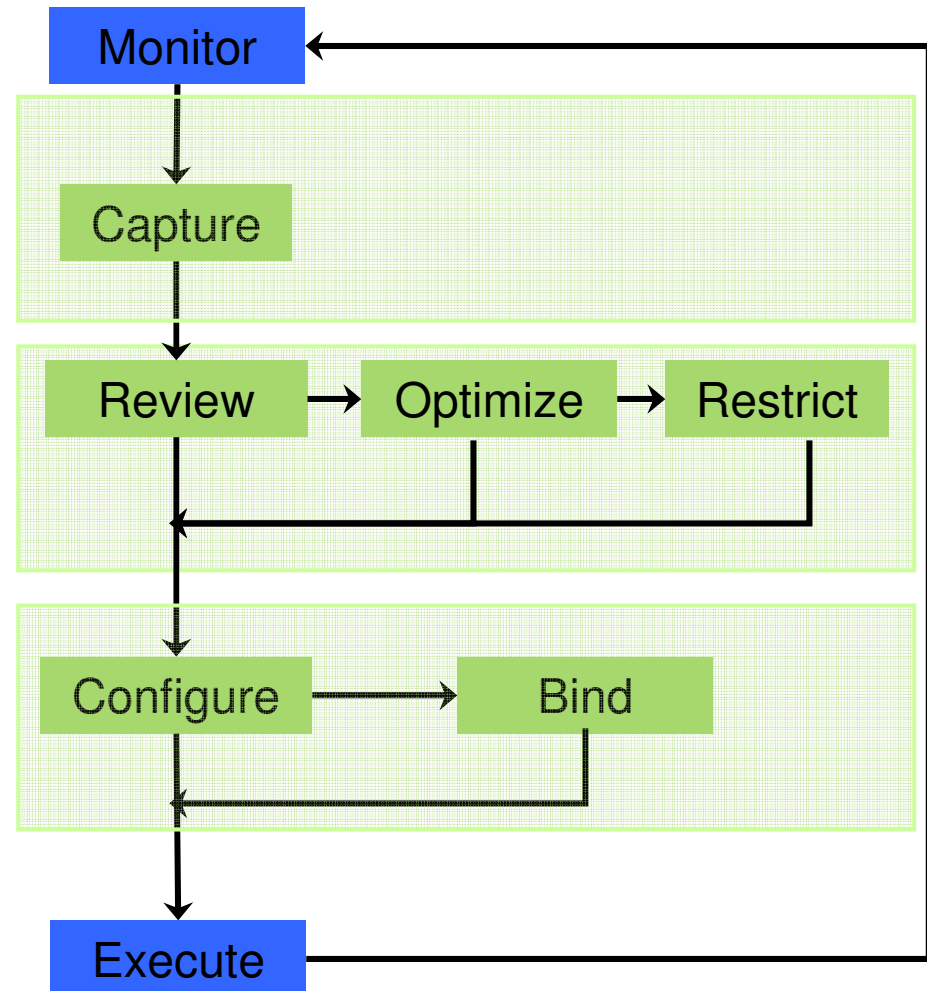
**Heterogenous Batching** – Multiple operations across different tables all execute as one batch



# Finding the SQL/Program 2...

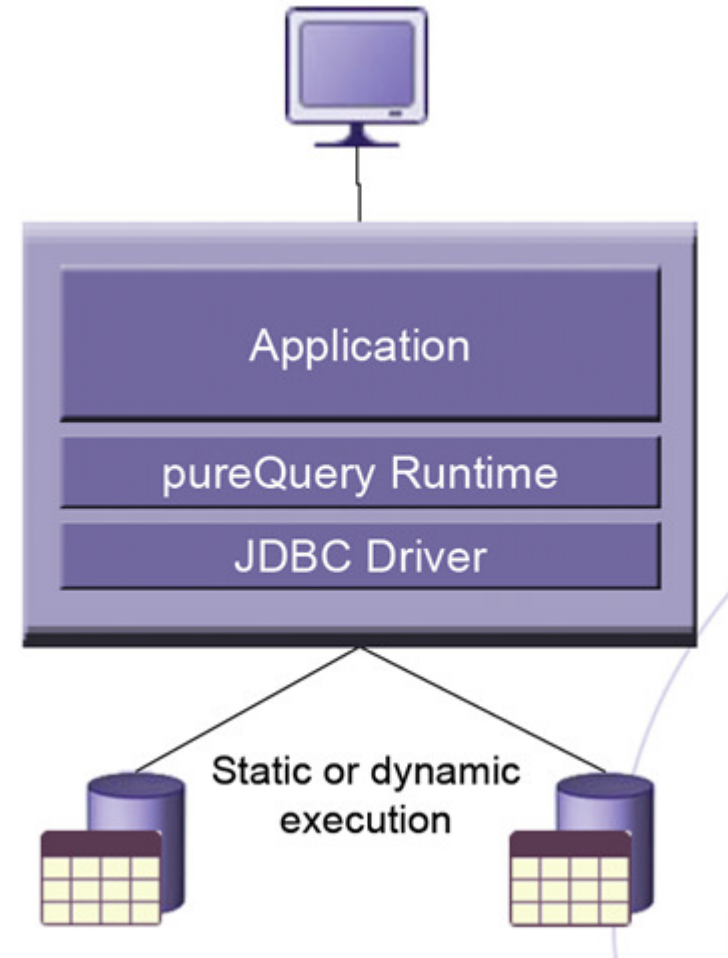
## Managing already developed Applications

- **Monitor**
  - Review performance
  - Analyze hotspots
- **Capture**
  - Capture performance and application metadata
- **Optimize**
  - Review and share SQL
    - Visualize hotspots
    - Analyze impacts from schema changes
    - Trace SQL to originating source
  - Optimize and replace SQL
  - Create approved SQL list
- **Deploy**
  - Configure execution properties
  - Optionally bind for static execution



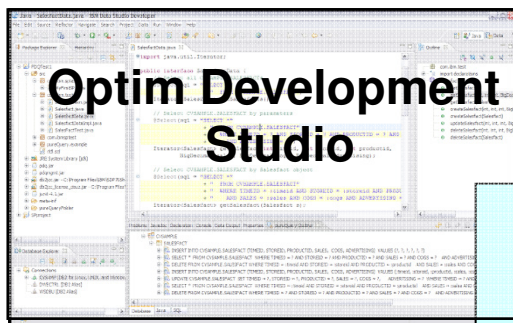
## Using pureQuery in development & Production...

- **Up front development vs. After the fact**
- **You decide at runtime....**
  - (Dynamic/Static/Mix)
- **DB2 / Oracle / Informix**



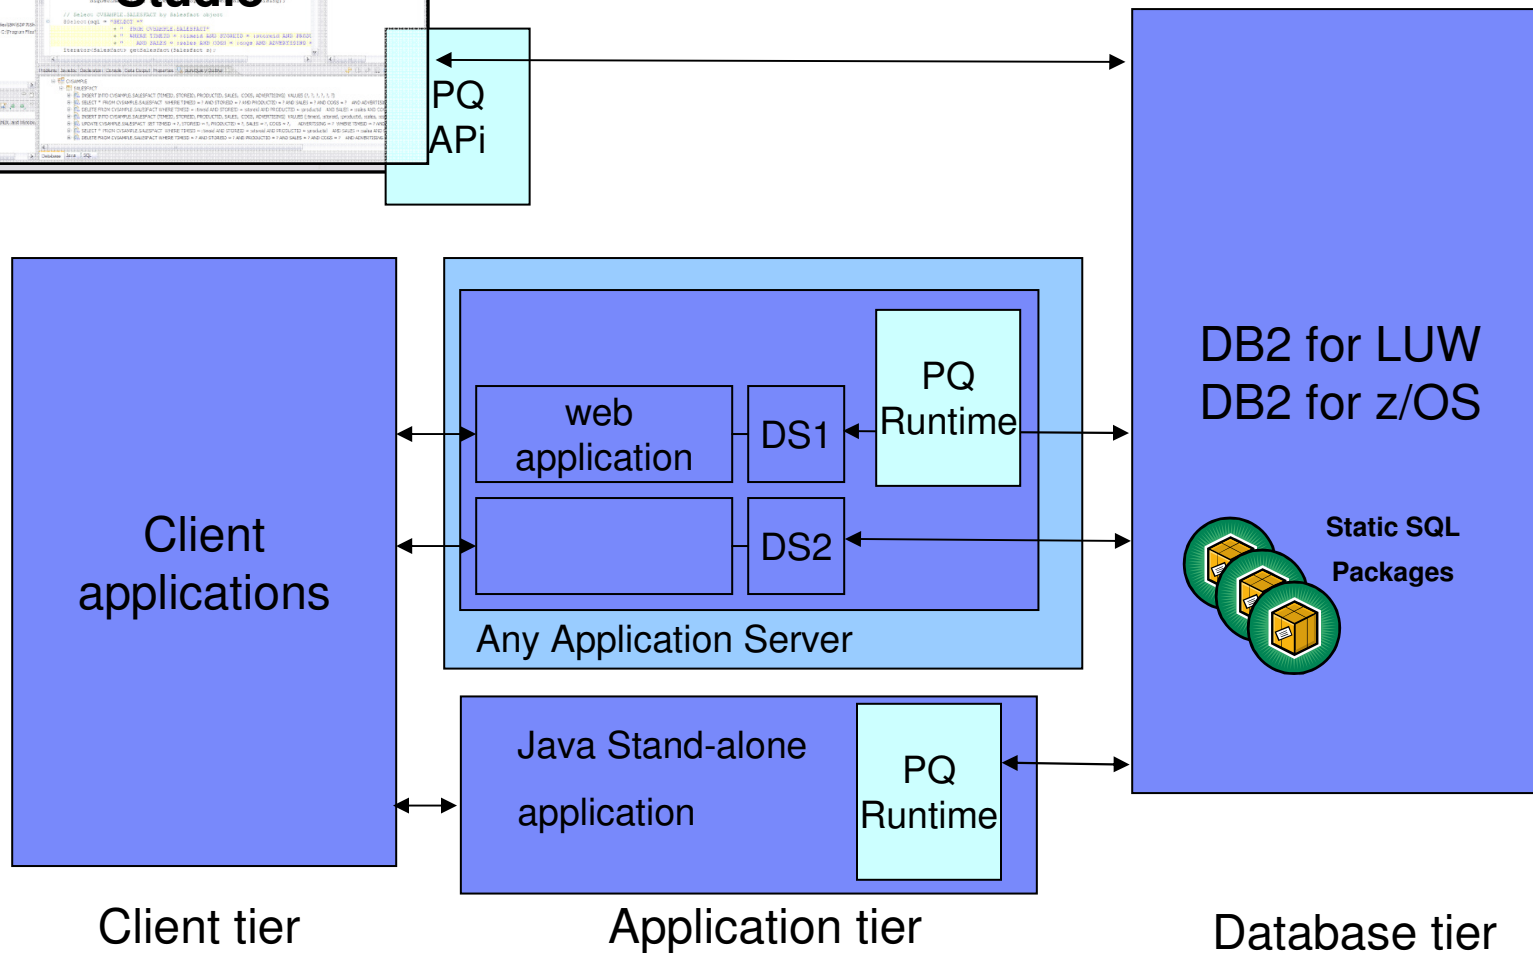
# Find the SQL and the program using pureQuery

pureQuery	pureQuery
AppServer	DB2 Connect



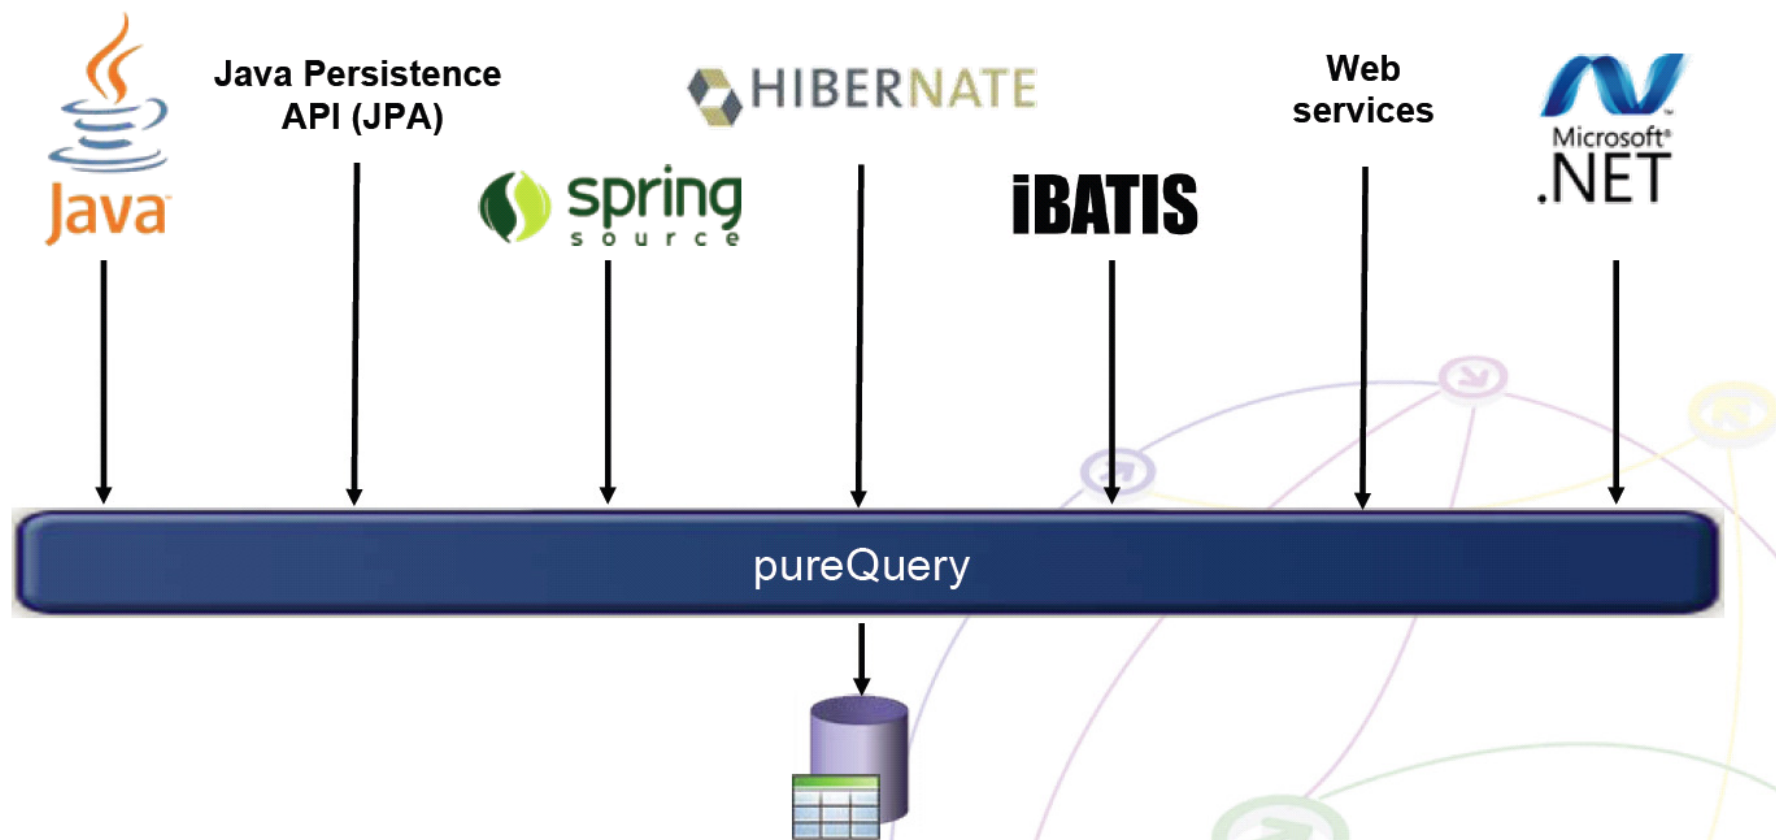
**Optim Development Studio**

PQ API





So – does this support .....



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



# Editing, Testing and Optimizing SQL with Development Studio

The screenshot displays the IBM Optim Development Studio interface. The main editor shows a Java file named `ReportApplication.java` with the following SQL code snippet:

```

resultSet.close();
statement.close();

statement = connection.createStatement();
if(statement.execute("SELECT CUST_CODE, STDDEV(ORD_STOT_COST) AS TOTAL, CUST_CODE FROM GOSALESCUST")) {
    resultSet = statement.getResultSet();
    while(resultSet.next()) {
        System.out.print("customer id: " + resultSet.getString("CUST_CODE") + ",");
        System.out.print("net sales:" + resultSet.getDouble("TOTAL"));
        System.out.println();
    }
}

```

Below the editor, the Performance Data Set table is visible, showing execution statistics for various SQL queries. The table has columns for Schemas, Number of Times Run, Total Time, Max Time, Average Time, and Min Time.

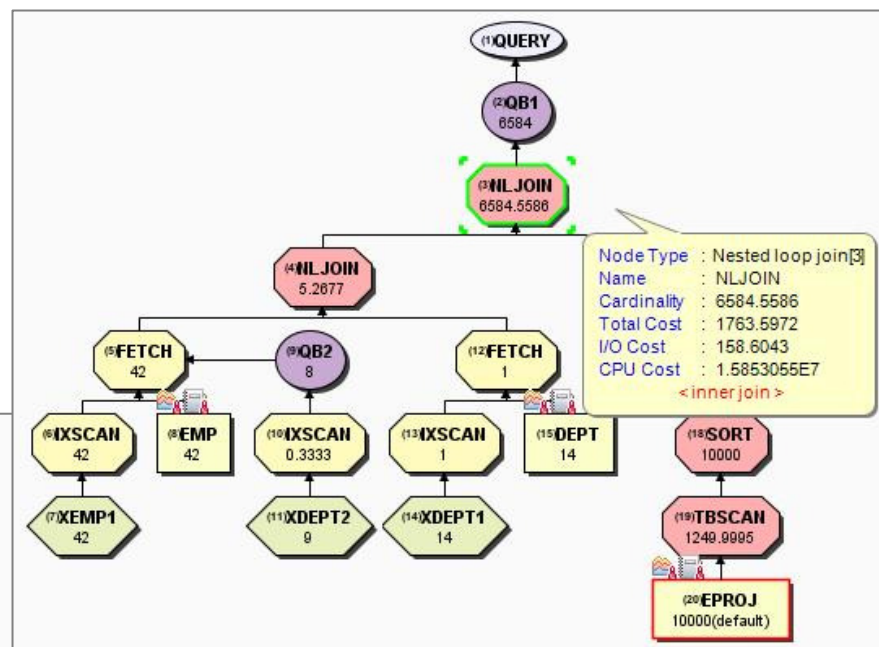
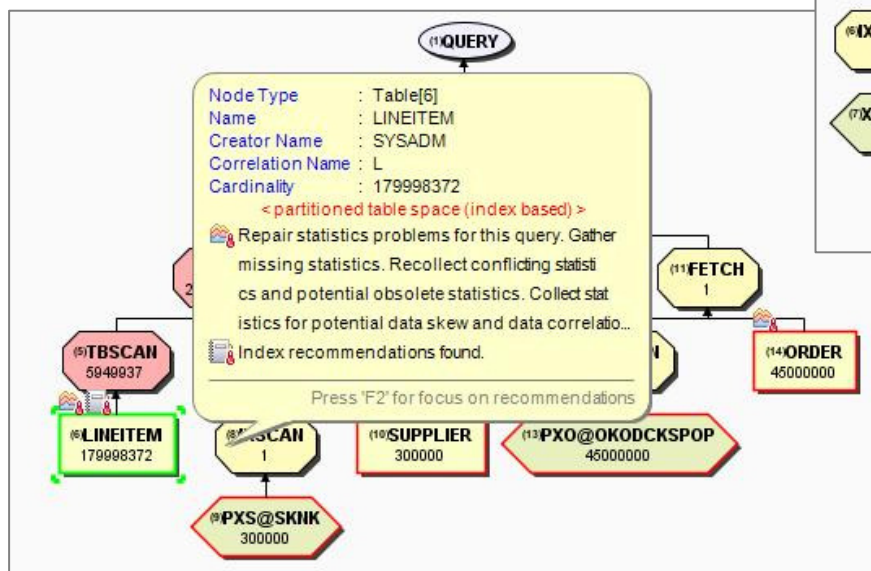
Schemas	Number of Times Run	Total Time	Max Time	Average Time	Min Time
GOSALESCUST					
CUST					
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	1	90.92	90.92	90.92	90.92
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	2	297.06	273.61	148.53	23.45
SELECT CUST_CODE, CUST_CITY FROM GOSALESCUST.CUST ORDE	3	20.09	7.99	6.70	4.97
SELECT CUST_CODE, CUST_FRST_NAME, CUST_LAST_NAME, CU	1	2.70	2.70	2.70	2.70
SELECT count(CUST_CODE) FROM GOSALESCUST.CUST	14	61.04	40.82	4.36	1.20
SELECT count(CUST_CODE) FROM GOSALESCUST.CUST where C	14	61.04	40.82	4.36	1.20
ReportGenerator	-	-	-	-	-
CUST_CODE					
CUST_ORD					
SELECT CUST_CODE, STDDEV(	4	7.11	2.18	1.78	1.35
SELECT ORD_NBR AS ORDER	5853	3335.63	46.98	0.57	0.22
SELECT CUST_CODE, SUM(ORD	3	5.80	2.38	1.93	1.63
SELECT CUST_CODE, AVG(ORD	4	5.17	1.40	1.29	1.14
SELECT CUST_CODE, CORRELA	3	3.24	1.20	1.08	1.00

A context menu is open over the table, listing actions such as "Find in Source", "Run SQL", "Show in SQL Editor", "Export SQL to File...", "Compare", "Retrieve EXPLAIN Data", "Open Visual Explain", "Tune SQL", and "Generate pureQuery Code...". A callout bubble points to this menu with the text: "Execute, tune, share, trace, explore SQL".

# Visual Explain – Pictures talks...

- Visualize access path

- See flow of query processing
- See indexes and operations
- See optimizer rationale



# Figure out what triggers Explain/DB2 Optimizer

Data Studio  
QueryTuner

**New!**  
Repair and Complete  
Control Statements

- Provides advice on
  - Missing statistics
  - Conflicting statistics
  - Out-of-date statistics

Statements | Recommendation Summary | **Workload Statistics Advisor** | Workload Query Advisor

Existing statistics status - 5 tables need repair out of the 7 tables that were checked

**Repair** Complete

This version of the RUNSTATS command repairs the problems that the Workload Statistics Advisor found. Run this version to conserve time

RUNSTATS Control Statements

```

RUNSTATS TABLESPACE DB4LINEI.TSLINEI
TABLE(SYSADM.LINEITEM) SAMPLE 5
COLGROUP(L_QUANTITY) FREQVAL COUNT 10
COLGROUP(L_DISCOUNT) FREQVAL COUNT 10
COLGROUP(L_DISCOUNT) HISTOGRAM NUMQUANTILES 20
COLGROUP(L_SUPPKEY) HISTOGRAM NUMQUANTILES 20
COLGROUP(L_SHIPDATE) FREQVAL COUNT 10
COLGROUP(L_SHIPDATE) HISTOGRAM NUMQUANTILES 20
COLGROUP(L_RECEIPTDATE) FREQVAL COUNT 10
COLGROUP(L_RETURNFLAG) FREQVAL COUNT 10
COLGROUP(L_TAX) FREQVAL COUNT 10
COLGROUP(L_RECEIPTDATE,L_RETURNFLAG,L_SHIPDATE,L_SHIPMODE)
COLGROUP(L_SHIPMODE) FREQVAL COUNT 10
COLGROUP(L_ORDERKEY,L_QUANTITY)
    
```

Generates  
RUNSTATS  
control  
statements

- 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*

Statistics Advisor report

Interesting columns:

S\_SUPPKEY  
 Cardinality: 10000.0  
 Uniform statistics collection time: 2008-09-29 16:06:48.376482  
 Uniform statistics status: OK  
 Frequency statistics collection time: 2008-09-29 16:06:48.376482  
 Frequency statistics status: OK  
 Histogram statistics collection time: null  
 Histogram statistics status: **missing**  
 Possibly point skewed: No  
 Possibly range skewed: No

S\_NATIONKEY  
 Cardinality: 25.0  
 Uniform statistics collection time: 2008-09-29 16:06:48.376482  
 Uniform statistics status: OK  
 Frequency statistics collection time: null  
 Frequency statistics status: **missing**  
 Histogram statistics collection time: null

Conflicts detail

TABLE SYSADM.LINEITEM  
 One of the frequency records (-1.0) of the L\_ORDERKEY column group is greater than the average frequency, or 1 divided by the tolerance.  
 Tolerance: 0.0010

The maximum frequency of the column group or column (L\_ORDERKEY), (0.0), is less than the average frequency, or 1 divided by the tolerance, unless only least-frequently occurring values are being collected.  
 Tolerance: 0.0010

Indicates  
conflicting and  
missing statistics

Conflicting  
statistics  
explanation



# Improve Query Design

Statements | Recommendation Summary | **Workload Query Advisor** | Workload Index Advisor | Workload Query Advis

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

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

View analysis summary

Filter recommendations by severity

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

## Guard against errors and oversights:

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

Highlights relevant components of the query

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

Recommendation and rationale

Selected Recommendation:

**Recommendation Details**

Provide a join predicate based on the referential constraint between tables DSN8910.EMP and DSN8910.DEPT.

There is a referential constraint between these two tables, but there is no join predicate that uses the foreign keys and their corresponding primary keys. Consider adding join predicates between columns EMPNO in table DSN8910.EMP and columns MGRNO in table DSN8910.DEPT to avoid fetching redundant data in the result set.

**Explanation**

If a referential constraint is defined between two tables, the queries that join the two tables generally have corresponding join predicates that map to the referential constraint.

# Visualize Queries and Costs to Speed Analysis

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

- **Accelerate analysis, reduce downtime**

- Spot human errors
- Identify where filtering should occur



## Streamlined Analysis

### Define or select a workload

1. Status: Text Sources, Input Text, File, SQL Category, Exported Workload

2. Capture: DB2 for z/OS Sources, Statement Cache, Catalog Plan or Package, QMF, QMF HPO, SQL Procedure, Plan Table, Statement Table, Function Table, View, Trigger, or SQL UDF

3. Manage: DB2 for Linux, UNIX, and Windows Sources, Package, SQL Procedure, View, Trigger, or SQL UDF

4. Invoke

5. Review

### Execute Advisors

1. Status: Single Query, Workload

2. Capture: Set Advisor Options, Run Workload Advisors, Generate Workload Table Report, Open Schedule

3. Manage: Select Query-Tuning Activities

4. Invoke:  Workload Statistics Advisor,  Workload Query Advisor,  Workload Index Advisor,  Workload Report

5. Review: Select All, Clear All

### Drill Down into advice

Advisor	Number	Priority	Description
Recommendations - 1 In			
Statistics Advisor	1	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect conflicting st...
Query Advisor	2	MEDIUM	Provide a join predicate based on the referential constraint between tables DSN8910.EM...
Query Advisor	3	MEDIUM	Provide a predicate on column WORKDEPT.
Access Path Advisor	4	LOW	Avoid reading all index keys on an index scan (QBLOOQVO = 1, PLANNO = 1).
Index Advisor	5	LOW	Index recommendations found.

### Validate improvement

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

# Optimizing programs & SQL in distributed world....

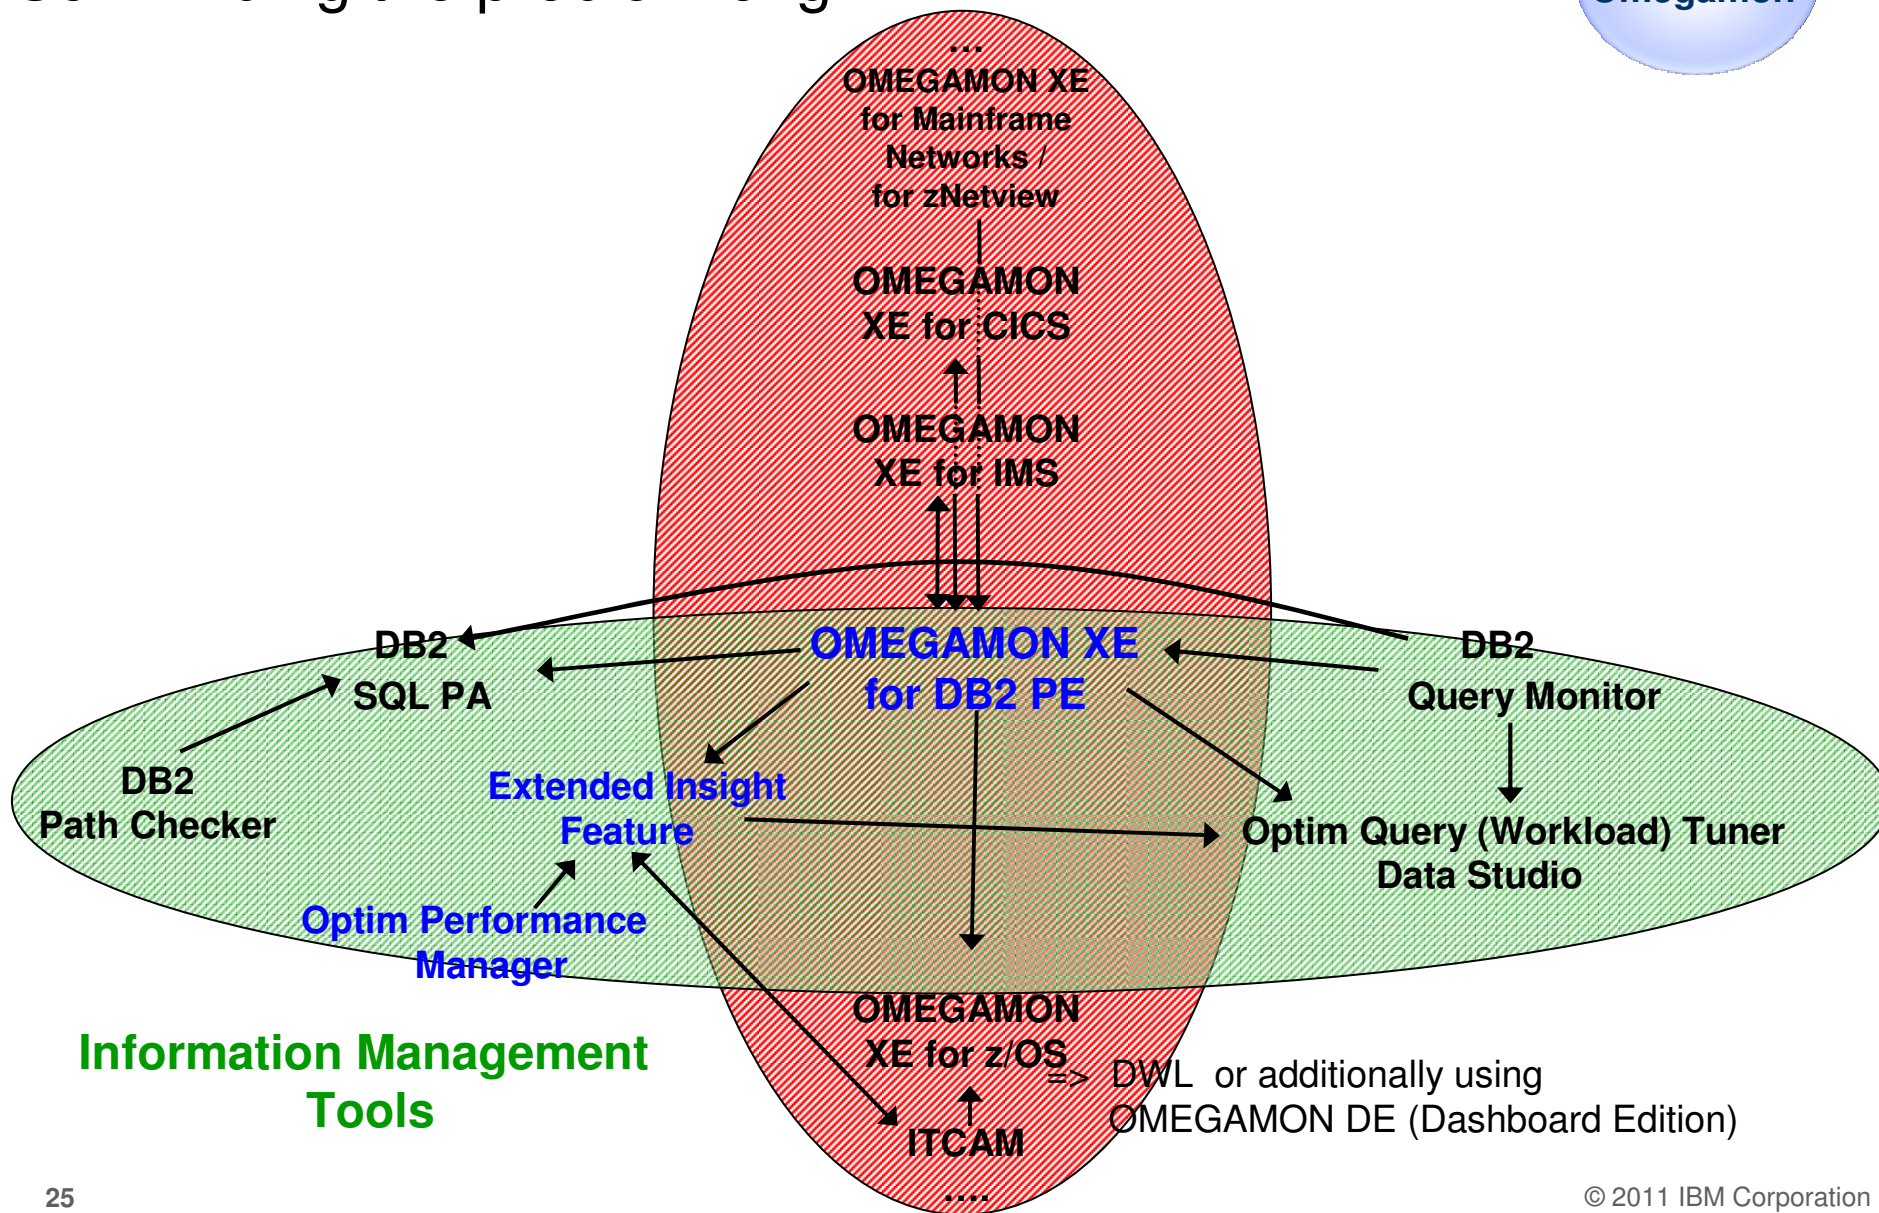
**A high-performance, data access platform to simplify developing, managing, securing, and optimizing data access.**



- **Infosphere Optim pureQuery - Simple and intuitive API**
  - Enables SQL access to databases or in-memory Java objects
  - Enables high Batching (fewer trips on the net)
  - Collects run statistics in central database
  - **Runtime option** enables flexible SQL deployment for DB2
- **Infosphere Optim Development Studio**
  - Integrated development environment with Java and SQL support
  - Improve problem isolation and impact analysis
- **Infosphere Optim Query Workload Tuner**
  - Optimize application performance and reduce costs
  - The power of advisors



So – finding the problem origin.....





# Performance Monitoring and Tuning

## OMEGAMON XE for DB2 Performance Expert on z/OS

**Norbert Jenninger**



# Omegamon for DB2 - Newest Versions

- **Version 510 - GA 10/22/2010**

- **OMEGAMON XE for DB2 Performance Expert on z/OS V510 \***
- **OMEGAMON XE for DB2 Performance Monitor on z/OS V510 \***
- **DB2 Buffer Pool Analyzer for z/OS V510**

- Key Content of these new versions

- Full “Exploitation” Support for DB2 10
- Additional customer-driven requirements (Improved Data Sharing Support)
- Reduced CPU and storage consumption
- **New Extended Insight allows monitoring an application and SQL “end-to-end”**

- (\*) We will use OMPE and OMPM as abbreviation in this presentation

## OMEGAMON V510 – complete DB2 10 support

- **Support for approximately 30 DB2 Line-items and change requests:**
  - SMF compression
  - ACCOUNTING
    - Separation of Lock and Latch wait times
    - More granularity on package level although ACCUMAC >1 is used
    - Accounting roll-up changes
    - Distributed threads accounting
  - STATISTICS
    - IFCID 225 (memory) changes
    - Multiple IFCID 2 for each 25 buffer pools (>25 buffer pool usage)
    - DSC enhancements including static SQL
    - EDM Pool and other working memory moved above the bar
  - New and updated ZPARMs
  - PERFORMANCE traces
  - Audit trace changes
    - Row-level and Column-level access control
    - New DBA privileges



**DB2 10 beta:** More than twenty DB2 10 beta customers have downloaded and used OM PE V510

**OMEGAMON beta:** OM PE V510 beta customers tested OM PE with DB2 V8, 9, and 10

## Latest News:

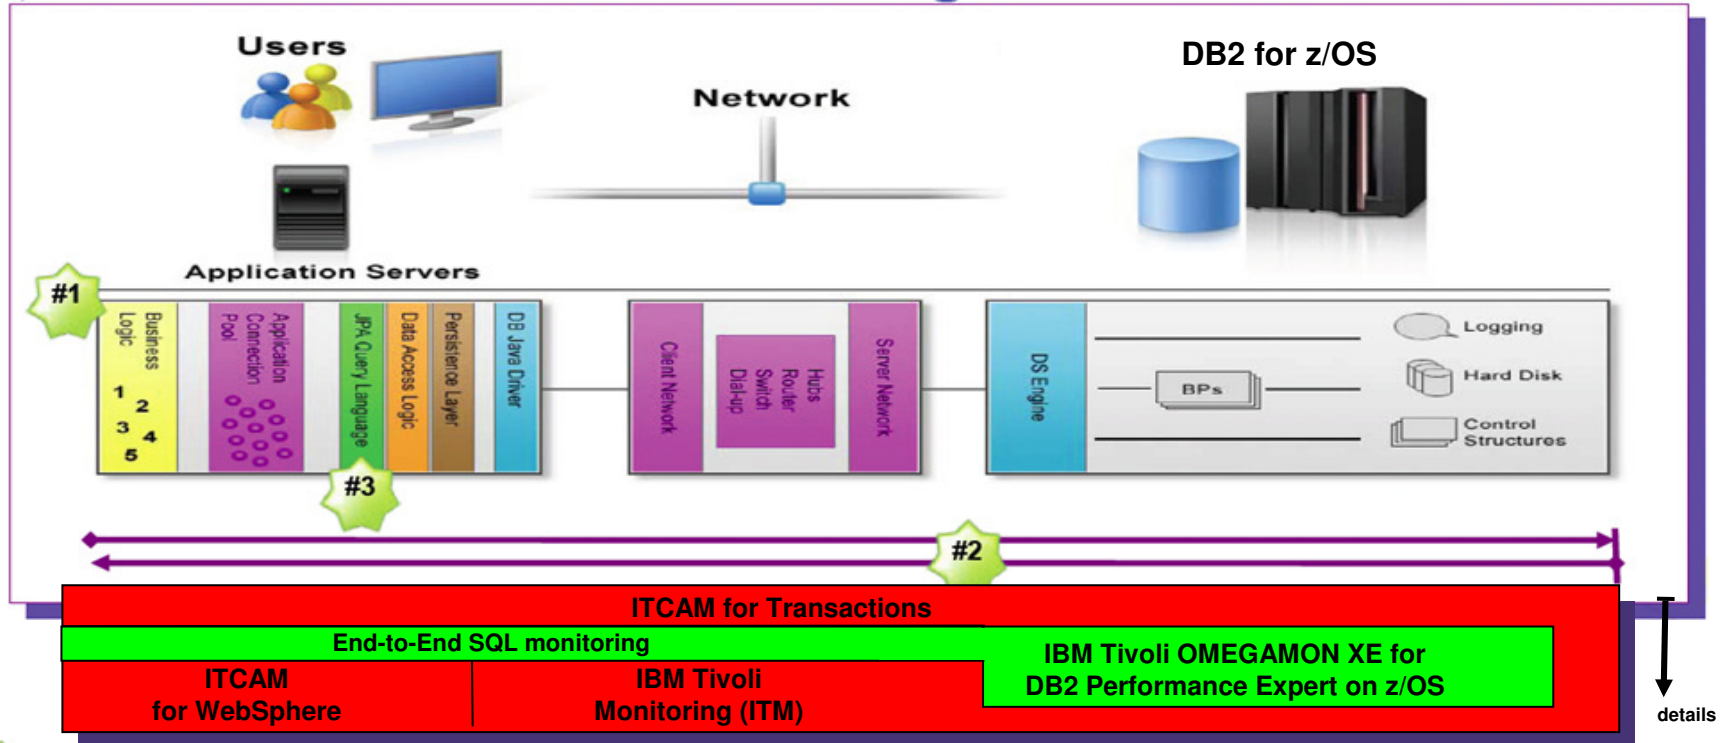
### New Function PTF UK65924 / APAR PM35049

- **Near-Term History now with zIIP off-load**  
plus improved precision in NTH CPU filter
- **ATF display new DB2 10 metrics for CLOSE CURSOR**
- **Reporting**
  - SQL Activity Trace shows new DB2 10 metrics for CLOSE CURSOR  
Static/Dynamic, Sync buff reads, Getpages, Indexscan, TS scan, rows exam., rows proc., sorts, buffer writes, ....
  - Audit TRACE and REPORT supports new IFCID 271. Data related to objects ROW PERMISSION and COLUMN MASK are shown
  - Record Trace show new V10 (after GA DCR) instrumentation
  - New System Parameter
- **Performance DB (INTEGER to DECIMAL, extend fields to avoid overflow)**
- **New PE client driver (Display Static SQL cache)**
- **plus couple of internal fixes**
- ...



# So - where is my problem?

=> End-to-end monitoring with SQL level deep dive - Extended Insight



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

# DB2 Performance metrics (Trace records – Trace types)

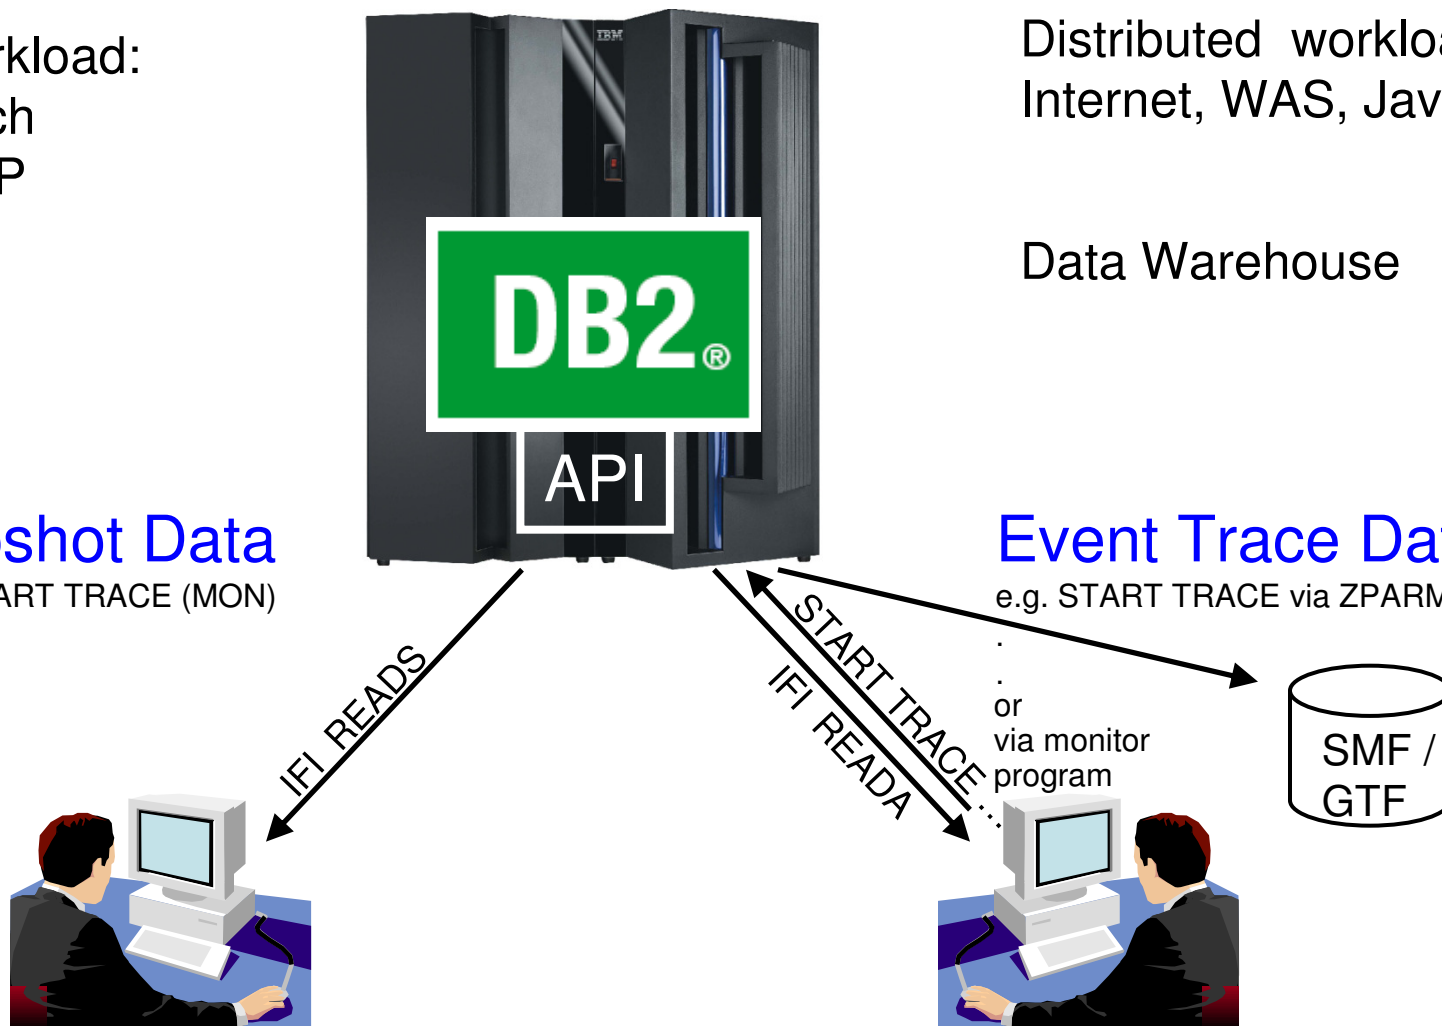
Legacy workload:  
 Cobol, Batch  
 CICS, OLTP

Distributed workload:  
 Internet, WAS, Java

Data Warehouse

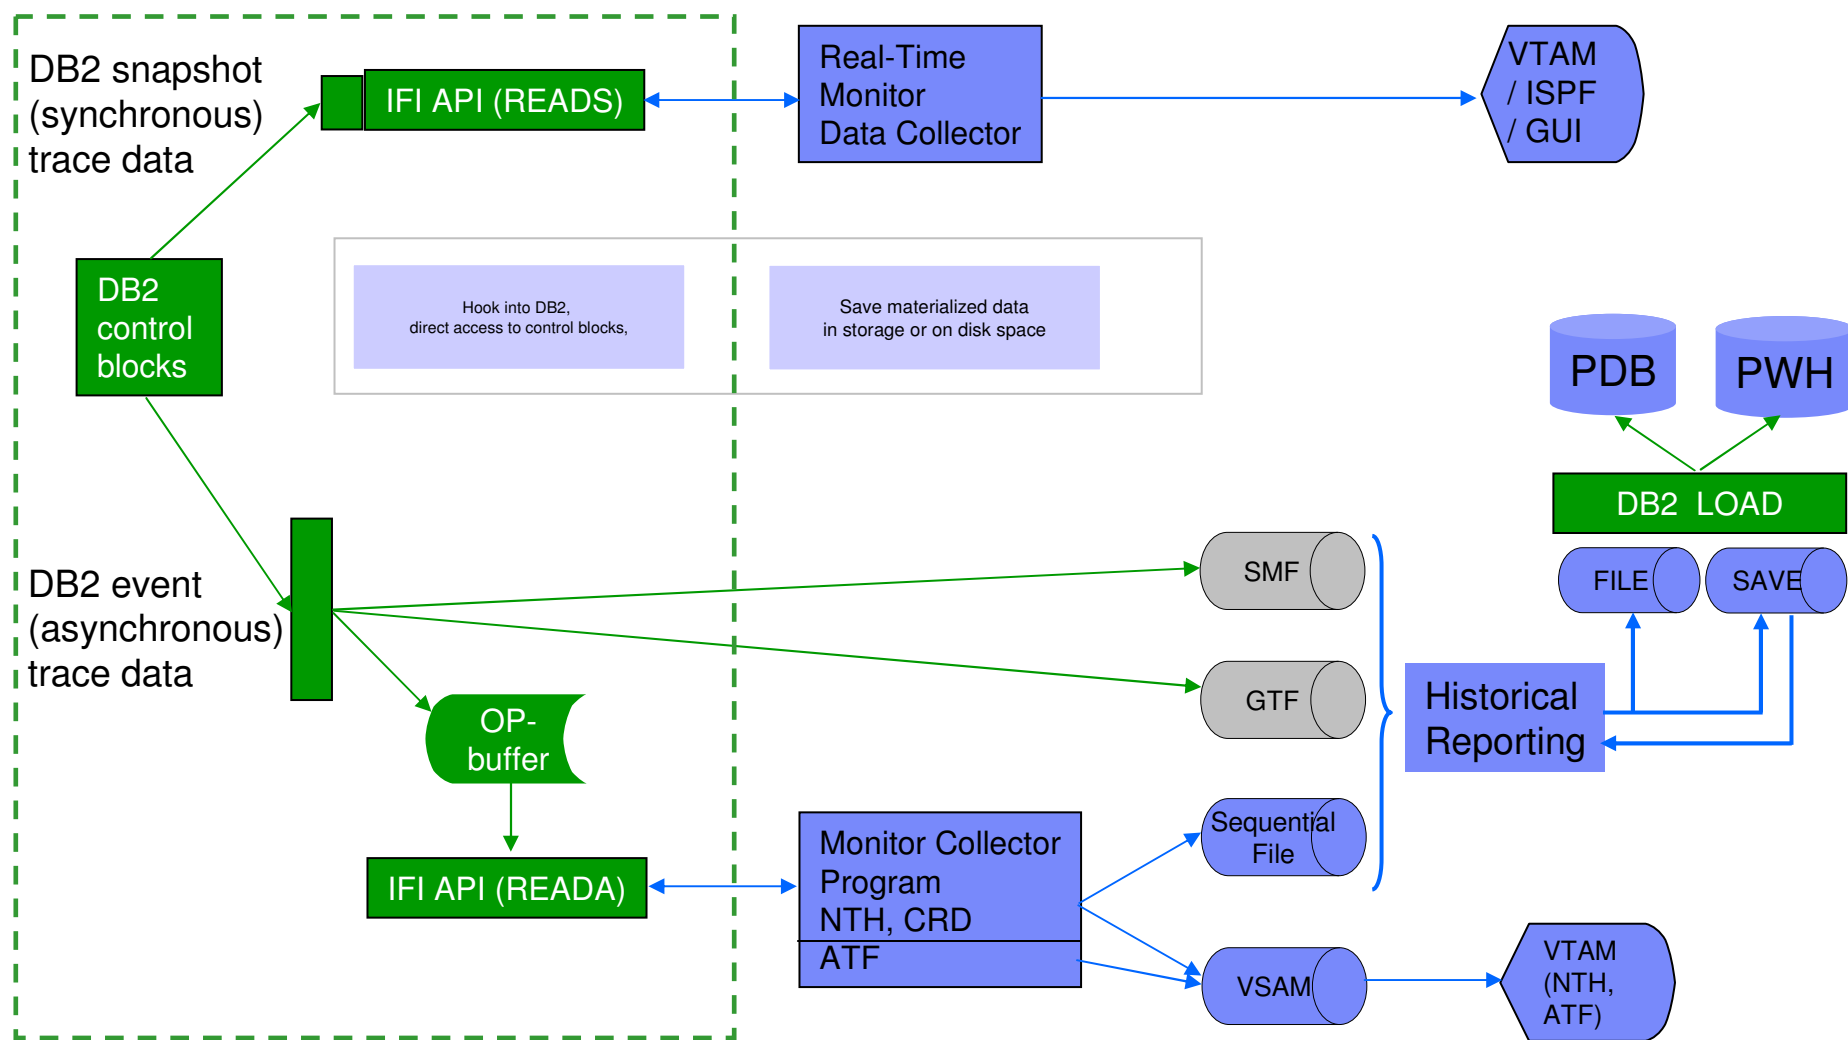
**Snapshot Data**  
 START TRACE (MON)

**Event Trace Data**  
 e.g. START TRACE via ZPARMS  
 .  
 .  
 or via monitor program

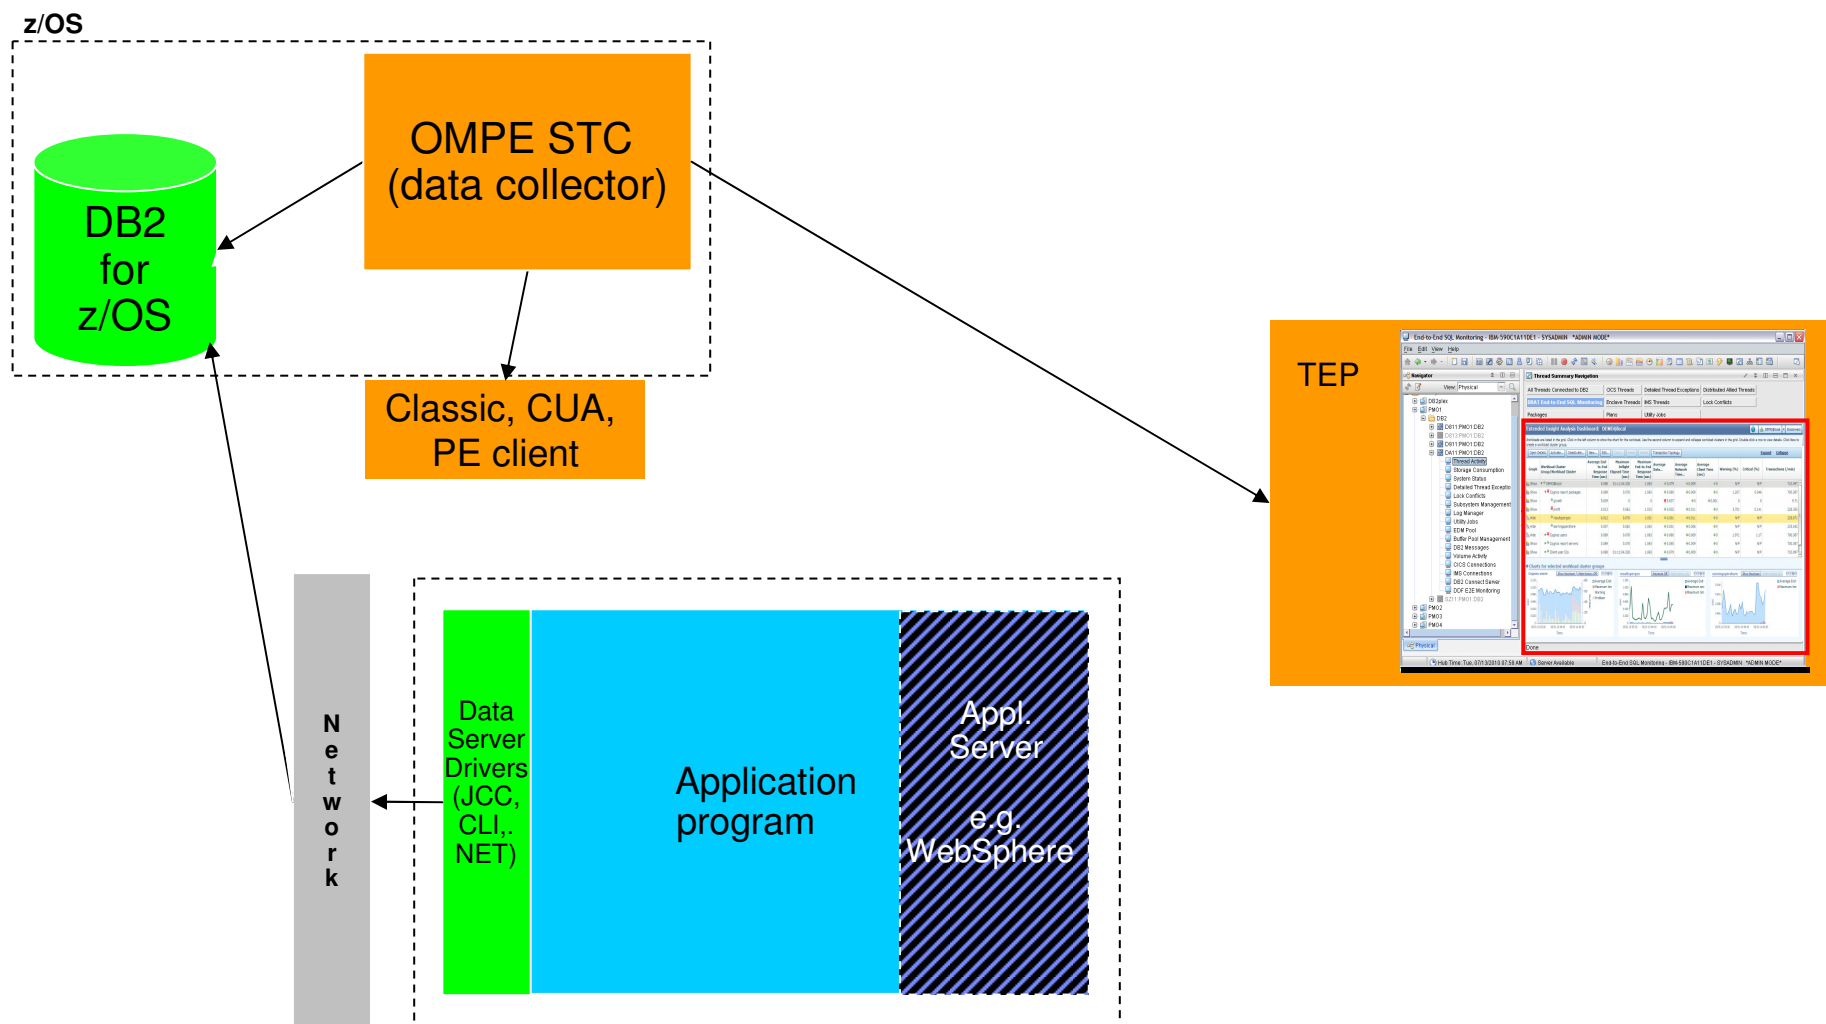




# DB2 Metrics / Instrumentation data and their usage



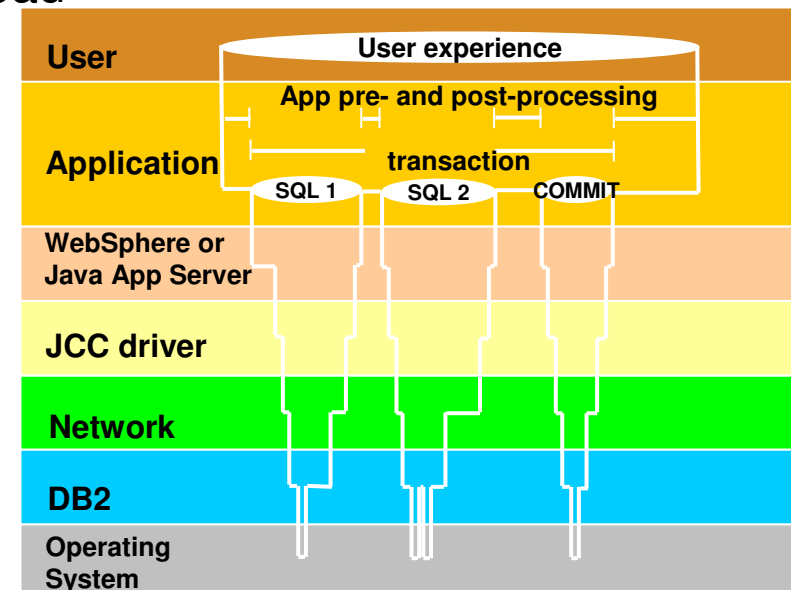
# Still one major pain point - What is happening outside of DB2? (e.g. with remote applications)



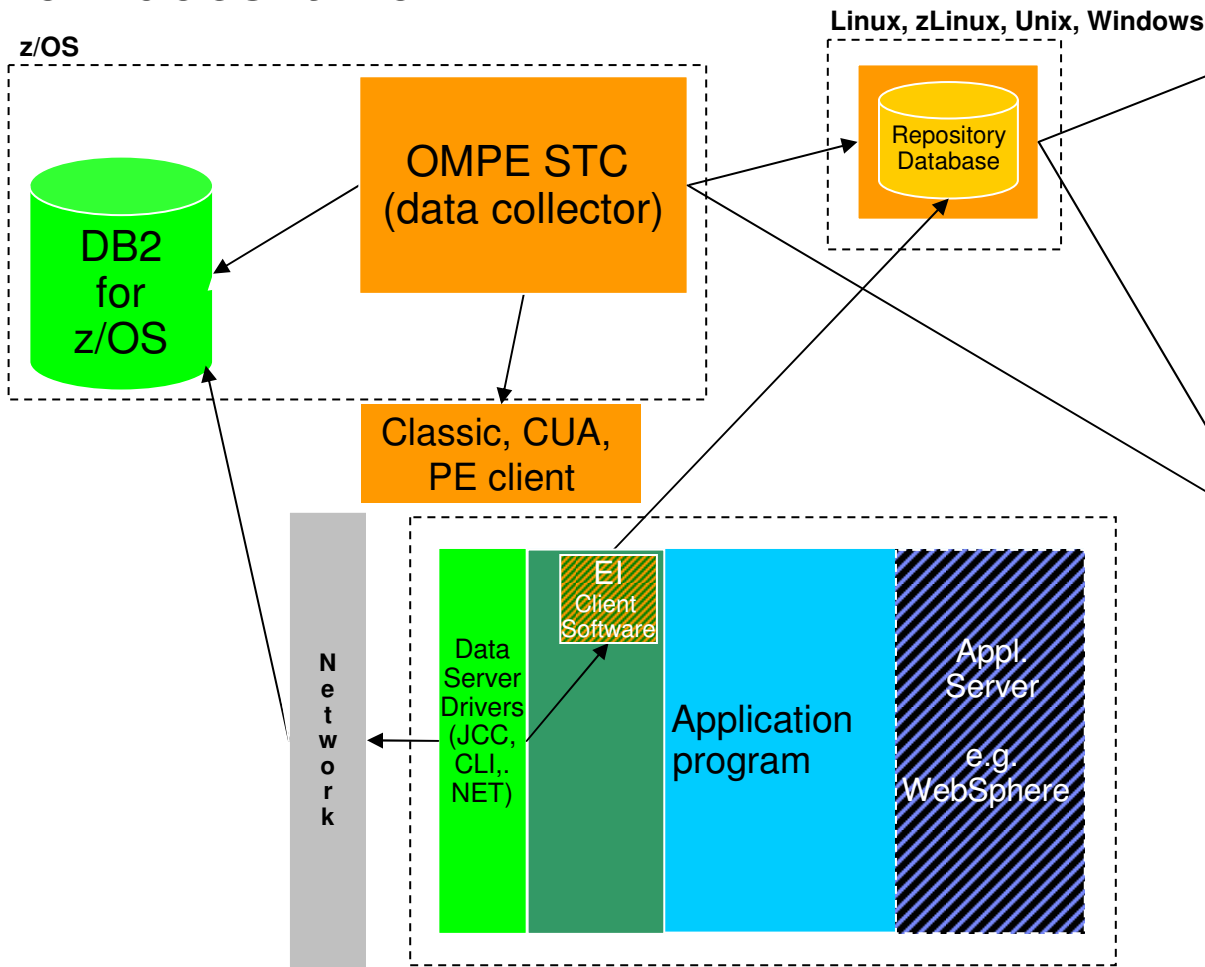
## Where is my DB2 application spending its time?

- **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

**Extended Insight** is available with the Performance **Expert** Offering only

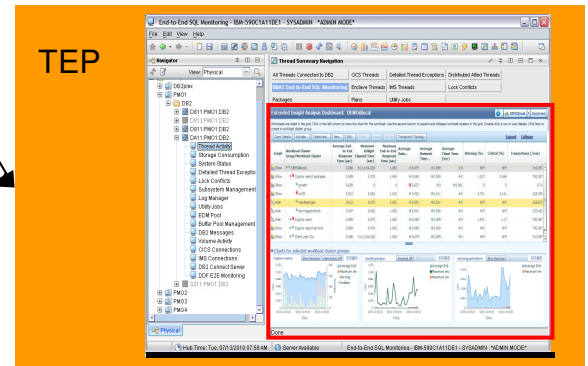
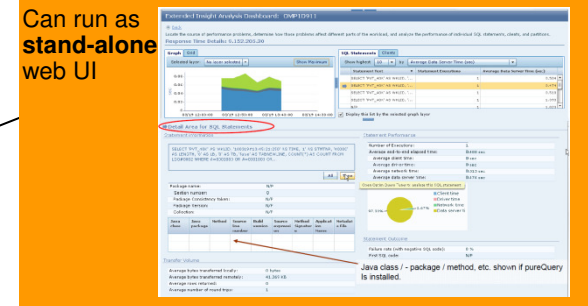


# How does it work?



Color coding:

- = OMEGAMON XE for DB2 PE V510
- = DB2 for z/OS and DB2 Data Server drivers



## 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

The screenshot shows the DBAT End-to-End SQL Monitoring application interface. The 'Navigator' pane on the left shows a tree view for 'DA11:PMO1:DB2' with 'Thread Activity' selected. The 'Thread Summary Navigation' pane contains several buttons, with 'DBAT End-to-End SQL Monitoring' circled in red. Below this, two bar charts show 'DDF DB Access Thread CPU%' for 'SUK'. The bottom pane, 'Distributed Database Access Thread Connection Summary', contains a table with columns for MAXDBAT, Current DBAT, DBAT Utilization, MAXDBAT HWM, CONDBAT, DBAT Connection, DBAT Conn Utilization, and DBAT Conn HWM. The 'End-to-End SQL Monitoring' row in this table is circled in red, with an arrow pointing to a text box.

Navigation to the OMPE workspace with the E2E SQL monitoring information

MAXDBAT	Current DBAT	DBAT Utilization	MAXDBAT HWM	CONDBAT	DBAT Connection	DBAT Conn Utilization	DBAT Conn HWM
				2	10000		

Time	Plan	DBRM (Unicode)	CP CPU Rate	Thread Status	Workstation ID	User ID	Transaction ID	Authorization ID (Unicode)	Conversations	SQL Received	Rows Sent	
00:01:07.1	DISTSERV	SQLC2F0A	0.0	IN-DB2	IBM-590C1A11DE1	suk	db2bp.exe	SUK	0	6	572	-DA
00:00:08.3	DISTSERV	SQLC2F0A	0.0	IN-DB2	IBM-590C1A11DE1	suk	db2bp.exe	SUK	0	6	572	-DA



# OMPE on z/OS TEP

## Navigation to the Extended Insight Analysis Dashboard

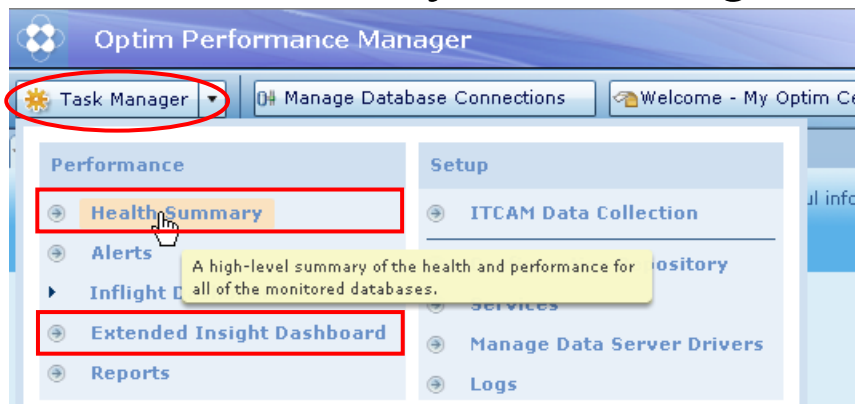
The screenshot displays the 'End-to-End SQL Monitoring - IBM-590C1A11DE1 - SYSADMIN \*ADMIN MODE\*' application window. The interface is divided into several sections:

- Navigator:** A tree view on the left showing the system hierarchy: DB2plex > PMO1 > DB2 > D811:PMO1:DB2, D813:PMO1:DB2, D911:PMO1:DB2, DA11:PMO1:DB2. The 'Thread Activity' option is selected under the DB2 nodes.
- Thread Summary Navigation:** A set of buttons at the top right for navigating between different thread-related views: All Threads Connected to DB2, CICS Threads, Detailed Thread Exceptions, Distributed Allied Threads, DBAT End-to-End SQL Monitoring (selected), Enclave Threads, IMS Threads, Lock Conflicts, Packages, Plans, and Utility Jobs.
- Extended Insight Analysis Dashboard: DEMO@local:** A central dashboard with a table of workload clusters and three charts below it.
 

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...	Average Network Time...	Average Client Time (sec)	Warning (%)	Critical (%)	Transactions (/min)
Show	DEMO@local	0.088	01:11:04.328	1.063	0.079	0.009	0	N/P	N/P	710.097
Show	Cognos report packages	0.089	0.078	1.063	0.080	0.009	0	1.207	0.046	700.387
Show	growth	5.639	0	0	5.637	0	0.001	0	0	9.71
Show	profit	0.013	0.062	1.015	0.002	0.011	0	3.701	0.141	228.355
Hide	resultspergeo	0.012	0.078	1.031	0.001	0.011	0	N/P	N/P	228.871
Hide	earningspershare	0.007	0.062	1.063	0.001	0.006	0	N/P	N/P	233.452
Hide	Cognos users	0.089	0.078	1.063	0.080	0.009	0	1.971	1.17	700.387
Show	Cognos report servers	0.089	0.078	1.063	0.080	0.009	0	N/P	N/P	700.387
Show	Client user IDs	0.088	01:11:04.328	1.063	0.079	0.009	0	N/P	N/P	710.097

Below the table, there are three charts for selected workload cluster groups: 'Cognos users', 'resultspergeo', and 'earningspershare'. Each chart shows 'Average End-End' (blue line), 'Maximum tim' (orange bars), 'Warning' (yellow bars), and 'Problem' (red bars) over time.
- Status Bar:** At the bottom, it shows 'Hub Time: Tue, 07/13/2010 07:58 AM', 'Server Available', and the application title.

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



Alert Severity: All

- Critical Alerts (3)
- Warning Alerts (0)
- Critical and Warning Alerts (3)

Data Source	Alerts		System										Database			Extended Insight				
	Critical	Warning	Data Server status	Monitoring Status	Memory Usage	CPU Utilization (%)	Connections	Storage	Recovery	Workload	Logging	I/O	Sorting	Locking	Transactions (/min)	Physical Page I/O (/min)	Lock Wait Time (sec)	Longest Running SQL (sec)	Average End-to-End Response Time (sec)	Maximum End-to-End Response Time (sec)
OMPD911	1	0	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	0	0	0s	0s	1.120s	4.625s
OMPDA11	1	0	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	0	0	0s	0s	0s	0s
SN51	1	0	OK	OK	0	0	OK	OK	OK	OK	OK	OK	OK	OK	0	0	0s	0s	9.692s	26.281s

3 total items

10 Items per page Page 1 of 1

Navigate to Extended Insight Analysis Dashboard for E2E SQL monitoring =>



# Extended Insight Analysis Dashboard

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

Optim Performance Manager | TSCHAFFL | Log out | About | ?

Task Manager | Manage Database Connections | Welcome - My Optim Central

Welcome - My Optim Central | Manage Database Connections | Health Summary | Workload | System | Overview | **Extended Insight Dashboard**

Learn about the time controls. | 03/17/10 16:00 | 03/26 09:06 - 03/26 12:06 | GMT +01:00 | End Time: 03/26/10 12:06 | Duration: 3 Hours

Extended Insight Analysis Dashboard: OMP1D911 | OMP1D911 | Disconnect

Workloads are listed in the grid. Click in the left column to show the chart for the workload. Use the second column to expand and collapse workload clusters in the grid. Double-click a row to view details. Click New to create a workload cluster group.

Open Details | Activate... | Deactivate... | New... | Edit... | Copy... | Reset | Delete | Transaction Topology | Expand | Collapse

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...	Average Client Time (sec)	Warning (%)	Critical (%)	Transactions (/min)
Sh...	▼ Host names/IP address	0.588	0	1.828	◆ 0.465	◆ 0.002	◆ 0.120	N/P	N/P	6
Sh...	◆ 9.152.205.30	0.588	0	1.828	◆ 0.465	◆ 0.002	◆ 0.120	N/P	N/P	6
Sh...	▼ Application Types	0.588	0	1.828	◆ 0.465	◆ 0.002	◆ 0.120	N/P	N/P	6
Sh...	◆ OTHER	0.588	0	1.828	◆ 0.465	◆ 0.002	◆ 0.120	N/P	N/P	6
Sh...	◆ WebSphere Applicati	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P
Sh...	◆ WebSphere applicati	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P	N/P

Charts for selected workload cluster groups

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

# Extended Insight Analysis Dashboard

Optim Performance Manager

Task Manager | Manage Database Connections | Welcome

Welcome - My Optim Central | Manage Database Connections | Learn about...

Recent | History | 34 sec | Aggregation level:1

03/17/10 16:00

### Extended Insight Analysis Dashboard: OMP1D91

Workloads are listed in the grid. Click in the left column to show the chart for the workload cluster group. Click New to create a workload cluster group.

Open Details | Activate... | Deactivate... | New... | Edit... | Copy

Graph	Workload Cluster Group/Workload Cluster	Average End-to-End Response Time (sec)	Maximum Inflight Elapsed Time (sec)
Sh...	Host names/IP address	0.588	0
Sh...	9.152.205.30	0.588	0
Sh...	Application Types	0.588	0
Sh...	OTHER	0.588	0
Sh...	WebSphere Application Servers	N/P	N/P
Sh...	WebSphere applications	N/P	N/P

Charts for selected workload cluster groups

### Activate Workload Cluster Groups

Select the workload cluster groups to activate. Only activated workload cluster groups are monitored and shown on the performance monitoring dialogs.

| Edit...

Workload Cluster Group	Description
<input type="checkbox"/> E2E_SAP_SAP_REPORT_ZOS_CL...	E2E_SAP_SAP_REPORT_ZOS_CL_GROUP_DESCRIPTION
<input type="checkbox"/> E2E_SAP_SAP_TRANSACTIONS_Z...	E2E_SAP_SAP_TRANSACTIONS_ZOS_CL_GROUP_DESCRIPTION
<input type="checkbox"/> SAP application servers	Contains a workload cluster for each SAP application server t...
<input type="checkbox"/> SAP users	Contains a workload cluster for each SAP end user that send...
<input type="checkbox"/> SQW application servers	Shows the response time of each InfoSphere Warehouse ap...
<input type="checkbox"/> SQW applications and flows	Shows the InfoSphere Warehouse applications and flows acc...
<input type="checkbox"/> DataStage jobs	Contains a workload cluster for each DataStage job that sen...
<input type="checkbox"/> DataStage servers	Contains a workload cluster for each DataStage server that ...
<input type="checkbox"/> Cognos users	Contains a workload cluster for each Cognos user that send...
<input type="checkbox"/> Cognos report packages	Contains a workload cluster for each Cognos report package...
<input type="checkbox"/> Cognos report servers	Contains a workload cluster for each Cognos report server t...
<input type="checkbox"/> WebSphere Application Servers	Contains a workload cluster for each WebSphere Application...
<input type="checkbox"/> WebSphere applications	Contains a workload cluster for each WebSphere application...
<input type="checkbox"/> Authentication IDs	Contains a workload cluster for each authorization ID that s...
<input type="checkbox"/> Client workstations	Contains a workload cluster for each workstation that sends ...

OK Cancel

Define your "Workload cluster groups" as you need

# Extended Insight Analysis Dashboard

### Activate Workload Cluster Groups

Select the workload cluster groups to activate. Only activated workload cluster groups are monitored and shown on the performance monitoring dialogs.

Edit...

#### Edit Workload Cluster Group SAP Transaction/Batch Job

**Step 2 of 3**

A workload cluster group can cover the entire workload of a database or only part of it. Specify connection attributes to generate workload clusters for this group. You can select one more attributes for clustering. You can also select connection attribute to reduce the workload that is covered. Click Browse (...) to view the available filter values. Click Refresh to generate the workload clusters.

**Connection Attributes and Filter Criteria**

Type of workload cluster group: SAP on System Z Sampling period: Current time s

Cluster by Connection Attribute	Filter the Workload		
<input type="checkbox"/> SAP Application Server	<input type="checkbox"/> SAP Application Server	=	
<input type="checkbox"/> SAP End User	<input type="checkbox"/> SAP End User	=	
<input type="checkbox"/> SAP Report/IDoc Message/BW Query	<input type="checkbox"/> SAP Report/IDoc Mes...	=	
<input checked="" type="checkbox"/> SAP Transaction/Batch Job	<input type="checkbox"/> SAP Transaction/Batc...	=	

**Generated Workload Clusters**

Workload clusters: 5 Transactions executed: 16.593 Refresh

Reset Cluster Name

Workload Cluster Name	SAP Application Server	SAP End User	SAP Report/IDoc Message/BW Query
session_manager	*	*	*
sap_collector_for_pe	*	*	*
sm12	*	*	*
	*	*	*
sap_rsicfdlt	*	*	*

#### Edit Workload Cluster Group SAP Transaction/Batch Job

**Step 3 of 3**

You can specify response time thresholds for the entire workload cluster group or for individual workload clusters. When the workload cluster group is activated for monitoring, you are informed if thresholds were violated. The thresholds that you set for workload clusters on the Extended Insight Dashboard do not apply to other dashboards.

Do not use default values. Specific thresholds can be entered in the table below.

Use default thresholds for all workload clusters(in addition,specify thresholds can be entered in the table).

Warning threshold: None

Critical threshold: None

Sampling period: < 03/16 08:49:01 - 03/16 09:49:01 >

Workload Cluster Name	Average End-to-End Response	Maximum End-to-End Response	Warning Thresholds	Critical Thresholds
session_manager	0.506	0.810	None	None
sap_collector_for_pe	0.280	4.425	None	None
sm12	0.240	1.771	None	None
--	0.065	7.460	None	None
sap_rsicfdlt	0.052	0.090	None	None

Click in the threshold column to apply a specific value to a workload cluster.

<Back Next> Finish Cancel

<Back Next> Finish Cancel

© 2011 IBM Corporation

## Setting Client information - Samples

```

DB2ConnectionPoolDataSource ds = new DB2ConnectionPoolDataSource();
ds.setUser("myuser");
ds.setPassword("mypass");
ds.setDatabaseName("mydb");
ds.setConnectionAttribute
  ("ClientAppName=WebSphere-Samples;
   ClientWrkstnName=WebSphere-Wkstn;
   ClientUserid=WebSphere-User
   ClientAcctStr=WebSphere-Acctstr");
PooledConnection poolconn = ds.getPooledConnection();
con = poolconn.getConnection();
  
```

WhitePaper: <http://www.ibm.com/developerworks/data/library/techarticle/0212shayer/0212shayer.html>

```

▶▶ CALL DSNRLI(—function, correlation-id, —accounting-token, accounting-interval)
└─, —retcode
  └─, —reascod
    └─, user
      └─, —appl
        └─, ws
          └─, —xid
  
```



## Reported on the DB2 for z/OS

Accounting report, for example

ACCOUNTING

REPORT FROM (04-01-13 ,13:42) TO (04-01-13 ,13:50)

ORDER (ENDUSER-WSNAME-TRANSACT)

INCLUDE (WSNAME(WORKSTATNAME))

EXEC

or in a DB2 “DISPLAY THREAD ....” command output

```

DSNV401I  -SGI2 DISPLAY THREAD REPORT FOLLOWS -
DSNV402I  -SGI2 ACTIVE THREADS -
NAME      ST A  REQ ID          AUTHID  PLAN      ASID  TOKEN
SERVER    RA *  663 DSNVE.exe    JEN     DISTSERV 0089  353
V437-WORKSTATION=JENNINGE, USERID=JEN,
APPLICATION NAME=DSNVE.exe
V445-G998C474.KA08.040112202345=353 ACCESSING DATA FOR 9.152.196.116
SERVER    RA *  44 db2bp.exe    JEN     DISTSERV 0089  351
V437-WORKSTATION=JENNINGE, USERID=jen,
APPLICATION NAME=db2bp.exe
V445-G998C474.E207.040112191547=351 ACCESSING DATA FOR 9.152.196.116

```

# Extended Insight Analysis Dashboard

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

## Extended Insight Analysis Dashboard: COH1 SN81

Workloads are listed in the grid. Click in the left column to show the chart for the workload. Use the second column to expand and collapse workload clusters in the group.

Graph	Workload Cluster Group/Workload Cluster	Average End-to-End Response Time	Maximum Inflight Elapsed Time	Maximum End-to-End Response Time	Average Data Server Time	Average Network Time	Average Client Time
Sh...	COH1 SN81	0.071	0.203	7.460	0.019	0.006	0.045
Sh...	SAP Transaction/Bal	0.071	0.203	7.460	0.019	0.006	0.045
Sh...	SAP Report/IDoc Me	0.071	0.203	7.460	0.019	0.006	0.045
Sh...	SAP application serv	0.071	0.203	7.460	0.019	0.006	0.045
Sh...	Application Types	0.071	0.203	7.460	0.019	0.006	0.045

Graph	Workload Cluster Group/Workload Cluster	Average End-to-End Response Time	Maximum Inflight Elapsed Time	Maximum End-to-End Response Time	Average Data Server Time	Average Network Time	Average Client Time	Transactions (/min)	St
Sh...	COH1 SN81	0.071	0.203	7.460	0.020	0.006	0.045	15.776	
Sh...	SAP Transaction/Bal	0.071	0.203	7.460	0.020	0.006	0.045	15.776	
Sh...	session_manager	0.506	0	0.810	0.387	0.078	0.041	3	
Sh...	sap_collector_for	0.280	0	4.425	0.072	0.019	0.188	21	
Sh...	sm12	0.275	0	1.771	0.122	0.130	0.024	1.182	
Sh...		0.062	0.203	7.460	0.016	0.004	0.042	15.103	
Sh...	sap_rsicfdlt	0.052	0	0.090	0.021	0.012	0.019	2	
Sh...	SAP Report/IDoc Me	0.071	0.203	7.460	0.020	0.006	0.045	15.776	
Sh...	rsstatph	0.493	0	0.493	0.010	0.473	0.010	1	
Sh...	saplscsm_downti	0.461	0	0.645	0.050	0.011	0.400	0.571	

# 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	Maximum Inflight Elapsed Time	Maximum End-to-End Response Time	Average Data Server Time	Average Network Time	Average Client Time
Sh...	▼ SAP application serv	0.071	0	7.460	◆ 0.019	◆ 0.007	◆ 0.046
Sh...	◆ 9.152.20.157	0.092	0	7.460	◆ 0.020	◆ 0.003	◆ 0.068
Sh...	◆ 9.152.20.159	0.069	0	5.095	◆ 0.022	◆ 0.012	◆ 0.036
Sh...	◆ 9.152.20.158	0.040	0	1.829	◆ 0.010	◆ 0.001	◆ 0.028

Show	▼ Client application name	9.692	35.875	26.281	◆ 0.943	◆ 0.034	◆ 8.715
Show	◆ testapplication_  2	12.421	11.671	21.234	◆ 0.845	◆ 0.043	◆ 11.532
Show	◆ testapplication_  5	11.128	20.407	26.281	◆ 0.979	◆ 0.034	◆ 10.114
Show	◆ testapplication_  4	10.641	12.343	21.344	◆ 0.815	◆ 0.033	◆ 9.792
Show	◆ testapplication_  3	9.601	10.375	24.219	◆ 0.916	◆ 0.031	◆ 8.654
Show	◆ testapplication_  0	9.232	35.875	15.375	◆ 0.991	◆ 0.033	◆ 8.208
Show	◆ testapplication_  8	8.381	4.328	14.703	◆ 1.024	◆ 0.037	◆ 7.320
Show	◆ testapplication_  7	8.258	6.422	18.703	◆ 0.993	◆ 0.028	◆ 7.236
Show	◆ testapplication_  6	8.095	6.843	22.672	◆ 0.985	◆ 0.029	◆ 7.082
Show	◆ testapplication_  1	8.050	6.656	16.828	◆ 0.991	◆ 0.031	◆ 7.028
Show	▼ Client user IDs	9.692	35.875	26.281	◆ 0.943	◆ 0.034	◆ 8.715
Show	◆ testuser_  0	10.601	12.343	24.219	◆ 0.933	◆ 0.035	◆ 9.633



# Extended Insight Analysis Dashboard

## Show additional graphs for selected workload clusters

Optim Performance Manager

jen | [Log out](#) | [About](#) | ?

Task Manager | Manage Database Connections | Welcome - My Optim Central

Welcome - My Optim Central | Manage Database Connections | **Extended Insight Dashboard** | Health Summary

Learn about the time controls.

Recent | History | Refresh

08/23/10 02:37 03:41 04:30 05:19 06:08 06:56 07:45 08:34 09:23 10:12 11:01 11:50 12:39 13:28 14:17 15:06 15:55 16:44 08/23/10 17:48

Aggregation level:1

08/23/10 15:56 - 08/23/10 16:16 GMT +02:00 End Time: 08/23/10 16:16 Duration: 20 Mins

Extended Insight Analysis Dashboard: SN51

Workloads are listed in the grid. Click in the left column to show the chart for the workload. Use the second column to expand and collapse workload clusters in the grid. Double-click a row to view details. Click New to create a workload cluster group.

Open Details | Activate... | Deactivate... | New... | Edit... | Copy... | Reset | Delete | View All Known Clients | Transaction Topology

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	9.501	10.156	21.234	0.961	0.032	8.508	--	--	3.095	0
Show	Client user IDs	9.501	10.156	21.234	0.961	0.032	8.508	--	--	3.095	0
Hide	Client application name	9.501	10.156	21.234	0.961	0.032	8.508	--	--	3.095	0
Show	testapplication_I 2	11.375	10.156	21.234	0.904	0.050	10.421	--	--	0.381	0
Show	testapplication_I 1	11.078	0	16.828	0.890	0.031	10.157	--	--	0.333	0
Show	testapplication_I 5	10.848	0	21.234	1.063	0.032	9.754	--	--	0.333	0
Hide	testapplication_I 3	10.483	8.531	14.484	0.838	0.030	9.614	--	--	0.429	0
Hide	testapplication_I 4	10.367	7.437	17.000	0.907	0.026	9.434	--	--	0.381	0
Show	testapplication_I 0	10.323	8.125	12.875	0.982	0.040	9.301	--	--	0.286	0
Show	testapplication_I 6	8.027	0	14.172	1.037	0.033	6.957	--	--	0.333	0
Show	testapplication_I 8	7.333	4.328	13.937	0.963	0.025	6.344	--	--	0.333	0
Show	testapplication_I 7	4.393	0	9.516	1.148	0.019	3.227	--	--	0.286	0

Charts for selected workload cluster groups

Client application nan

testapplication\_I 3

testapplication\_I 4

# Extended Insight Analysis Dashboard

## Show response time histogram for selected workload

Optim Performance Manager TSCHAFFL | Log out | About | ?

Task Manager | Manage Database Connections | Welcome - My Opti

Host names/IP addresses

Response Time

Dashboard

GMT +01:00  
 03/26 12:18  
 03/26/10  
 12:18  
 Duration:  
 4 Hours

OMP1D911 | Disconnect

double-click a row to view details.

Expand Collapse

Graph Workload Cluster Group/Workload Average End-to-End Response Time Maximum Inflight Elapsed... Maximum End-to-End Response Time

Graph	Workload Cluster Group/Workload	Average End-to-End Response Time	Maximum Inflight Elapsed...	Maximum End-to-End Response Time
Sh...	OMP1D911	0.650	0	2.000
Hide	Host names	0.650	0	2.000
Sh...	Application	0.650	0	2.000
Sh...	WebSphere	N/P	N/P	N/P
Sh...	WebSphere	N/P	N/P	N/P

Overall average response time per transaction: 0.650  
 Maximum response time: 2.094  
 Maximum Time of running transactions: 0  
 Warning threshold: None  
 Critical threshold: None

Alert History

Response Time Histogram

Go to the workload cluster details graphs.

OK Apply Cancel

Click to show response time histogram

Charts for selected workload cluster groups

Host names/IP addresses | Show Maximum | Alert History On

Average End-to-End Response Time  
 Maximum time of running transaction

sec

Time

# Extended Insight Analysis Dashboard

## Zoom into selected workload and see the TOP SQL list

Optim Performance Manager TSCHAFFL | Log out | About | ?

Task Manager | Manage Database Connections | Welcome - My Optim Central

Welcome - My Optim Central | Manage Database Connections | Health Summary | Workload | System | Overview | **Extended Insight Dashboard**

### Extended Insight Analysis Dashboard: OMP1D911

[Back](#)

Locate the source of performance problems, determine how those problems affect different parts of the workload, and analyze the performance of individual SQL statements, clients, and partitions.

**Response Time Details: 9.152.205.30**

**Graph** | Grid

Selected layer: Average End-to-End Response Time Show Maximum

**SQL Statements** | Clients

Show highest 10 by Average Data Server Time (sec)

Statement Text	Statement Executions	Average Data Server Time (sec)
SELECT 'PVT_40K' AS WKLID, '...	1	0.504
SELECT 'PVT_40K' AS WKLID, '...	1	0.474
SELECT 'PVT_40K' AS WKLID, '...	1	0.518
SELECT 'PVT_40K' AS WKLID, '...	1	1.393
N/P	1	1.023

Display this list by the selected graph layer

**Detail Area for Average End-to-End Response Time**

**End-to-End Response Time**

Overall average response time per transaction:	0.075 sec
Maximum response time:	15.282 sec
Maximum Time of running transactions	10.688 sec
Number of transactions:	61,245
Statements:	65,344

**Time Distribution (%)**

**Transaction Throughput**

**Statement Throughput**

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

Extended Insight Analysis Dashboard: OMP1D911

[Back](#)

Locate the source of performance problems, determine how those problems affect different parts of the workload, and analyze the performance of individual SQL statements, clients, and partitions.

Response Time Details: 9.152.205.30

**Graph** | Grid

Selected layer: No layer selected Show Maximum

**SQL Statements** | Clients

Show highest 10 by Average Data Server Time (sec)

Statement Text	Statement Executions	Average Data Server Time (sec)
SELECT 'PVT_40K' AS WKLID, '...	1	0.504
SELECT 'PVT_40K' AS WKLID, '...	1	0.474
SELECT 'PVT_40K' AS WKLID, '...	1	0.518
SELECT 'PVT_40K' AS WKLID, '...	1	1.393
N/P	1	1.023

Display this list by the selected graph layer

**Detail Area for SQL Statements**

Statement information

```
SELECT 'PVT_40K' AS WKLID, '100319#13:45:21:250' AS TIME, '1' AS STMTNR, '40000' AS LENGTH, '0' AS LB, '0' AS TB, 'false' AS TABNEWLINE, COUNT(*) AS COUNT FROM LGQ#0002 WHERE A=0001000 OR A=0001000 OR...
```

All Type

Package name:	N/P
Section number:	0
Package Consistency token:	N/P
Package Version:	N/P
Collection:	N/P

Java class	Java package	Method	Source line number	Build version	Source expression	Method Signature	Application Name	Metadata File

Transfer Volume

Average bytes transferred locally:	0 bytes
Average bytes transferred remotely:	41,369 KB
Average rows returned:	0
Average number of round trips:	1

Statement Performance

Number of Executions:	1
Average end-to-end elapsed time:	0.488 sec
Average client time:	0 sec
Average driver time:	0 sec
Average network time:	0.013 sec
Average data server time:	0.474 sec

Open Optim Query Tuner to analyze this SQL statement.

Legend: Client time (blue), Driver time (red), Network time (green), Data server time (yellow)

Statement Outcome

Failure rate (with negative SQL code):	0 %
First SQL code:	N/P

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

# Extended Insight Analysis Dashboard

## Select SQL from list and zoom into SQL level details

Extended Insight Analysis Dashboard: OMP1D911

Back

Locate the source of performance problems, determine how those problems affect different parts of the workload, and analyze the performance of individual SQL statements, clients, and partitions.

Response Time Details: 9.152.205.30

Graph Grid

Selected layer: No layer selected Show Maximum

SQL Statements Clients

Show highest 10 by Average Data Server Time (sec)

Statement Text	Statement Executions	Average Data Server Time (sec)
SELECT 'PVT_40K' AS WKLD, '...	1	0.504
SELECT 'PVT_40K' AS WKLD, '...	1	0.474
SELECT 'PVT_40K' AS WKLD, '...	1	0.518
SELECT 'PVT_40K' AS WKLD, '...	1	1.393
N/P	1	1.023

Display this list by the selected graph layer

**Detail Area for SQL Statements**

Statement information

```
SELECT 'PVT_40K' AS WKLD, '100319#13:45:21:250' AS TIME, '1' AS STMTNR, '40000' AS LENGTH, '0' AS LB, '0' AS TB, 'false' AS TABNEWLINE, COUNT(*) AS COUNT FROM LGQ#0002 WHERE A=0001000 OR A=0001000 OR...
```

All Type

Package name: N/P  
Section number: 0  
Package Consistency token: N/P  
Package Version: N/P  
Collection: N/P

Java class	Java package	Method	Source line number	Build version	Source expression	Method Signature	Application Name	Metadata File

Transfer Volume

Average bytes transferred locally: 0 bytes  
Average bytes transferred remotely: 41.369 KB  
Average rows returned: 0  
Average number of round trips: 1

Statement Performance

Number of Executions: 1  
Average end-to-end elapsed time: 0.488 sec  
Average client time: 0 sec  
Average driver time: 0 sec  
Average network time: 0.013 sec  
Average data server time: 0.474 sec

Open Optim Query Tuner to analyze this SQL statement.

Statement Outcome

Failure rate (with negative SQL code): 0 %  
First SQL code: N/P

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 ...

```

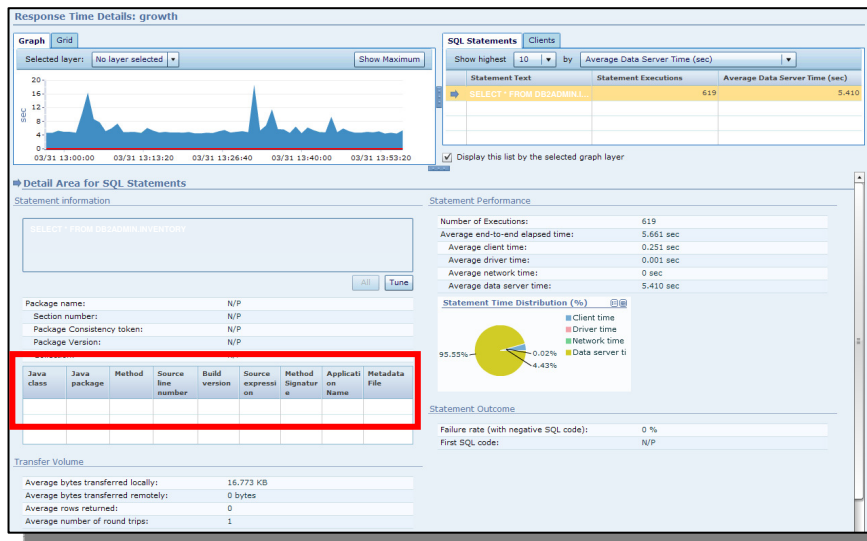
Application source
public class TestOPM {
    public static void main(String [] args) throws Exception{
        String url = "jdbc:db2://sv1-imtestg12.svl.ibm.com:50000/SAMPLE";
        Connection con = SampleUtil.getConnection(url, "db2admin", "hot6cold");
        ((com.ibm.db2.jcc.DB2Connection) con).setDB2ClientApplicationInformation("blah");

        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

Transfer Application Metadata

Choose to transfer the metadata to a repository or to save the metadata to the file system.

Destination:

- Transfer to a metadata repository (selected)
- Save to file system

Application Information:

Runtime group name: blah

Version: blahVer

OK Cancel

# How pureQuery Runtime can help !

## Extended Insight feature can tell you where the query came from ...

Application source

```

public class TestOPM {
    public static void main(String [] args) throws Exception{
        String url = "jdbc:db2://sv1-intestg12.svl.ibm.com:50000/SAMPLE";
        Connection con = SampleUtil.getConnection(url, "db2admin", "not6cold");
        ((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
                
```

Java class	Java package	Method	Source line number	Build version	Source expression	Method Signature	Application Name	Metadata File
TestOPM	my.test	main	13	blahVer	N/P	N/P	blah	capture...

Area for SQL Statements

Statement: SELECT \* FROM DB2ADMIN.INVENTORY

Statement class: com.ibm.db2.jcc.DB2Statement

Statement type: PreparedStatement

Statement SQL: SELECT \* FROM DB2ADMIN.INVENTORY

Statement SQL code: SELECT \* FROM DB2ADMIN.INVENTORY

Statement class: com.ibm.db2.jcc.DB2Statement

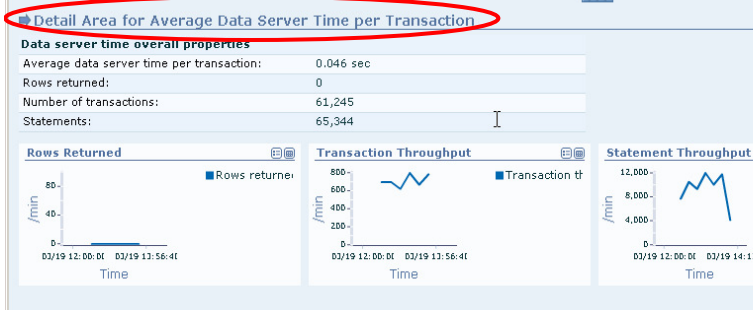
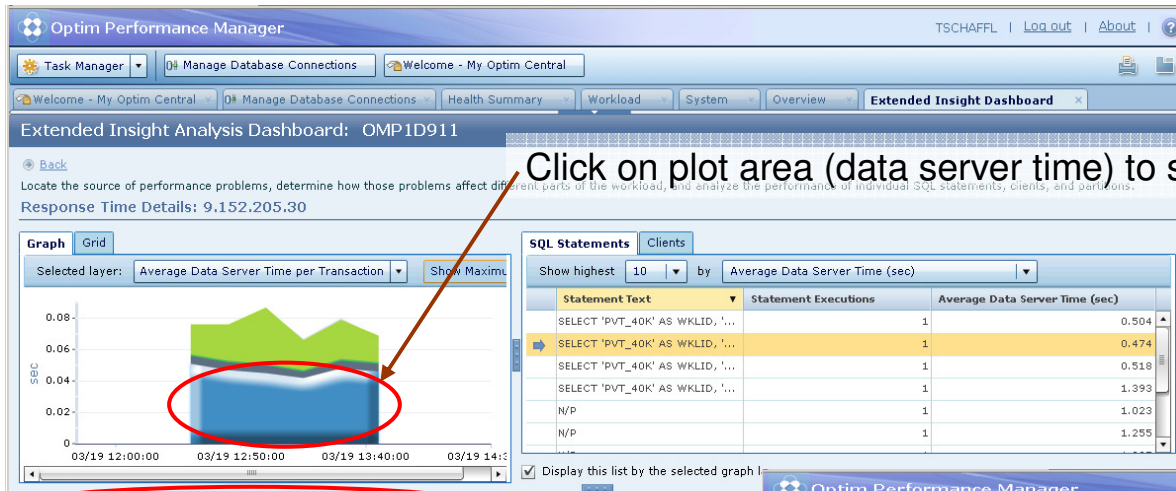
Statement type: PreparedStatement

Statement SQL: SELECT \* FROM DB2ADMIN.INVENTORY


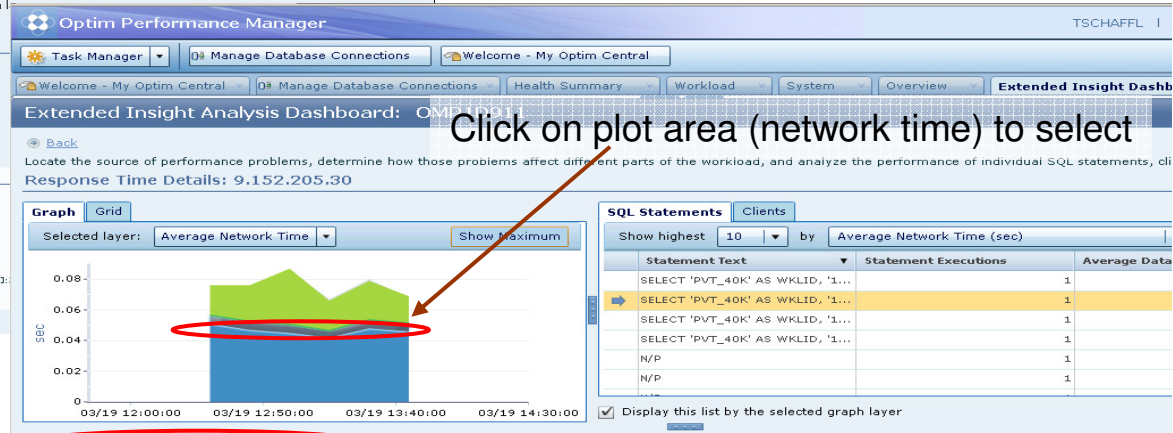
Statement SQL code: SELECT \* FROM DB2ADMIN.INVENTORY

# Extended Insight Analysis Dashboard

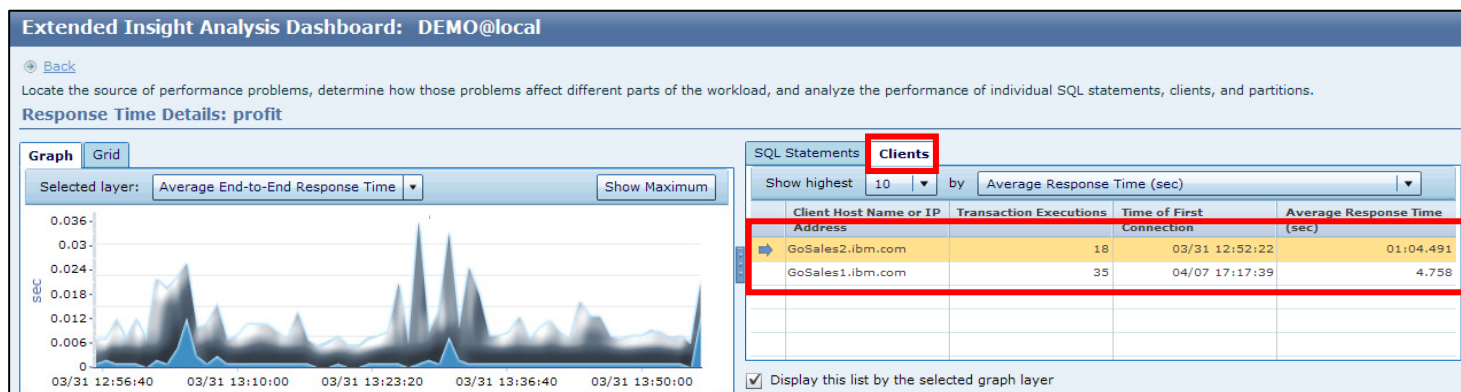
## Select a plot area and see the corresponding performance metrics



... for network times and throughput



# WebSphere – another area to be monitored in context



pureQuery level: 2.15.14  
 JRE vendor: Sun Microsystems Inc.  
 JRE version: 16.2-b04  
 JVM properties: java.runtime.name=Java(TM) SE Runtime E...

WebSphere Application Server data source name: GSDB  
 WebSphere Application Server server name: GoSales2  
 WebSphere Application Server version: 7.0.1

WAS Connection Pool

Connection pool size: 50  
 Average connections in use: 17  
 Maximum connection wait time:

Pool Usage

Client Time Distribution (%)

Client Comparison

**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	25.33%		100		16.2-b04	Window...	3.58.82

# Where is my DB application spending its time? e.g. in application, SQL, and network?

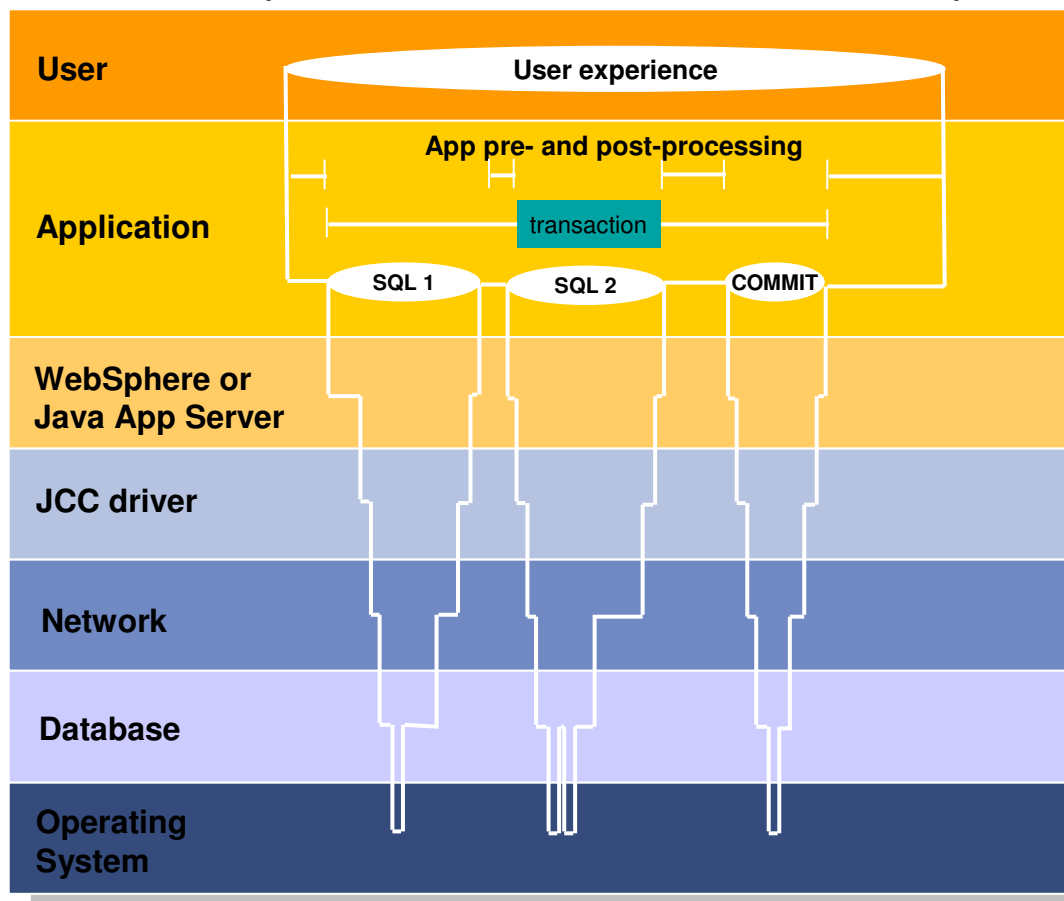
## ITCAM for Transactions

- **Deep** on end-to-end transaction path
- **Hand-off** to Optim, ITCAM, ITM for diagnose / repair

Drill down to  
Performance Manager

Drill down to  
ITCAM WR / ITM

- OM DB2 PE with Extended Insight**
- **Deep** on Database, JCC driver, etc
  - **Shallow** on network, OS, app server for sensing problems then **hand-off** to ITM for Servers to diagnose / repair



## ITCAM for WebSphere

- **Deep** on application / web server, J2EE / web app, web resources
- **Shallow** on JCC driver, database app for sensing problems, then **hand-off** to Optim Performance Manager for diagnose / repair

## ITM for Servers

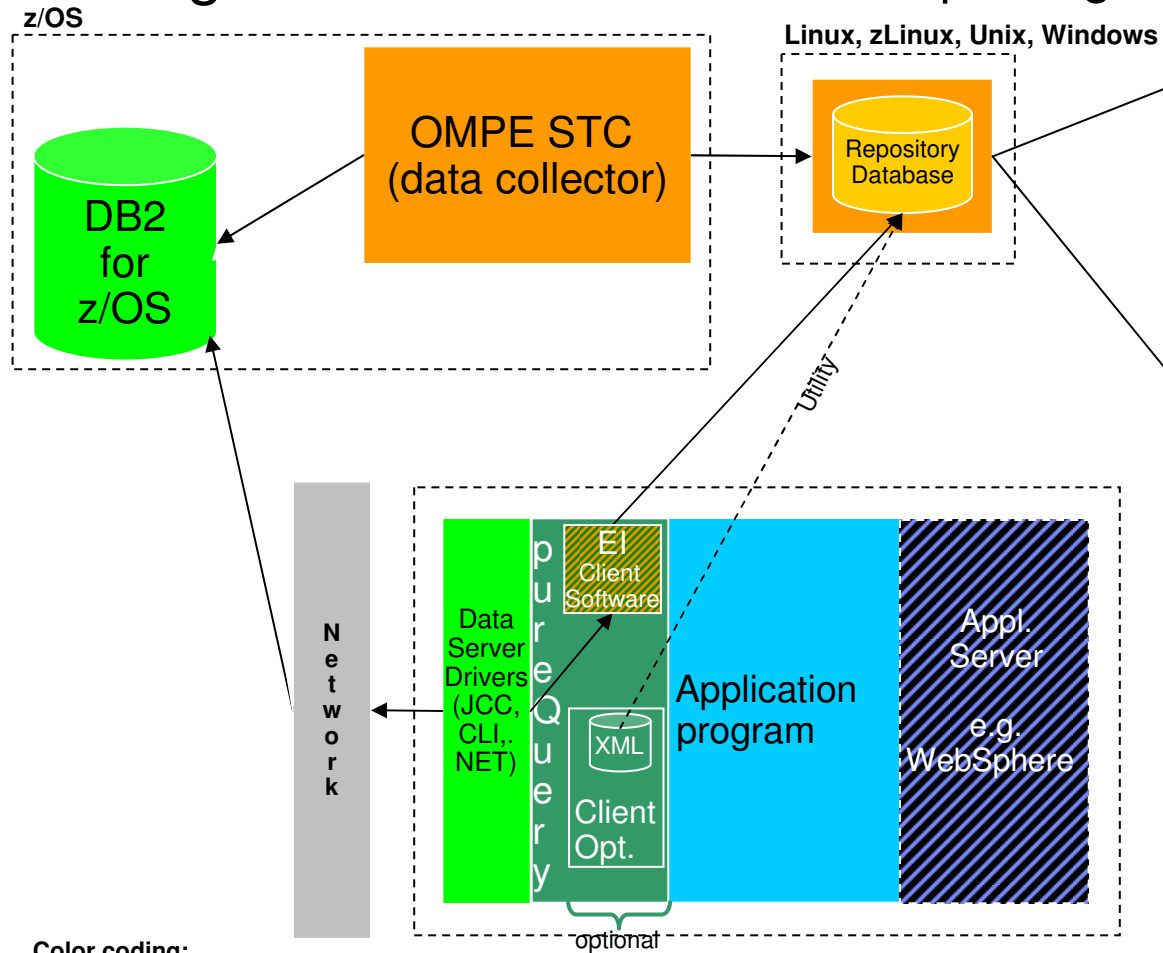
- **Deep** on OS, Network, Storage, I/O, etc
- **Shallow** on database for sensing problems then **hand-off** to Optim Performance Manager for diagnose / repair

Drill down to ITCAM WR / ITM

Drill down to Optim Performance Manager



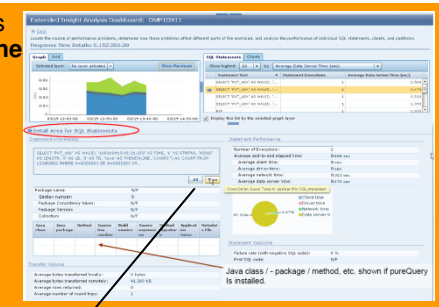
# Omegamon for DB2 Architecture Summary: An integrated solution and how it is packaged



Color coding:

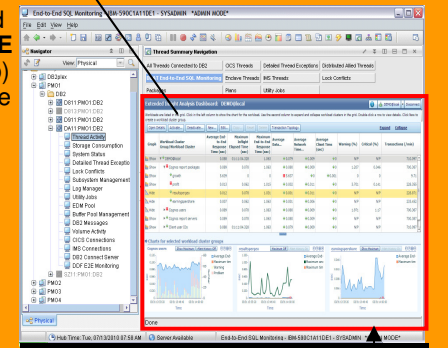
- = OMEGAMON XE for DB2 PE V510
- = DB2 for z/OS and DB2 Data Server drivers
- = Tools (Optim pureQuery runtime and Optim Query [Workload] Tuner)
- = IBM Tivoli Composite Application Monitoring

Can run as **stand-alone web UI**

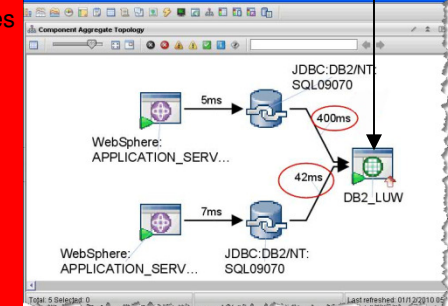


**Optim Query (Workload) Tuner**

Integrated into **OMPE TEP (web) workspace**



Integrates with **ITCAM for TT**



## Extended Insight Feature Summary

- **Advanced way to monitor the database workload (SQL) of your applications and solutions**
  - Get response times and time breakdown (appl, driver, network, data server) per defined workload/cluster, e.g. per system, per application, per 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
  - Zoom into the various layers
  
- **Optional integration (- advantages) with**
  - Optim Query (Workload) Tuner / Data Studio
  - pureQuery (Runtime while using Data capturing)
  - ITCAM for WebSphere applications accessing DB2 via JDBC



Thank You !



© 2011 IBM Corporation



# Backup































# Resources

- pureQuery for DB2 for z/OS video
  - **Improving ROI for existing applications**
- <http://www.ibm.com/developerworks/offers/lp/demos/summary/im-purequery4zos.html>
  
- IBM InfoSphere Optim Development Studio  
<http://www-01.ibm.com/software/data/optim/development-studio/>
  
- IBM InfoSphere Optim pureQuery Runtime web page  
[www.ibm.com/software/data/studio/purequery/](http://www.ibm.com/software/data/studio/purequery/)
  
- Articles and tutorials covering the InfoSphere Optim portfolio  
<http://www.ibm.com/developerworks/spaces/optim>
  
- DB2 Connect Advanced edition  
[https://www.ibm.com/developerworks/mydeveloperworks/blogs/idm/entry/db2\\_connect\\_advanced\\_edition1?lang=en\\_us](https://www.ibm.com/developerworks/mydeveloperworks/blogs/idm/entry/db2_connect_advanced_edition1?lang=en_us)



## Contents of Eclipse-based Query Tuning offerings for z/OS

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

## OM PE V510 customer-driven requirements



- Cancel remote threads
- Identify CPU utilization for remote threads
- Report on DSN Activity for remote threads
- See DB2 Connect Server details for a distributed thread originating on a remote LPAR
- See statement text for static SQL in Application Trace
- Support for SQL/PA V410
- Launch “explain” tools: Optim Query Workload Tuner as well as Data Studio from OM PE

## Improvements in OMPE, e.g. Reduced Overhead

- **V510 shows strong improvements compared to V420**
  - Moved more storage areas above the bar, resulting in relief below the bar (→ 31bit private and ECSA)
  - Better management of background processing (code path reductions, better stack implementations)
  - Reduced overhead - using extended filtering and qualification (Classic and Tivoli Enterprise Portal)
  - Reduced the number of internal TCBs to lower private storage usage (→ Tivoli Enterprise Monitoring Agent)
  - New “out of the box VTAM profile” (basic monitoring to get started)

# Extended Insight Analysis Dashboard

## Page down to review the host Dynamic SQL statement cache metrics

Detail Area for SQL Statements

General Information | **Statement Server Execution Details**

Most Recent Identification

Package name:	--
Statement Type:	--
Package Version:	--
Cache Insert time stamp:	07/28 10:37:30
Last execution:	--

Most Recent Compilation

Isolation level:	CS
First referenced table:	--
Authorization ID:	JEN
CURRENT SQL ID:	JEN
Client user ID:	
Client workstation name:	--
Client accounting string:	--
Object qualifier:	JEN
Literal replacement:	--
CURRENTDATA bind option:	N
DYNAMICRULE bind option:	R
CURRENT RULES:	D
CURRENT PRECISION:	N
Cursor WITH HOLD:	Y

Data Server Execution Time

Number of transactions:	1,029
Average execution time:	0 sec
Average CPU time:	0 sec

Overall Time Distribution

Transaction Logging Statistics

Average log writer wait time:	0 sec
-------------------------------	-------

Locking Statistics

Average local lock wait time:	0 sec
Average global lock wait time:	0 sec
Average latch request wait time:	0 sec
Average page latch wait time:	0 sec
Average drain lock wait time:	0 sec
Average drain claim release wait time:	0 sec
Average read wait time for another thread:	0 sec
Average write wait time for another thread:	0 sec

Statement Row and Sort Details

Average rows read:	1
Average rows returned or modified:	1
Average index scans:	0
Average table space scans:	1
RID usage failures due to internal DB2 limits:	
Average RID usage failures due to RID list storage:	0
Average Number of Parallel Groups:	0
Total sorts:	0

Row Efficiency

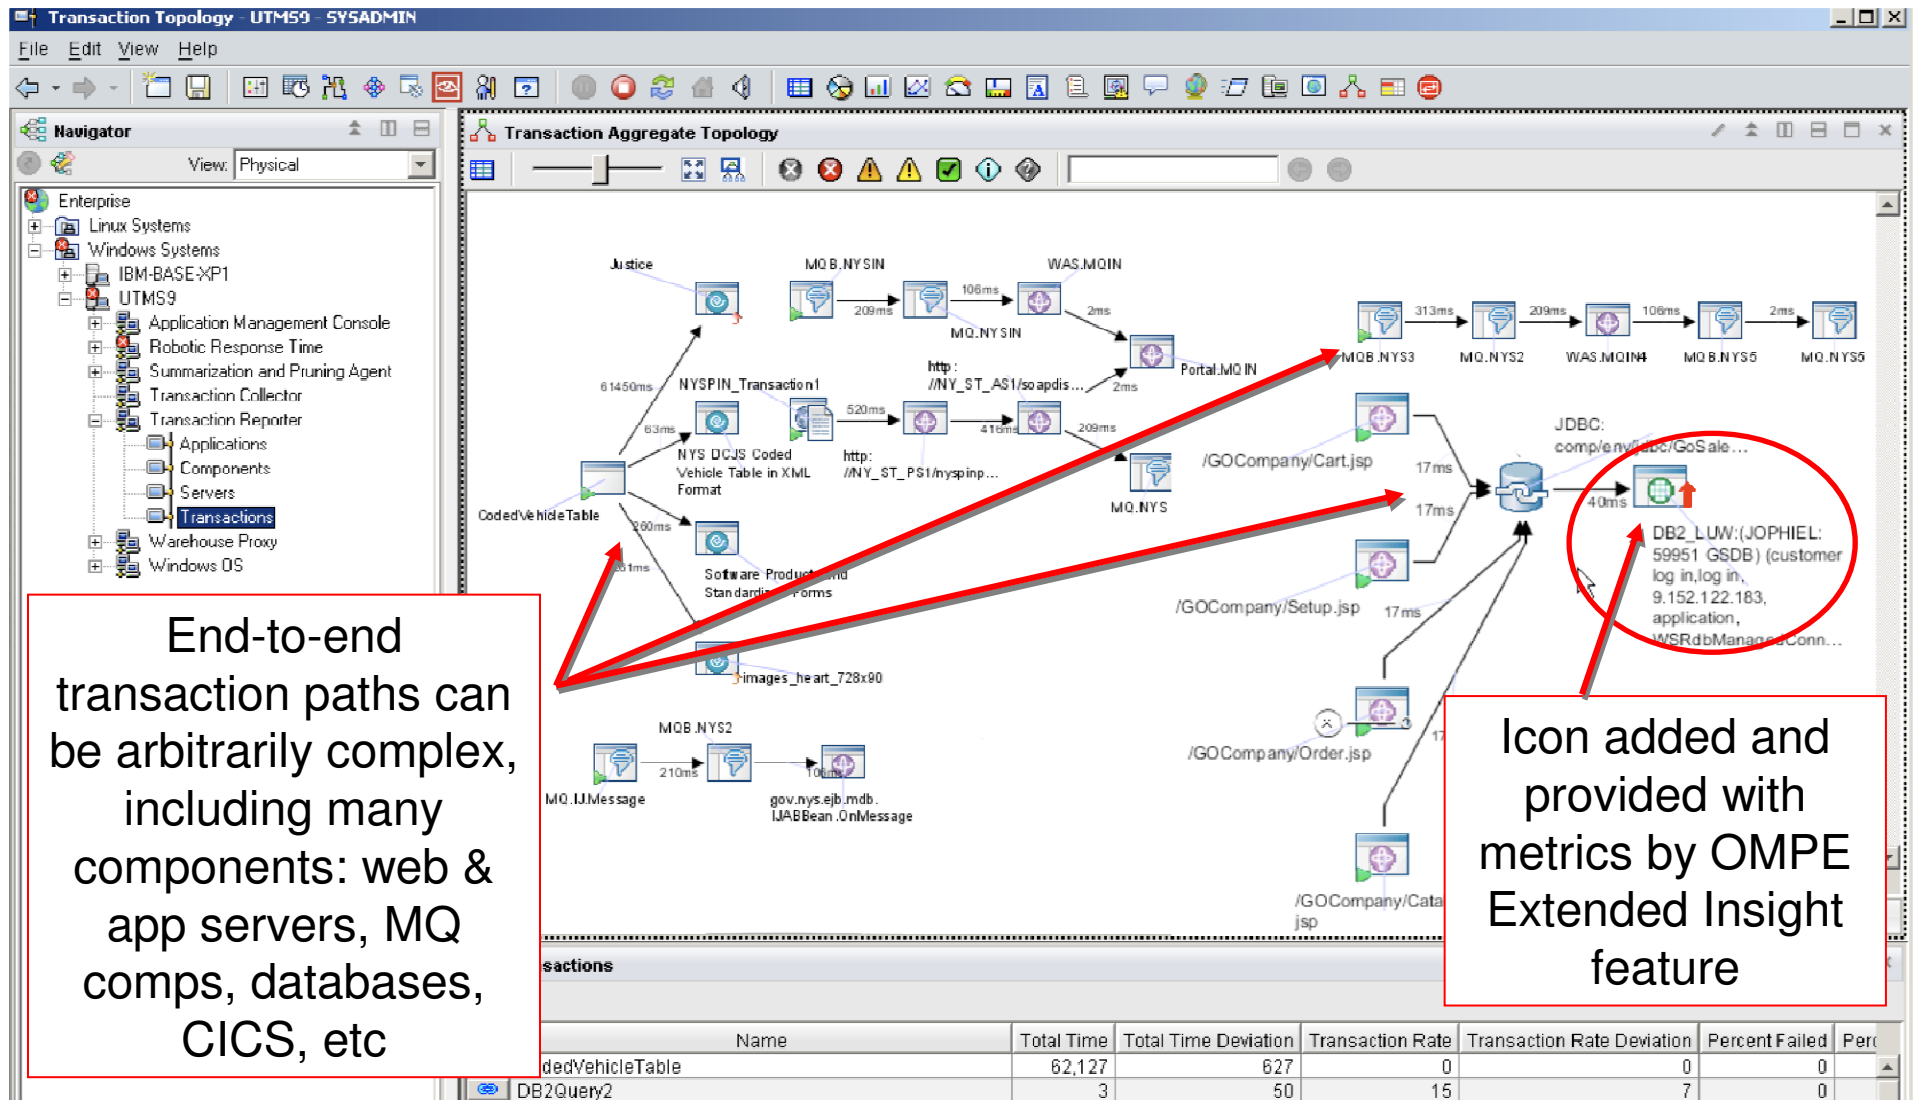
I/O Statistics

Buffer Pool Hit Ratio:	99.903 %
Logical page I/O:	2,058
Physical page I/O:	2

Data retrieved from the host Dynamic SQL Statement Cache

Transferring data from localhost...

# Seek out any problems in end-to-end transaction path → use ITCAM for Transactions in the TEP console

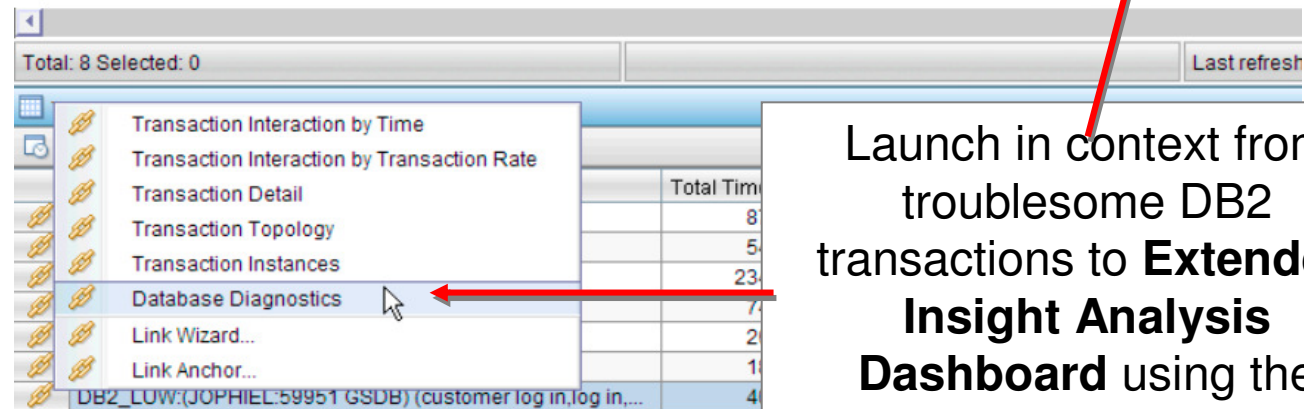
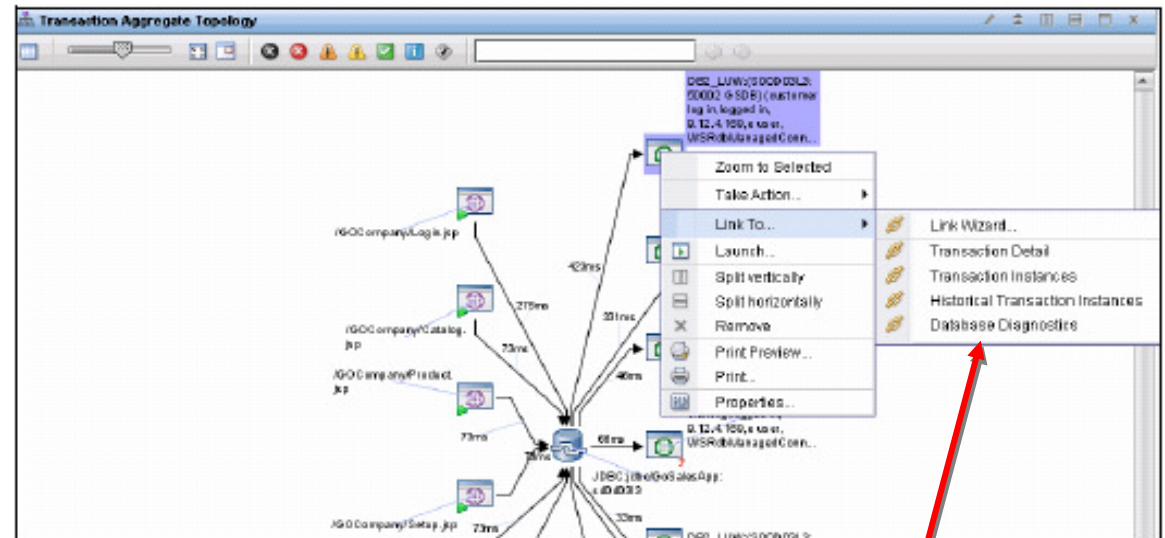


End-to-end transaction paths can be arbitrarily complex, including many components: web & app servers, MQ comps, databases, CICS, etc

Icon added and provided with metrics by OMPE Extended Insight feature

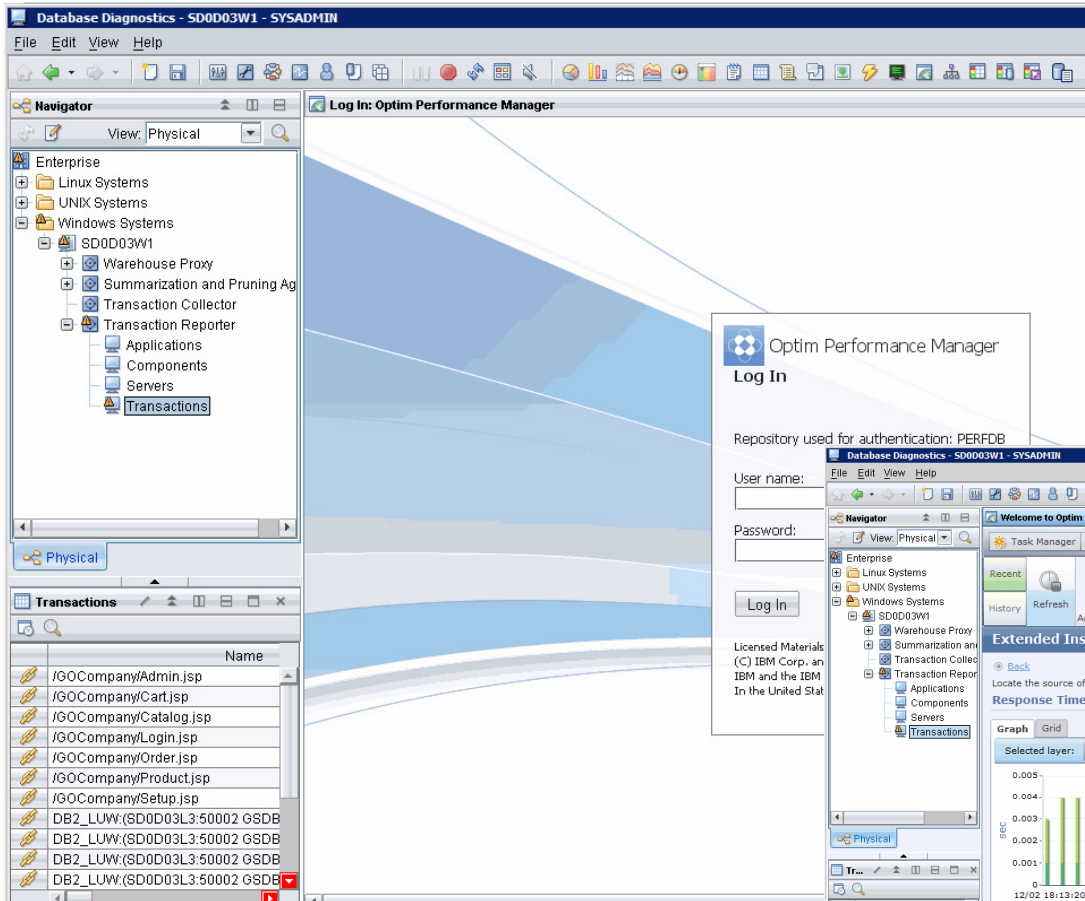


Launch Extended Insight Analysis dashboard GUI in context of the troublesome database transaction either from topology view or list box



Launch in context from troublesome DB2 transactions to **Extended Insight Analysis Dashboard** using the **“Database Diagnostics”** context link

# “Launch in context” workspace with Extended Insight dashboard



The browser is automatically positioned to the Extended Insight details dashboard for the specific transaction you selected in TEP. It knows the context because of the connection attributes

