



A Peek at IBM Optim Query Tuner

Hong S. Tie

*Senior Development Manager
Optim Query Tuner Development
IBM Silicon Valley Lab
tiehs@us.ibm.com*

DB2 for z/OS Technical Conference

October 5-6, 2009

Taipei, Taiwan

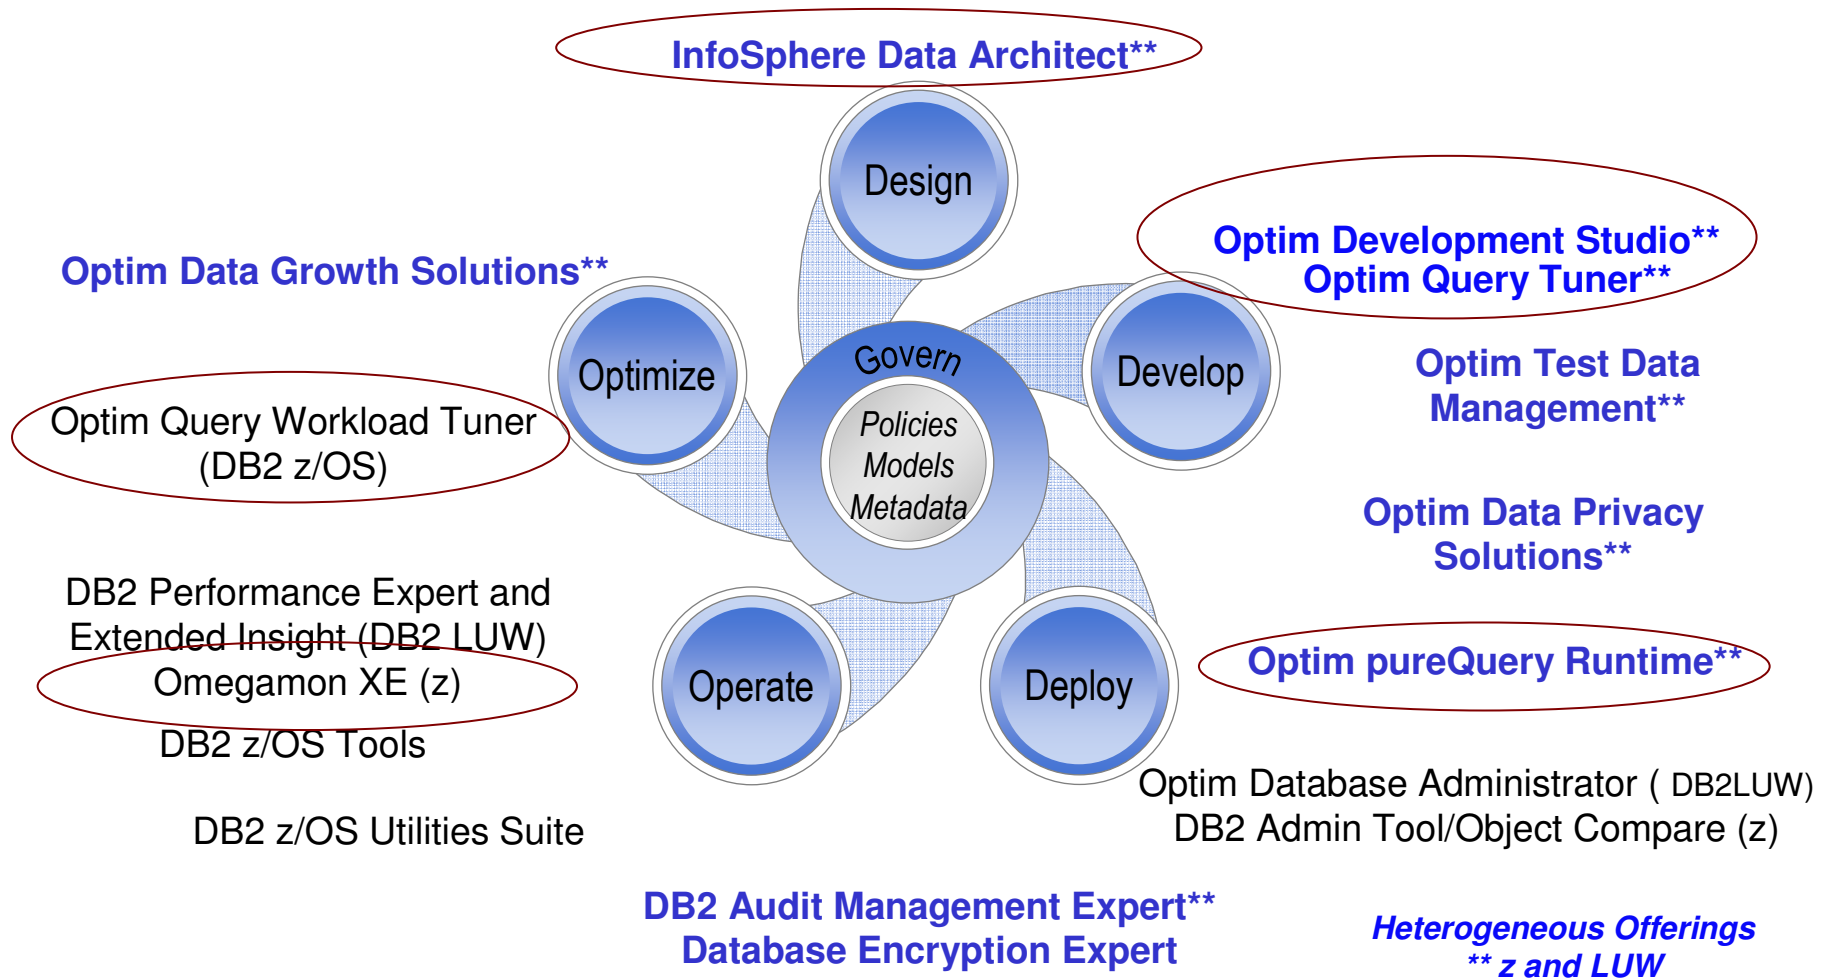
Agenda

- ***What is Optim Query Tuner***
- ***What role it plays in the Integrated Data Management Solutions***
- ***Business Value of Optim Query Tuner***
- ***Key Features/Functions Overview***
 - *Single Query tuning features*
 - *Workload Tuning features*
- ***Product Integration support***
- ***What's New for V2.2***
- ***Roadmap for the future***
- ***Summary***

What is *Optim Query Tuner*?

- *Integrated Data Management Solution that provides tools and advisors to optimize query performance and physical database design*
 - *(a.k.a DB2 Optimization expert)*
- *Reduces the complexity and manual effort required to perform query tuning*
- *Decreases the reliance on specialized Query Optimization and SQL tuning skills*
- *Provides a rich set of SQL tuning tools and design advisors that can be used Cross-Data Lifespan*
- *Eclipse Based with integration and common shell sharing with:*
 - *InfoSphere data Architect*
 - *Optim Development Studio (with pureQuery)*
 - *Optim Database Administrator*
 - *and more...*
- *Support Both DB2 for z/OS V8 and V9 and DB2 for LUW*

Optim Integrated Data Management Products – Overview



Optim Query Tuner Key Features at a Glance

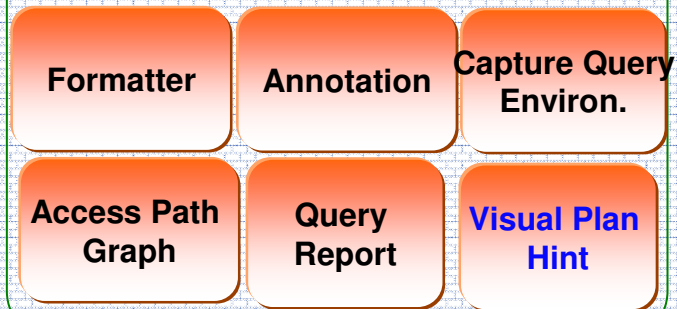
Query Tuner User Interface

Eclipse



Query Based Tools and Advisors




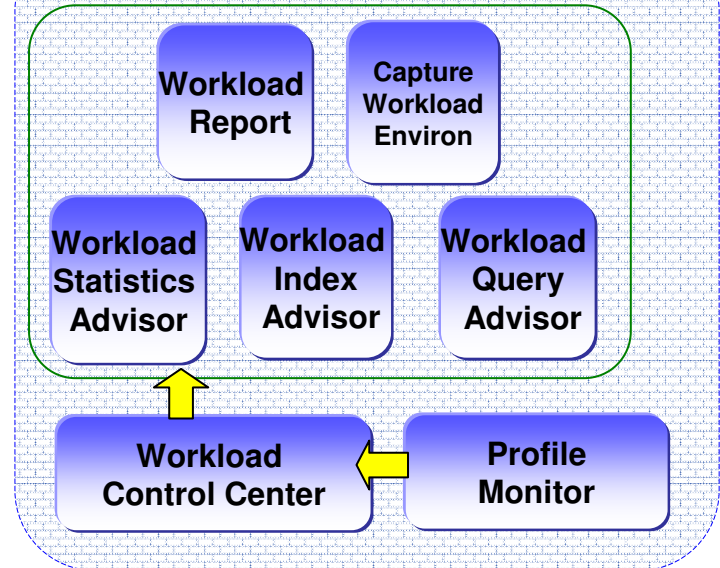
Query Tools



Query Advisors



Workload Advisors



IT Challenges Today

- **Applications are designed and implemented very quickly today**
- **Insufficient skill and resources to perform an adequate review of SQL performance and database physical design**
- **Entire applications can be developed and/or enhanced with performance "surprises" discovered in production**
- **Tuning an entire workload requires**
 - Analyzing each query for Access Path in the workload,
 - the frequency of query execution,
 - The expected CPU/ET of an individual query
- **Impact to business:**
 - Significant efforts and resources required to perform the SQL performance review
 - SQL Performance analysis may not be done, or done incompletely
 - *High cost of CPU and resources consumption*

Use IBM Optim Query Tuner/Workload Query Tuner Solutions

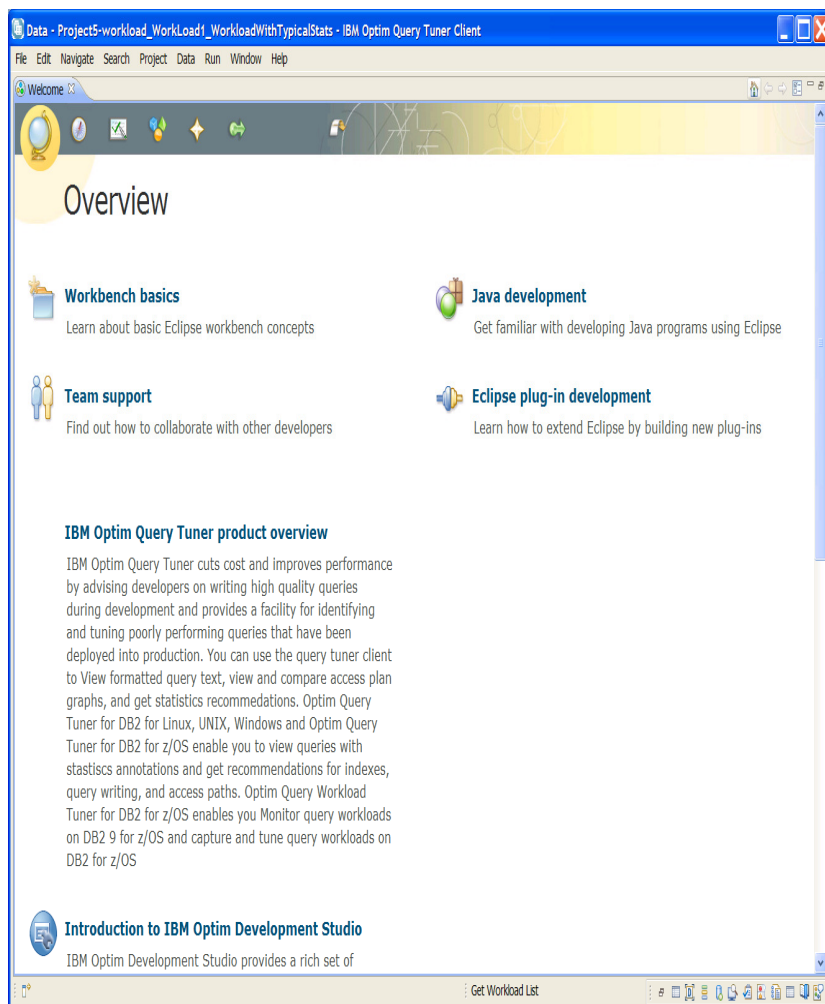
- **Identify query/workload candidates**
 - DB2 Catalog,
 - Dynamic statement Cache,
 - Text, File, package, QMF and more
- **Prevent problems before they impact the business**
 - Get early warning from Statistics advice, Access path advice, Query advice, index advice of emergent problems
 - Isolate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



IBM Optim Query Tuner

SQL Tuning for Single Query

QT Welcome Page – Product Overview and First Steps



Data - Project5-workload_WorkLoad1_WorkloadWithTypicalStats - IBM Optim Query Tuner Client

File Edit Navigate Search Project Data Run Window Help

Welcome

Overview

Workbench basics
Learn about basic Eclipse workbench concepts

Java development
Get familiar with developing Java programs using Eclipse

Team support
Find out how to collaborate with other developers

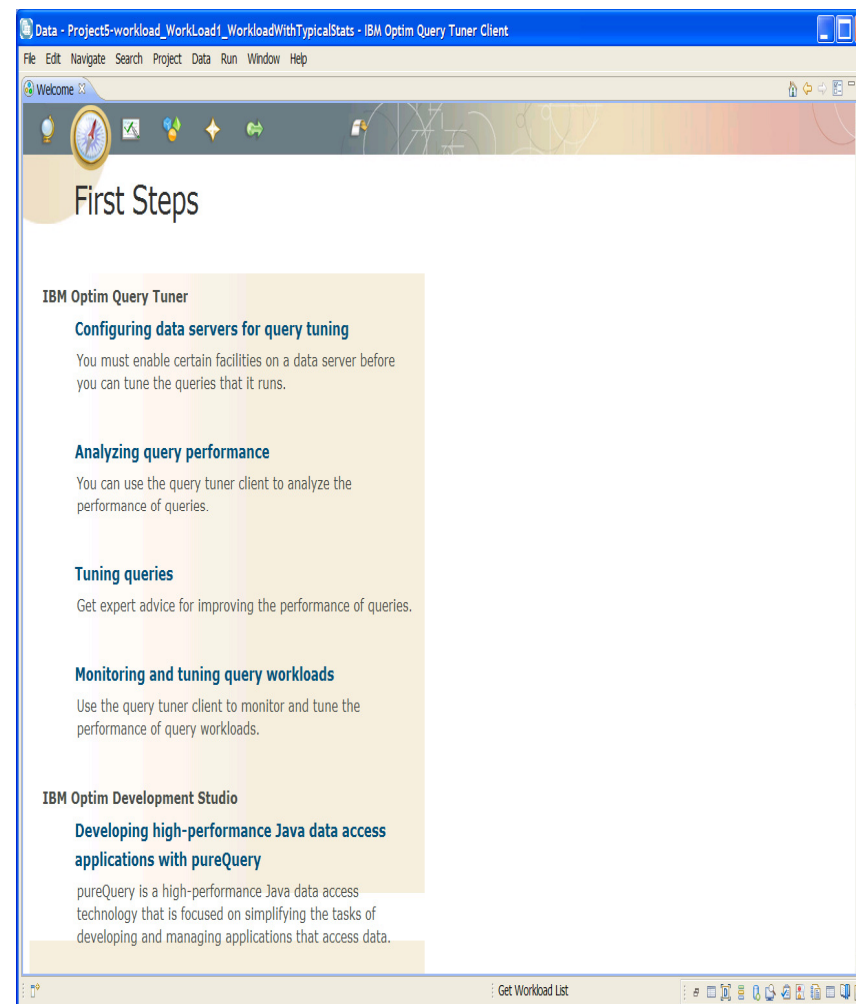
Eclipse plug-in development
Learn how to extend Eclipse by building new plug-ins

IBM Optim Query Tuner product overview

IBM Optim Query Tuner cuts cost and improves performance by advising developers on writing high quality queries during development and provides a facility for identifying and tuning poorly performing queries that have been deployed into production. You can use the query tuner client to View formatted query text, view and compare access plan graphs, and get statistics recommendations. Optim Query Tuner for DB2 for Linux, UNIX, Windows and Optim Query Tuner for DB2 for z/OS enable you to view queries with statistics annotations and get recommendations for indexes, query writing, and access paths. Optim Query Workload Tuner for DB2 for z/OS enables you Monitor query workloads on DB2 9 for z/OS and capture and tune query workloads on DB2 for z/OS

Introduction to IBM Optim Development Studio
IBM Optim Development Studio provides a rich set of

Get Workload List



Data - Project5-workload_WorkLoad1_WorkloadWithTypicalStats - IBM Optim Query Tuner Client

File Edit Navigate Search Project Data Run Window Help

Welcome

First Steps

IBM Optim Query Tuner

Configuring data servers for query tuning
You must enable certain facilities on a data server before you can tune the queries that it runs.

Analyzing query performance
You can use the query tuner client to analyze the performance of queries.

Tuning queries
Get expert advice for improving the performance of queries.

Monitoring and tuning query workloads
Use the query tuner client to monitor and tune the performance of query workloads.

IBM Optim Development Studio

Developing high-performance Java data access applications with pureQuery
pureQuery is a high-performance Java data access technology that is focused on simplifying the tasks of developing and managing applications that access data.

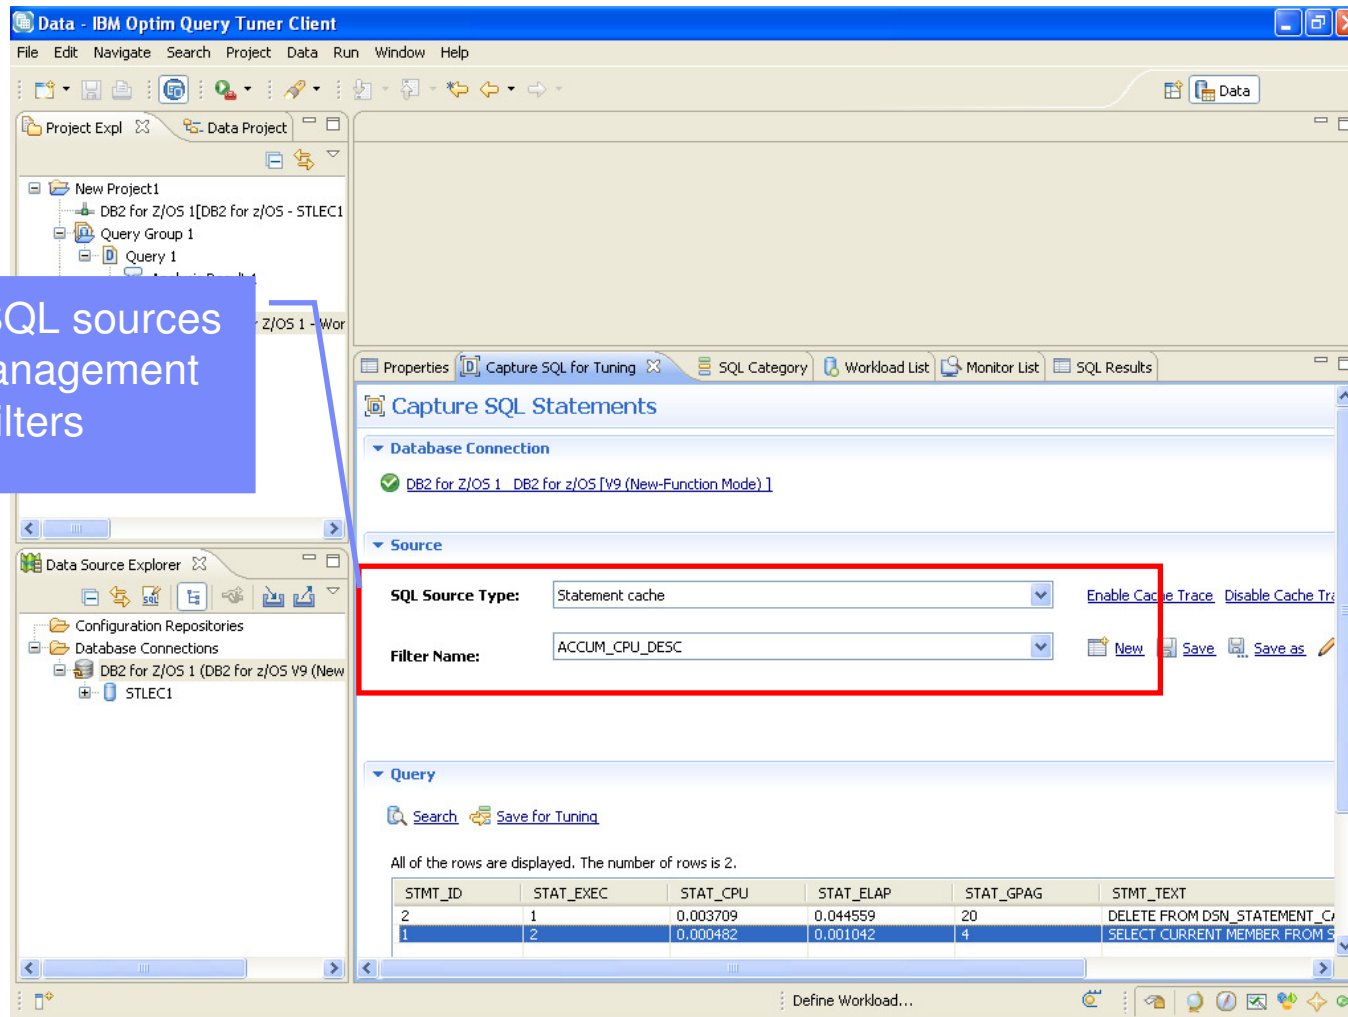
Get Workload List

Capture SQL for Tuning

- **Identify SQL with performance issues**
- **Capture SQL for Tuning**
 - *DB2 for z/OS*
 - *Cache,*
 - *Catalog,*
 - *QMF, QMFHPO*
 - *Text, File, Category*
 - *PLAN Table/Statement Table/Function Table*
 - *Workload/Monitor*
 - *DB2 for Linux, Unix and Windows*
 - *Text,*
 - *File,*
 - *Category*
 - *Package*

Capture SQL for Tuning

Select SQL sources and management filters



Tuning for Single Query – Input SQL Text

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Data - New Project1_Query Group 1_Query 1 - IBM Optim Query Tuner Client". The interface includes a menu bar (File, Edit, Navigate, Search, Project, Data, Run, Window, Help), a toolbar, and several panes:

- Project Explorer:** Shows a tree view with "New Project1" containing "DB2 for Z/OS 1[DB2 for z/OS - STLEC1]" and "Query Group 1" containing "Query 1".
- Optim Query Tuner:** The central pane, divided into sections:
 - Server Activation:** States "IBM Optim Query Workload Tuner for DB2 for z/OS is activated on the data server. All tuning functions are enabled." with a "Learn More" link.
 - Query Text:** A text area containing a complex SQL query. A blue callout box labeled "Input SQL Text" points to this area. The query is:


```
SELECT S_SUPPKEY, S_NAME, SUM(L_DISCOUNT) AS REVENUE FROM ORDER,
LINEITEM, SUPPLIER WHERE L_ORDERKEY = O_ORDERKEY AND L_EXTENDEDPRICE
> 200.0 AND O_ORDERDATE >= DATE('1993-10-01') AND O_ORDERDATE < DATE
('1993-10-01') + 3 MONTH AND S_NATIONKEY IN (1,2,3,4,5) AND S_SUPPKEY <
10000 AND O_ORDERKEY IN ( SELECT LLL.L_ORDERKEY FROM LINEITEM LLL WHERE
LLL.L_ORDERKEY = 20000000 AND LLL.L_EXTENDEDPRICE > 3000.0) GROUP BY
S_SUPPKEY, S_NAME HAVING SUM(L_DISCOUNT) > (SELECT MIN(L_DISCOUNT)
FROM lineitem, order WHERE l_orderkey = o_custkey AND o_orderdate = '1998-01-
01' and o_orderkey > 10000000 ) ORDER BY REVENUE DESC
```
 - Tuning Options:** A list of actions including "Tune Query...", "Choose Tuning Activities...", "Generate Report", "Query Environment Capture", "Test Candidate Indexes...", and "Create Plan Hint". A blue callout box labeled "Start to Tune" points to the "Tune Query..." button.
 - EXPLAIN options:** Radio buttons for "Re-capture EXPLAIN information" (selected) and "Use existing server EXPLAIN information".
- Data Source Explorer:** Shows "Configuration Repositories" and "Database Connections" including "DB2 for Z/OS 1 (DB2 for z/OS V9 (New STLEC1)".
- Properties Table:** A table with columns "Property" and "Value" is visible at the bottom.

Single Query Tuning Result

Maximum the editor

Access Plan Graph

Query Annotation

Advisor Recommendation Overview

Advisor Detail

The screenshot displays the IBM Optim Query Tuner Client interface. The main window shows the query editor with the following SQL:

```
SELECT SYSADM.SUPPLIER.S_SUPPKEY
       , SYSADM.SUPPLIER.S_NAME
       , SUM( SYSADM.LINEITEM.L_DISCOUNT ) AS REV
FROM SYSADM.SUPPLIER
```

The **Access Plan Graph** shows a query plan with a root node **QUERY** (2500) and a child node **WFSCAN** (2500).

The **Advisor Recommendation Overview** table lists the following recommendations:

Advisor	Priority	Description
Recommendations		
Statistics Advisor	HIGH	Repair statistics problems for
Access Path Advisor	HIGH	The optimizer chooses a join
Access Path Advisor	HIGH	The SYSADM.ORDER table is

The **Statistics Recommendation Detail** panel shows the following RUNSTATS Control Statements:

```
RUNSTATS TABLESPACE DB4SUPPLY.TSSUPPLY
TABLE(SYSADM.SUPPLIER)
COLUMN(S_NATIONKEY)
INDEX(SYSADM.PX5@SKNK HISTOGRAM NUMCOLS 1
NUMQUANTILES 25)
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
```

Why Query Formatting/Annotation?

A peek at an unformatted Query

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

Query Annotation - Annotate a statement

The screenshot shows the IBM Optim Query Tuner Client interface. At the top, there are tabs for 'Original' and 'Transformed'. A red dashed box highlights these tabs, with a callout box stating 'Click to show the tab of original query'. Below the tabs is a toolbar with buttons for 'Expand All', 'Collapse All', 'Customize', 'Save', 'Print', and 'Clear Highlights'. The main area is divided into two sections: 'Formatted Query' on the left and 'Annotation' on the right. The 'Formatted Query' section contains a SQL query snippet. The 'Annotation' section contains a table of annotations. A callout box labeled 'Tool bar' points to the toolbar. A callout box labeled 'Formatted Query Area' points to the SQL query text. A callout box labeled 'Annotation Area' points to the table of annotations.

Annotation		Additional Inf
CARDF=25	QUALIFIED_ROWS=1	
COLCARDF=(not applicable)	MAX_FRE	
COLCARDF=25	MAX_FREQ=	
COLCARDF=25/(not applicable)	MAX_FF	
COLCARDF=25	MAX_FREQ=	

Query Annotation - Transformed Query Annotation

The screenshot shows the 'Query Format and Annotation' window in the IBM Optim Query Tuner Client. At the top, there are two tabs: 'Original' and 'Transformed'. Both tabs are circled in red. A red arrow points from the 'Transformed' tab to a blue callout box containing the text 'Click to show the tab of transformed query'. Below the tabs, there is a dropdown menu for 'Annotation to display' set to 'All'. A toolbar includes buttons for 'Expand All', 'Collapse All', 'Customize', 'Save', 'Print', and 'Clear Highlights'. The main area is divided into three columns: 'Formatted Query', 'Annotation', and 'Additional Info'. The 'Formatted Query' column contains a SQL query starting with '(SELECT * FROM SYSADM.NATION'. The 'Annotation' column shows various statistics such as 'CARDF=2', 'CARDF=1', and 'COLCARDF=45,000,000'. The 'Additional Info' column contains dropdown menus for each annotation. At the bottom, there is a 'Context' section with 'Query' and 'Analysis Result 3' tabs.

Query Annotation - Select Annotation to display

The screenshot shows the IBM Optim Query Tuner Client interface. The main window displays a query with various annotations. A dropdown menu is open, showing options for 'Annotation to display': All, Catalog Statistics, and Cost Estimation. A red dashed box highlights this dropdown and the 'Expand All' button. A blue callout box points to the dropdown with the text 'Select which annotation to be displayed'. Another blue callout box points to a specific annotation 'COLCARDF (not applicable)' in the table with the text 'Select which annotation to be displayed in a line'. A third blue callout box points to a node in the query tree with the text 'Click a node to expand or collapse it'.

Formatted Query	Annotation	Additional Info
SELECT *		
FROM SYSADM.NATION	CARDF=25	QUALIFIED_ROWS=1
, SYSADM.LITEM_UIV_NOGBY		
WHERE (SYSADM.LITEM_UIV_NOGBY.ORDERKEY BETWEEN 6000001 AND 45000000	COLCARDF=(not applicable)	MAX_FREQ=
AND SYSADM.NATION.N_NATIONKEY > 1	COLCARDF=25	MAX_FREQ=
AND SYSADM.NATION.N_NATIONKEY = SYSADM.LITEM_UIV_NOGBY.ORDERKEY	COLCARDF=25/(not applicable)	MAX_FREQ=
AND SYSADM.NATION.N_NATIONKEY IN (COLCARDF=25	MAX_FREQ=
)		

Stats Advisor (SA)

Optim Query Tuner

Server Activation
IBM Optim Query Workload Tuner for DB2 for z/OS is activated on the data server. All tuning functions are enabled. [Learn More.](#)

Query Text
Specify the query text for query tuner.

```
select S.DBNAME as SDBNAME , S.NAME as S
S.BPOOL as SBPOOL , S.PARTITIONS as S
SLOCKRULE , S.PGSIZE as SPGSIZE , S.EF
S.STATUS as SSTATUS , S.IMPLICIT as SI
SNTABLES , S.CLOSERULE as SCLOSERUI
S.STATSTIME as SSTATSTIME , S.LOCKM
SLOCKPART , S.TYPE as STYPE , S.ENCO
SENCODING_SCHEME , S.SBCS_CC SID as
SDBCS_CC SID , S.MIXED_CC SID as SMIX
```

EXPLAIN options: Re-capture EXPLAIN inform

Context

Choose Tuning Activities

- Query Format and Annotation
- Access Plan Graph
- Statistics Advisor
- Query Advisor
- Access Path Advisor
- Index Advisor

Select All Clear All

OK Cancel

Tune Query...
Choose Tuning Activities...
Generate Report
Query Environment Capture
Test Candidate Indexes...
Create Plan Hint

Stats Advisor – Recommendation Summary

Advisor Recommendation Overview

Advisor	Priority	Description
Recommendations		
Statistics Advisor	HIGH	Repair statistics problems for thi
Statistics Advisor	MEDIUM	Determine the access path agair
Statistics Advisor	MAINTENANCE	Gather and recollect all of relev

Context
Query Analysis Result 1

Statistics Recommendation Detail

Recommendations

RUNSTATS Control Statements

```
RUNSTATS INDEX("SYSADM", "PX5@SKNK" FREQVAL NUMCOLS 1
COUNT 15)
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE

RUNSTATS TABLESPACE "DB4CUST", "TSCUST"
TABLE("SYSADM", "CUSTOMER") SAMPLE 40
COLUMN("C_MKTSEGMENT")
COLGROUP("C_MKTSEGMENT") FREQVAL COUNT 15
SORTNUM 4
```

Previous RUNSTATS statement stored in database

```
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE

RUNSTATS TABLESPACE "DB4CUST", "TSCUST"
TABLE("SYSADM", "CUSTOMER") SAMPLE 40
COLUMN("C_MKTSEGMENT")
COLGROUP("C_MKTSEGMENT") FREQVAL COUNT 15
SORTNUM 4
INDEX("SYSADM", "PXC@CKNKMS" KEYCARD,
"SYSADM", "UXC@NKCK")
```

Table, index, column, and column group details

Search: search key word in report

Statistics Advisor Summary Report
Analysis start time: 2009-06-17 12:58:10.298
Analysis end time: 2009-06-17 12:58:12.97

TABLE SYSADM.SUPPLIER
Table type: Table

Recommendations with different priority

Stats Advisor – Report

▶ Recommendations

▼ Table, index, column, and column group details

Search

search key word in report

search key word

key word

 Case-sensitive

(% = any...: % ? {)

 Regular expression

Highlight All

Filter

Statistics Advisor Detail Report

Analysis start time: 2009-06-17 12:58:10.298

Analysis end time: 2009-06-17 12:58:12.97

TABLE SYSADM.SUPPLIER

Table type: Table

Cardinality: 300000.0

Collection time: 0001-01-01 00:00:00.0

Statistics status: OK

INDEXES:

SYSADM.PXS@SKNK (S_SUPPKEY,S_NATIONKEY)

First key cardinality: 300000.0

Full key cardinality: 300000.0

Data repetition factor: -1.0

Collection time: 0001-01-01 00:00:00.0

Statistics status: missing

▼ Conflicts detail

TABLE SYSADM.SUPPLIER

One of the frequency records (-1.0) of the S_SUPPKEY column group is out of range [0,1]

Tolerance: 0.0010

The maximum frequency of the column group or column (S_SUPPKEY), (0.0), is less than the average frequency, or 1 divided by the cardinality for the column group or column (3.3333333333333333E-6). The maximum frequency is expected to be greater than the average unless only least-frequently occurring values are being collected.

Tolerance: 0.0010

Detail Report

Conflict Report

Stats Profile Support – new for V2.2

- **Produce consolidated recommendation**
 - SA will intersect/merge previous saved RUNSTATS command with the new RUNSTATS command to provide a consolidated RUNSTATS command.
 - For data server without statistical profile support, server side stats profile table will be created while creating the EXPLAIN tables to store information about the table (name, schema) and the current stats profile for the table.
 - User can save current or load previous SA recommendations (REPAIR or CONSOLIDATE) .

Statistics Advisor – A case of Stats Profile

Statistics Recommendation Detail

Recommendations

RUNSTATS Control Statements

```
RUNSTATS INDEX("SYSADM", "PXS@SKNK" FREQVAL NUMCOLS 1 COUNT 15)
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE

RUNSTATS TABLESPACE "DB4CUST", "TSCUST"
TABLE("SYSADM", "CUSTOMER") SAMPLE 40
COLUMN("C_MKTSEGMENT")
COLGROUP("C_MKTSEGMENT") FREQVAL COUNT 15
SORTNUM 4
INDEX("SYSADM", "PXC@CKNKMS" KEYCARD,
```

Previous RUNSTATS statement stored in database

```
RUNSTATS TABLESPACE "DB4CUST", "TSCUST"
TABLE("SYSADM", "CUSTOMER") SAMPLE 40
COLUMN("C_MKTSEGMENT")
COLGROUP("C_MKTSEGMENT") FREQVAL COUNT 15
SORTNUM 4
INDEX("SYSADM", "PXC@CKNKMS" KEYCARD,
"SYSADM", "UXC@NKCK")
SHRLEVEL CHANGE REPORT YES UPDATE ALL HISTORY NONE
```

Table, index, column, and column group details

Search

search key word in report

search key word

key word

Case-sensitive

Regular expression

(% = any...: % ? {)

Highlight All

Filter

Statistics Advisor Summary Report

Analysis start time: 2009-06-17 12:58:10.298

Analysis end time: 2009-06-17 12:58:12.97

TABLE SYSADM.SUPPLIER

Table type: Table

Cardinality: 300000.0

Collection time: 0001-01-01 00:00:00.0

Statistics status: OK

INDEXES:

SYSADM.PXS@SKNK (S_SUPPKEY,S_NATIONKEY)

First key cardinality: 300000.0

Full key cardinality: 300000.0

Data repetition factor: -1.0

Collection time: 0001-01-01 00:00:00.0

Statistics status: missing

Show Current and previous RUNSTATS commands

Stats Advisor – Integration with APG (new)

Slow Query - First run

Advisor Recommendation Overview

Priority	Advisor	Recommendation Description
High	Statistics Advisor	Repair statistics problems for this query. Gather missing statistics. R...
High	Index Advisor	Create index LINEITEM_VIRT_IDX_1181023618203 L_RETURNFLAG(ASC), L...
Low	Statistics Advisor	Gather and recollect all of relevant statistics for this query.
Low	Query Advisor	Consider replacing the asterisk (*) or the long column list of table ...
Low	Access Path Advisor	The table LINEITEM is accessed by a relational scan. When a large number of ...
Low	Index Advisor	Create index ORDER_VIRT_IDX_12130012345567 RECLLEN(ASC).

[Details](#)
[Generate report](#)
[Service SQL](#)

Formatted Query

```

SELECT X.DISCOUNT
, SUM( X.QUANTITY ) AS SUMQ
, AVG( X.QUANTITY * X.PRICE ) AS AVGP
FROM SYSADM.LINEITEM_UIV_NOGBY AS X
WHERE ( X.QUANTITY <= 5
AND X.DISCOUNT > 0.08
AND X.ORDERKEY > 30000000
)
GROUP BY X.DISCOUNT
    
```

Annotation

QUALIFIED_ROW
 LOW2KEY (not a
 COLCARDF=(not
 MAX_FREQ (not :

Execution Plan

Node Type : Table[16]
 Table Name : LINEITEM
 Creator Name : SYSADM
 Cardinality : 4254339
 < partitioned table space (index based) >
 < The RUNSTATS command was not run >

Context

Analysis_Result_1 2008/05/29 2:43pm

Explain Options

Query number: 101
 SQLID : SYSADM
 Optimization hint:
 Group member: 5

Application Environment

Schema: SYSADM
 Current degree: ANY
 Current refresh age: ANY
 Current maintained table types: ALL

Query Analysis_Result_1

Access Path Advisor – Access Path Warning

The screenshot displays the IBM Optim Query Tuner Client interface. At the top, the title bar reads "Data - New Project1_Query Group 1_Query 1 - IBM Optim Query Tuner Client". The main window is titled "Advisor Recommendation Overview" and contains a table of recommendations. A red box highlights the "Access Path Advisor" entries. A callout box labeled "Recommendations List" points to this table. Below the overview, the "Advisor Detail" window is open, showing "Path Recommendation Detail". Two callout boxes, "Warning Description" and "Warning Explanation", point to the respective text areas in this window. At the bottom, a "plan table related with this warning" is shown, with a callout box labeled "Corresponding records in PLAN_TABLE" pointing to it.

Advisor	Priority	Description
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect con
Access Path Advisor	HIGH	The optimizer chooses a join sequence that contains a new table SYSADM.ORDER
Access Path Advisor	HIGH	The SYSADM.ORDER table is accessed by a non-matching index scan (QBLOCKM
Query Advisor	MEDIUM	No join predicate exists between table SYSADM.SUPPLIER and other tables. Co
Access Path Advisor	LOW	The SYSADM.SUPPLIER table is accessed by a relational scan (QBLOCKM = 1

Description	Explanation
The optimizer chooses a join sequence that contains a new table SYSADM.ORDER without join predicates to any of the composite tables. When a large number of records are returned, DB2 might be using an inefficient access path. Check the explanation for this warning for more details about the possible cause and solution.	The optimizer chooses a join sequence that contains a Cartesian join. A Cartesian join is a form of nested loop join in which there are no join predicates between the two tables. However, a Cartesian join is generally not the best access path if the number of qualified rows from the outer table or the inner table is large. This problem is often caused by lack of applicable joined predicates.

ID	MIXOPSEQ	METHOD	CREATOR	TNAME	CORREL...	ACCESS...	ACCESS...	ACCESS...	MATCHC...	MERGE_...	INDEXO...	PREFE
1	5	0	3						0		N	
1	4	0	3						0		N	
1	3	0	2	SYSADM	LINEITEM	I	SYSADM	PXL@OK...	0	1	Y	S
1	2	0	1	SYSADM	ORDER	I	SYSADM	PXO@OK...	0		Y	S
1	1	0	0	SYSADM	SUPPLIER	R			0		N	S

Access Path Advisor and its recommendations

- *A example Query with Access plan of Non-matching index access*

Selected Recommendation:

Description	Explanation
The DSN8910.EMP table is accessed by a non-matching index scan (QBLOCKNO = 1, PLANNO = 1). If a table is accessed by non-matching index scan, then all the index keys and their RIDs are read. When a large number of keys and RIDs are accessed, DB2 might be using an inefficient access path. Check the explanation for this warning for more details about the possible cause and solution.	In a non-matching index scan, no matching columns are in the index. Consequently, all of the index keys must be examined. Because a non-matching index usually provides no filtering, very few cases provide an efficient access path. This problem is often caused by the lack of an appropriate index. Run the index advisor to determine whether creating an index might improve the access path.

PLAN_TABLE record

The following row in the plan table related with this warning.

QBLO...	PLANNO	MIXOP...	METHOD	CREA...	TNAME	CORR...	ACCE...	ACCE...	ACCE...	MATC...	MERG...	INDEX...	PREFE...	SORT...	SORT...	SOI
1	3	0	1	DSN89...	EPROJ	C	R			0		N		N	Y	N
1	2	0	1	DSN89...	DEPT	B	I	DSN89...	XDEPT1	1		N		N	N	N
1	1	0	0	DSN89...	EMP	A	I	DSN89...	XEMP1	0		N		N	N	N

Access Path Advisor – Hide/Show Recommendations

Recommendations List

Advisor	Priority	Description
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect con
Access Path Advisor	HIGH	The optimizer chooses a join sequence that contains a new table SYSADM.ORDER
Access Path Advisor	HIGH	The SYSADM.ORDER table is accessed by a non-matching index scan (QBLOCKM
Query Advisor	MEDIUM	No join predicate exists between table SYSADM.SUPPLIER and other tables. Co
Access Path Advisor	LOW	The SYSADM.SUPPLIER table is accessed by a relational scan (QBLOCKNO = 1

Show/Hide Recommendations

Hide Recommendation

Selected Recommendation:

Description	Explanation
The optimizer chooses a join sequence that contains a new table SYSADM.ORDER without join predicates to any of the composite tables. When a large number of records are returned, DB2 might be using an inefficient join sequence. For more details about the explanation for this warning for more details about the solution.	The optimizer chooses a join sequence that contains a Cartesian join. A Cartesian join is a form of nested loop join in which there are no join predicates between... a Cartesian join is generally not the best access path for a table or the inner table is... by lack of applicable joined predicates.

PLAN_TABLE record

The following row in the plan table related with this warning.

QBLOCKNO	PLANNO	MIXOPSEQ	METHOD	CREATOR	TNAME	CORREL...	ACCESS...	ACCESS...	ACCESS...	MATCHC...	MERGE_...	INDEXO...	PREFE
1	5	0	3							0		N	
1	4	0	3							0		N	
1	3	0	2	SYSADM	LINEITEM		I	SYSADM	PXL@OK...	0	1	Y	S
1	2	0	1	SYSADM	ORDER		I	SYSADM	PXO@OK...	0		Y	S
1	1	0	0	SYSADM	SUPPLIER		R			0		N	S

Query Advisor - Recommendations

Advisor Recommendation Overview

Advisor	Priority	Description
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Recollect con
Access Path Advisor	HIGH	The optimizer chooses a join sequence that contains a new table SYSADM,ORDE
Access Path Advisor	HIGH	The SYSADM.ORDER table is accessed by a non-matching index scan (QBLOCKM
Query Advisor	MEDIUM	No join predicate exists between table SYSADM.SUPPLIER and other tables. Co
Access Path Advisor	LOW	The SYSADM.SUPPLIER table is accessed by a non-matching index scan (QBLOCKM

Advisor Detail

Query Recommendation Detail

SQL Text

```

SELECT SYSADM.SUPPLIER.S_SUPPKEY
      , SYSADM.SUPPLIER.S_NAME
      , SUM( SYSADM.LINEITEM.L_DISCOUNT ) AS REVENUE
FROM SYSADM.SUPPLIER
      , SYSADM.ORDER
      , SYSADM.LINEITEM
WHERE ( SYSADM.LINEITEM.L_EXTENDEDPRICE > 200.0
      AND SYSADM.ORDER.O_ORDERDATE >= DATE( '1993-10-01' )
      AND SYSADM.ORDER.O_ORDERDATE < ( DATE( '1993-10-01' ) + 3 MONTHS )
      AND SYSADM.ORDER.O_ORDERKEY IN (
      SELECT LLL.L_ORDERKEY
      FROM SYSADM.LINEITEM AS LLL
    )
  
```

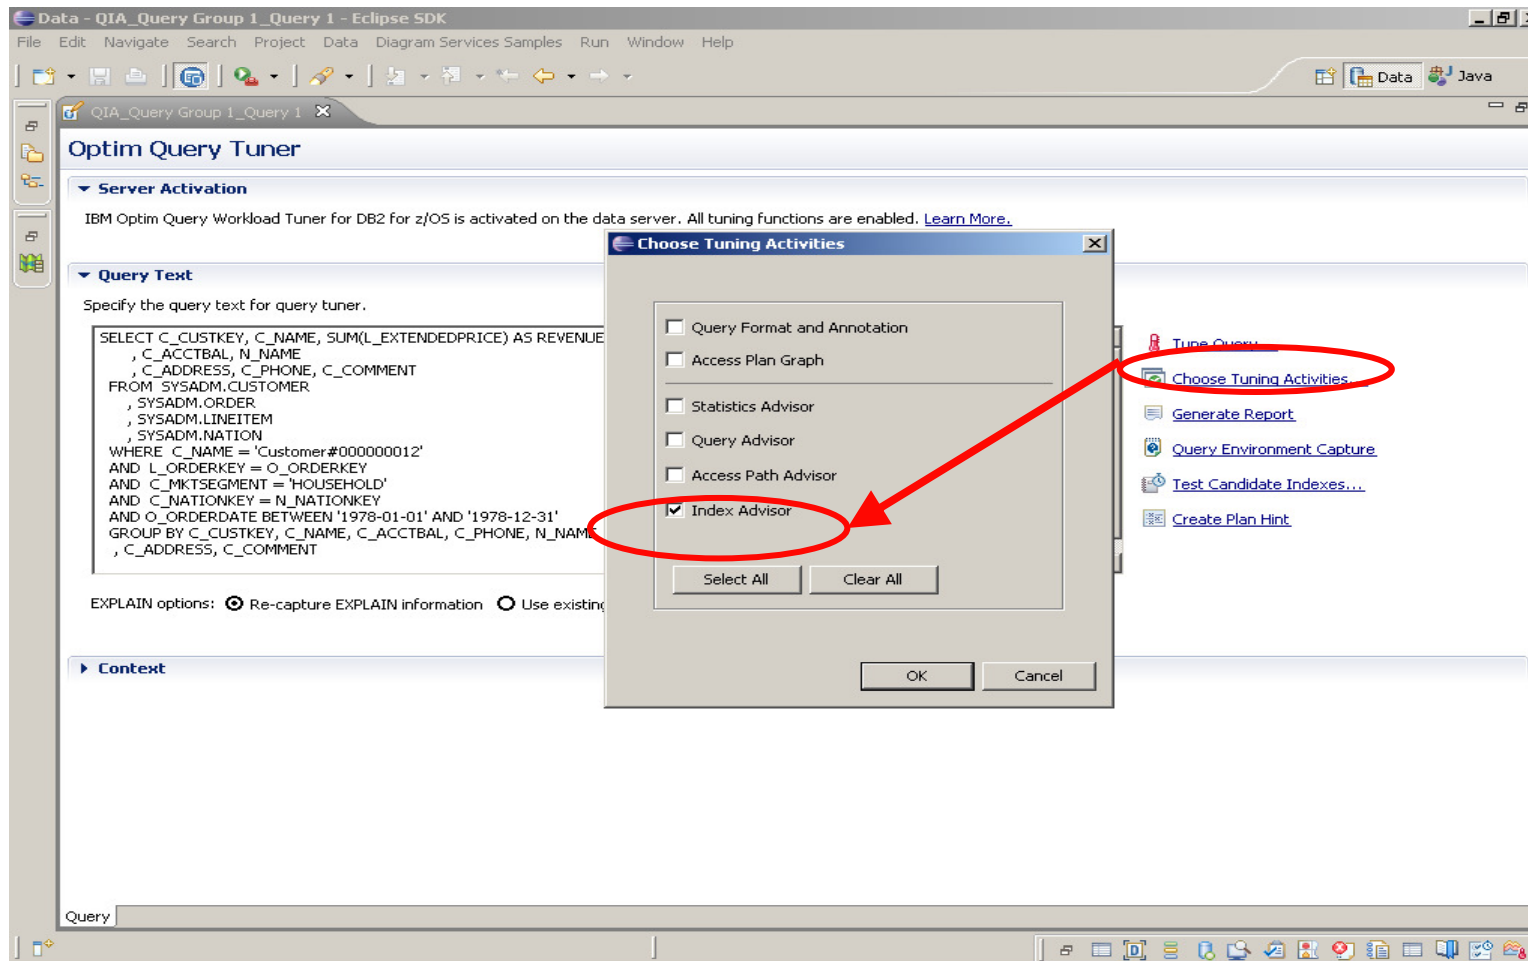
Selected Recommendation:

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

Index Advisor - Business Value

- **Reduce the cost of index design and evaluation**
- **Facilitate users to design appropriate indexes at development phase**
- **Lower the effort to diagnose and resolve performance problems caused by inappropriate indexes**
- **Identify, evaluate and drop inefficient indexes to recycle disk space**

Query Index Advisor – Launch for Indexes Tuning and recommendations



Index Advisor with Recommendations - Scenario 1

- ❑ Query Tuner Index Advisor Recommended with a list of indexes with estimated performance improvement below.

The screenshot shows the 'Advisor Recommendation Overview' window with the following table:

Advisor	Priority	Description
Recommendations		
Index Advisor	LOW	Index recommendations found.

The 'Advisor Detail' window shows the 'Recommendation Detail' view with the following data:

Performance Improvement:
 Estimated performance improvement: 98.74 %
 Disk space required (DASD space): 132.36 MB

Customized and Recommended Indexes

Feature Details	Creator	Object Name	Columns	Estimated Disk Space
<input checked="" type="checkbox"/> ORDER <input checked="" type="checkbox"/> Index	DB2OE	ORDER_VIRT_IDX_124417...	O_ORDERDATE(ASC),O_...	27.62109375 M
<input checked="" type="checkbox"/> LINEITEM <input checked="" type="checkbox"/> Index	DB2OE	LINEITEM_VIRT_IDX_1244...	L_ORDERKEY(ASC),L_EXT...	96.83203125 M
<input checked="" type="checkbox"/> CUSTOMER <input checked="" type="checkbox"/> Index	DB2OE	CUSTOMER_VIRT_IDX_12...	C_NAME(ASC),C_MKTSEG...	7.8828125 M
<input checked="" type="checkbox"/> NATION <input checked="" type="checkbox"/> Index	DB2OE	NATION_VIRT_IDX_12441...	N_NATIONKEY(ASC),N_N...	0.0234375 M

Existing indexes

Feature Details	Object Name	Columns
<input type="checkbox"/> ORDER Index	PXO@OKODCKS	O_ORDERKEY(ASC),O_O...

Recommended Indexes

Index Advisor - Run what-if with APG integration - Scenario 2 (new)

Check this option to show how virtual indexes are used in APG

Advisor	Priority
Index Advisor	LOW

Index	First key cardinality	Full key cardinality	PCTFREE	FREEPAGE	Cluster ratio
NATION_VIRT_IDX_1244171523296	-1	-1	0	10	0,5

Existing indexes	Feature Details	Obj
INDEX	INDEX	INDEX
CUSTOMER	INDEX	CUSTOMER
NATION	INDEX	NATION
NATION_VIRT_IDX_12441...	N_NATIONKEY(ASC),N_N...	0.0234375 M

Note: for this case, user can define index attributes based on the IA recommended index and run What If this index is selected for access path and the performance gain of it.

APG display IA recommend index and index stats

Selected Node Descriptor: EMP_VIRT_IDX_1245107511046

Description of Selected Node
Displays information about the node that is highlighted in the diagram.

- Index(EMP_VIRT_IDX_1245107511046)
 - Indexkeys
 - indexkey(EMP_VIRT_IDX_1245107511046 key1:SALARY)
 - Table

Attributes

Name	Value
Table Name	EMP
Name	EMP_VIRT_IDX_1245107511046
Creator	DB2OE
Unique Rule	D
Clustering	N
Cluster Ratio	0(default)
First Key Cardinality	25(default)
Full Key Cardinality	25(default)
Leaf Pages	34(default)
Levels	2(default)
Clustered	N
Type	2
Extension Type	Padded
Compress Index	N
Data Repeat Factor	
Timestamp	0001-01-01 00:00:00.0
Explain Time	2009-06-15 16:14:30.7

Description of the Selected Attribute

Index Advisor recommended index with Stats

Data - Project-443_Query Group 1_Query 4 - IBM Optim Query Tuner Client

File Edit Navigate Search Project Data Run Window Help

Manage Licenses

Project-443_Query G Project-443_Query Gr Projecta_Query Grou *Project5_Query Grou *Project-443_Query G

Query Format and Annotation

Original Transformed

Annotation to display: All

Formatted Query	Annotation	Additional Inform
SELECT * FROM DSN8910.EMP AS A WHERE A.SALARY > ? ORDER BY A.LASTNAME ASC	CARDF=42 QUALIFIED_ROI COLCARDF=(missing) MAX	

Access Plan Graph

Advisor Recommendation Overview

Advisor	Priority	Description
Statistics Advisor	HIGH	Repair statistics problems for this query. Gather missing statistics. Reco...
Statistics Advisor	MEDIUM	Determine the access path again at runtime. This query contains host ...
Query Advisor	LOW	Consider replacing the asterisk (*) or the long column list of table DSN...
Access Path Advic	LOW	A sort (QBLOCKNO = 1, PLANNO = 2) is used. When a large number o...
Statistics Advisor	MAINTENANCE	Gather and recollect all of relevant statistics for this query.

Choose Tuning Activities...
Generate Report
Query Environment Capture
Create Plan Hint

Context
Query / Analysis Result 1

Refresh Cache

Test Candidate Index – with user defined index Scenario 3 ((New))

Click Test Candidate Index in the Query page and choose/create an analysis result

```
FROM SYSADM.CUSTOMER
, SYSADM.ORDER
, SYSADM.LINEITEM
, SYSADM.NATION
WHERE C_NAME = 'Customer#00000012'
AND L_ORDERKEY = O_ORDERKEY
AND C_MKTSEGMENT = 'HOUSEHOLD'
AND C_NATIONKEY = N_NATIONKEY
AND O_ORDERDATE BETWEEN '1978-01-01' AND '1978-12-31'
GROUP BY C_CUSTKEY, C_NAME, C_ACCTBAL, C_PHONE, N_NAME
, C_ADDRESS, C_COMMENT
```

Specify an Analysis Result

Please specify an Analysis Result to test the candidate indexes.

Use existing local EXPLAIN Information from: Analysis Result 1

Create a new set of analysis results and collect EXPLAIN information: Analysis Result 2

OK Cancel

Test Candidate Index – with user defined index (New in 2.2)...

Virtual Index Definition

Customized and Recommended Indexes

[Add Index](#) [Edit Index](#) [Remove Index](#)

Feature Details	Creator	Object Name	Columns	Estimated Disk Space
<input checked="" type="checkbox"/> CUSTOMER <input checked="" type="checkbox"/> Index	MY	TESTIDX	C_CUSTKEY(ASC),C_ADD...	0.0 M

Existing indexes

Feature Details	Object Name	Columns
<input type="checkbox"/> CUSTOMER Index	IDX_PM_CUS_KEY	C_CUSTKEY(ASC)
<input type="checkbox"/> NATION Index	IDX_PM_NATI...	N_NATIONKEY(ASC)
<input type="checkbox"/> ORDERS Index	IDX_PM_ORDE...	O_ORDERKEY(ASC)
<input type="checkbox"/> LINEITEM Index	IDX_PM_ORDER...	L_ORDERKEY(ASC),L_LIN...

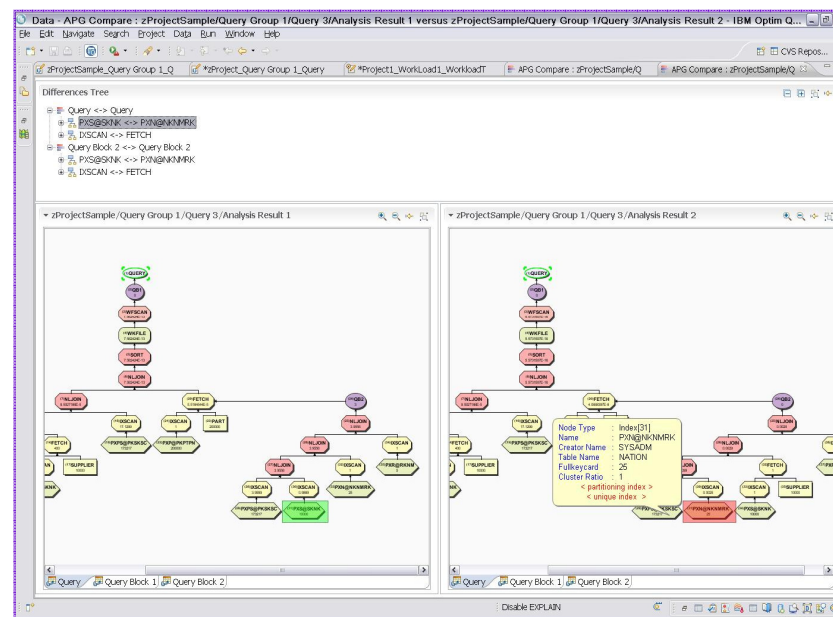
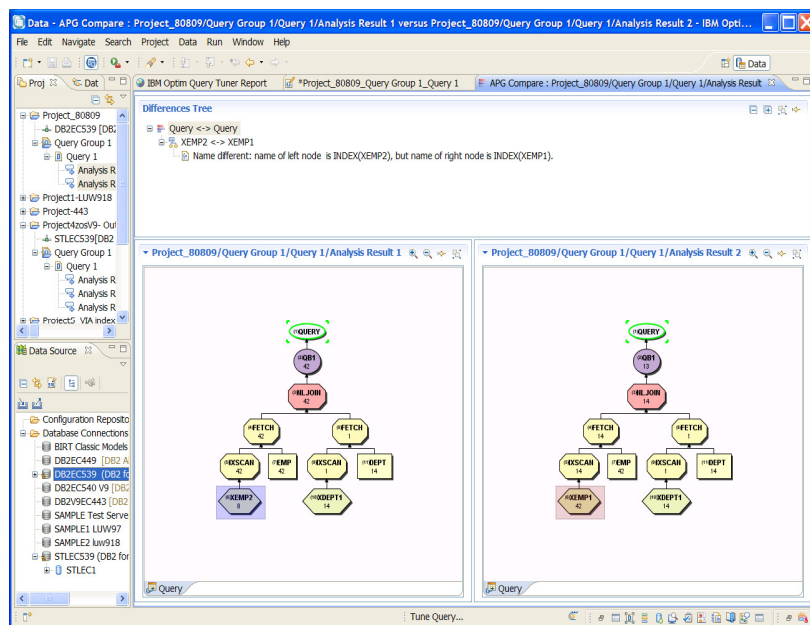
[Show DDL](#)
[Run DDL](#)
[Select All](#)
[Deselect All](#)
[Run What-If](#)

Virtual Index Definition

Define your own virtual indexes and click Run what-if to test it

Access Path Graph Comparison

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



How to Compare Access Plan Graph

The screenshot displays the IBM Optim Query Tuner Client interface. The Project Explorer on the left shows a tree view of a project named 'New Project1'. Under 'Query Group 1', two 'Analysis Result' items are selected. A context menu is open over these items, with the 'Compare Access Plan Graphs' option highlighted. The main window shows an 'Access Plan Graph' for a query, with a blue callout box pointing to the context menu. The callout box contains the text: 'In Project Explorer, select two analysis result, right click and select “Compare Access Plan Graph” in the context menu'. The bottom of the interface shows the 'Statistics Recommendation Detail' panel, which includes sections for 'Recommendations', 'Table, index, column, and column group details', 'Conflicts detail', and 'RUNSTATS Result'.

APG Comparison Results

The screenshot displays the 'Data - APG Compare' window. At the top, the title bar reads 'Data - APG Compare : New Project1/Query Group 1/Query 1/Analysis Result 1 versus New Project1/Query Group 1/Query 2/Analysis Result 1 - IBM ...'. The interface includes a menu bar (File, Edit, Navigate, Search, Project, Data, Run, Window, Help) and a toolbar. Below the toolbar, there are tabs for the compared queries.

The 'Differences Tree' on the left side is highlighted with a red box. It lists the following differences:

- Query <-> Query
 - WFSCAN <-> NLJOIN
 - WFSCAN <-> IXSCAN
 - UXO@CKOKODSP <-> SXL@PKSKOKEPDSQN
 - PXL@OKSDRFSKEPDC <-> UXO#CLOKOD
 - QB3 <-> No right related node
 - Query Block 1 <-> Query Block 1
 - Query Block 2 <-> Query Block 2
 - Query Block 3 <-> Query Block 3

The two main panes show Access Plan Graphs (APGs). The left pane, titled 'New Project1/Query Group 1/Query 1/Analysis Result 1', shows a query plan starting with a join node (25)NLJOIN 1, which branches into two IXSCAN nodes: (27)IXSCAN 380.0481 and (29)IXSCAN 400. The right pane, titled 'New Project1/Query Group 1/Query 2/Analysis Result 1', shows a different plan starting with a join node (22)QB2 1, which branches into (24)IXSCAN 1 and (28)IXSCAN 10000. Below these are nodes (25)SXL@PKSKOKEPDSQN 0 and (27)UXO#CLOKOD 25(default). A third join node (28)QB3 1 is visible at the bottom of the right pane.

List Differences in a tree. After clicking on it, related nodes are highlighted in below Access Plan Graphs

Visual Plan Hint – A example Query with Join seq. change (emp->dept to dept-> emp)

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Visual Plan Hint" and shows a "Join Graph - based on the existing access plan" and a "Join Sequence Definition Editor".

Join Graph: Shows a query plan with three tables: (2) DEPT, (1) EMP, and (3) EPROJ. The join sequence is DEPT -> EMP -> EPROJ. The join between DEPT and EMP is labeled "A. WORKDEPT=B.DEPTNO". The join between EMP and EPROJ is labeled "B.DEPTNO=C.DEPTNO".

Join Sequence Definition Editor: Shows a "Default Join Sequence" with three tables: (2) DEPT, (1) EMP, and (3) EPROJ. The join sequence is DEPT -> EMP -> EPROJ. The join between DEPT and EMP is labeled "IXSCAN". The join between EMP and EPROJ is labeled "IXSCAN".

Summary for Hint Definition: Shows a table with columns: QBLOCKNO, TABLE_CR..., TABLE_NAME, CORRELAT..., TABNO, JOIN_MET..., ACCESS_T..., ACCESS_C..., ACCESS_N..., SORTN_JOIN, SORTC_JOIN. The table contains one row for QBLOCKNO 1, TABLE_NAME EMP, CORRELAT... A, TABNO 1, JOIN_MET... NLJOIN.

QBLOCKNO	TABLE_CR...	TABLE_NAME	CORRELAT...	TABNO	JOIN_MET...	ACCESS_T...	ACCESS_C...	ACCESS_N...	SORTN_JOIN	SORTC_JOIN
1	DSN8910	EMP	A	1	NLJOIN					



Optim Query Tuner – SQL Tuning For workload

Tune Workload

- **Identify workload with performance issues**
- **Capture workload from Source in :**
 - Statement cache
 - Catalog
 - plan
 - package
 - QMF
 - QMF HPO
 - Category
 - File, text
 - Other workloads

View captured statements

The screenshot shows the 'Workload Tuning Editor' interface in the IBM Optim Query Tuner Client. The window title is 'Data - Project1_Workload Group 1_Workload_0 - IBM Optim Query Tuner Client'. The interface includes a menu bar (File, Edit, Navigate, Search, Project, Data, Run, Window, Help) and a toolbar. The main area is titled 'Workload Tuning Editor' and contains a 'Workload Statements' section. This section displays the workload name 'Workload_0', owner 'SYSADM', and status 'CAPTURED'. Below this, there are instructions and a toolbar with buttons for 'Capture', 'Workload Tools', 'Query Tools', 'Schedule', 'History', and 'Remove'. A filter name 'DEFAULT_VIEW' is shown, along with 'View', 'Customize', and 'Refresh' options. A table of captured statements is displayed, with 1-6 rows shown. The table has columns for Execution Count, Source, Accumulated Elapsed Time (sec), Average Elapsed Time (sec), Accumulated CPU Time (sec), Average CPU Time (sec), and Statement Text. The statements are all 'DECLARE' statements for cursors. A red box highlights the first six rows of the table.

Execution Count	Source	Accumulated Elapsed Time (sec)	Average Elapsed Time (sec)	Accumulated CPU Time (sec)	Average CPU Time (sec)	Statement Text
2	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR1 C
1	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR2 C
2	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR3 C
2	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR4 C
2	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR5 C
1	CATALOG	0	0	0	0	DECLARE DB2JCCCURSOR2 C

Workload Stats Advisor - Report

The screenshot displays the 'Workload Tuning Editor' window. On the left, there is a search panel with the text 'search key word in report' and a 'search key word' input field. The main area is divided into two sections:

- Statistics Advisor Detail Report:** This section, highlighted with a red box and labeled 'Detail Report', contains the following text:


```

            Statistics Advisor Detail Report
            Analysis start time: 2009-06-16 19:58:33.729513
            Analysis end time: 2009-06-16 19:59:38.32689

            -----
            TABLE SYSADM.CUSTOMER
            Table type: Table
            Cardinality: 4500000.0
            Collection time: 0001-01-01 00:00:00.0
            Statistics status: OK

            INDEXES:
            SYSADM.PXC@CKNKMS (C_CUSTKEY,C_NATIONKEY,C_MKTSEGMENT)
            First key cardinality: 4500000.0
            Full key cardinality: 4500000.0
            Data repetition factor: -1.0
            Collection time: 0001-01-01 00:00:00.0
            Statistics status: missing
            
```
- Conflicts detail:** This section, highlighted with a red box and labeled 'Conflict Report', contains the following text:


```

            TABLE SYSADM.LINEITEM
            One of the frequency records (-1.0) of the L_ORDERKEY column group is out of range [0,1]
            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 cardinality for the column group or column
            (2.2222222222222224E-8). The maximum frequency is expected to be greater than the average unless only least-frequently occurring values are being collected.
            Tolerance: 0.0010

            TABLE SYSADM.SUPPLIER
            One of the frequency records (-1.0) of the S_SUPPKEY column group is out of range [0,1]
            Tolerance: 0.0010
            
```

The bottom of the window shows a tabbed interface with 'Statistics Advisor Details' selected, and a 'Get Statistics Recommendation' button.

Workload Query Advisor summary

Summary Report

Workload Tuning Editor

Workload Query Advisor Recommendations Summary

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

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

View statements that meet the following criteria:

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

Filter Options of Statements Report

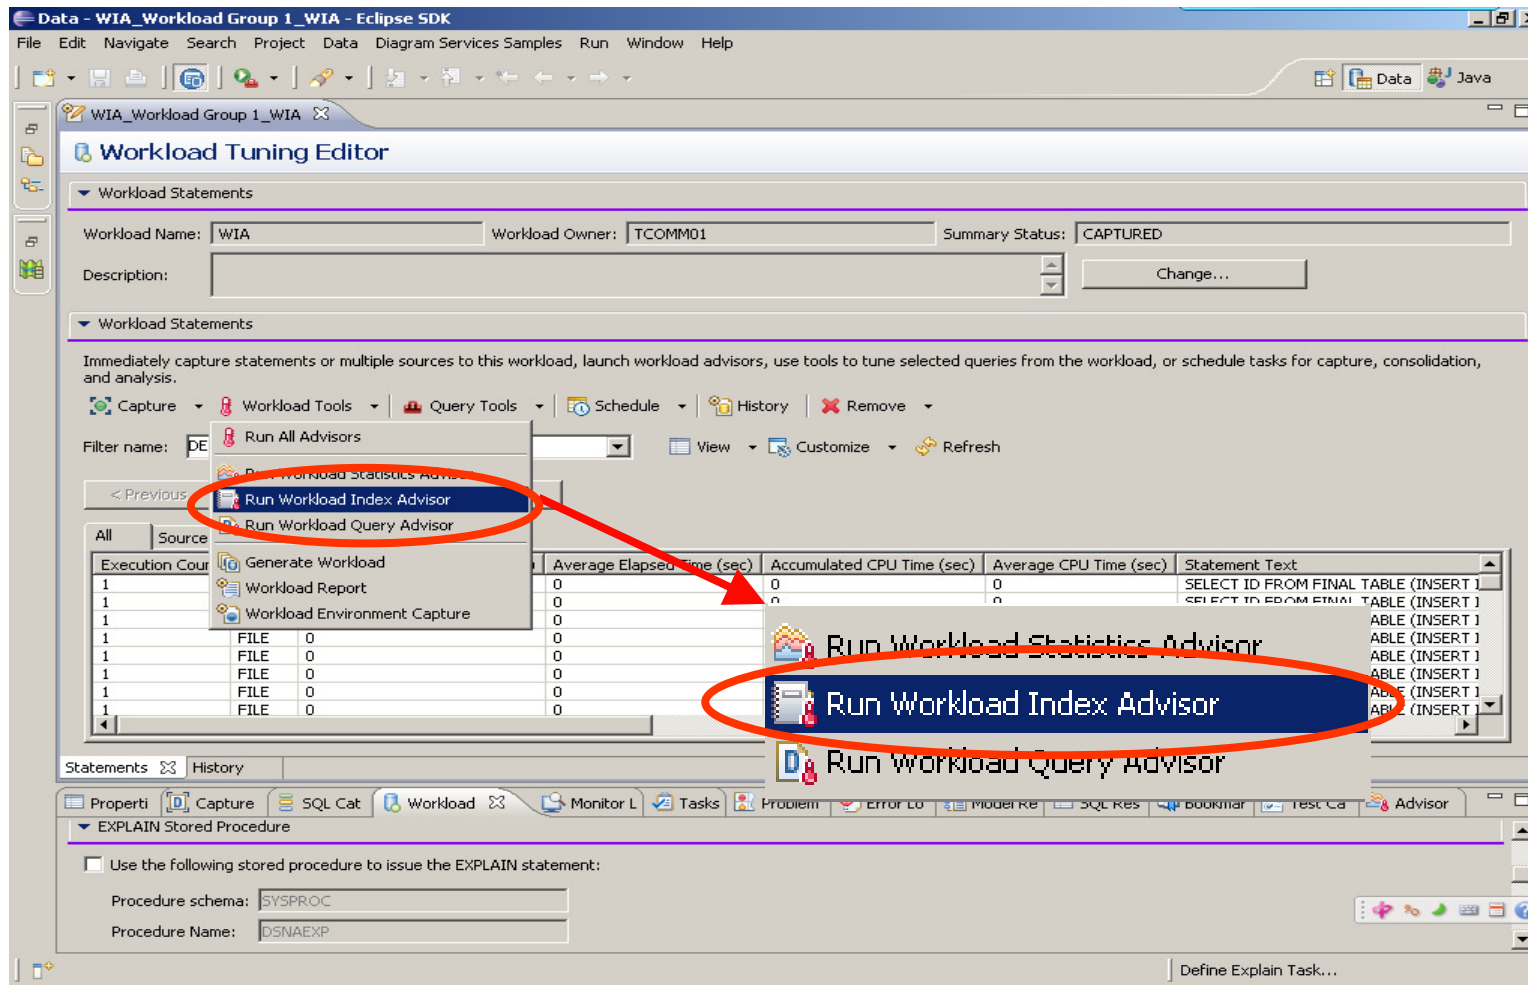
Query Advisor view statements

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Workload Tuning Editor" and shows a "Query List" tab. Below the tab, there is a section for "Workload Query Advisor Recommendations" with a descriptive text: "This table shows the filtered view of queries based on the confidence level of the recommendation and the degree of warning severity as specified in the filter." Below this text are "Details" and "Query Tools" buttons. A table of recommendations is displayed, showing columns for State, Status, Warnings, High Severity, Medium Severity, Low Severity, Execution Count, Accumulated Execution Time, and Accumulated Cost. The table contains 22 rows of data. At the bottom of the window, there are tabs for "Statements", "Advisors", "Query Advisor Summary", "Statistics Advisor", "Statistics Advisor Details", and "Query Advisor".

State...	Status	Warnings	High Severi...	Medium Se...	Low Severit...	Execution Count	Accumulated E...	Accumu
3530	Warnings w...	4	0	0	4	5	105.00051	0.06765
3531	Warnings w...	1	0	0	1	5	114.31219	2.25334
3547	Warnings w...	1	0	0	1	13	354.38287	6.41448
3540	Warnings w...	1	0	0	1	11	249.98885	4.91882
3541	No Warning	0	0	0	0	11	208.49591	0.15119
3544	No Warning	0	0	0	0	11	36.727528	2.33960
3542	No Warning	0	0	0	0	11	209.34059	0.08491
3548	No Warning	0	0	0	0	16	539.7503	46.9283
3546	No Warning	0	0	0	0	13	244.81737	1.26801
3545	No Warning	0	0	0	0	11	206.56808	0.05985
3553	No Warning	0	0	0	0	16	0.50924844	0.17154
3552	No Warning	0	0	0	0	16	299.74493	0.08627
3550	No Warning	0	0	0	0	16	11.583381	0.78896
3549	No Warning	0	0	0	0	16	13.630047	0.37115
3539	No Warning	0	0	0	0	11	191.87997	14.2829
3538	No Warning	0	0	0	0	9	170.33287	0.12218
3537	No Warning	0	0	0	0	9	81.64089	0.31584
3536	No Warning	0	0	0	0	5	61.39365	2.39194
3535	No Warning	0	0	0	0	5	497.97882	0.60133

Workload Index Advisor (WIA)

– SQL Tuning for workload indexes



WIA with Index recommendations

The screenshot displays the Workload Tuning Editor interface. At the top, the title bar reads "Data - WIA Workload Group 1_WIA - Eclipse SDK". The main window is titled "Workload Tuning Editor" and contains a section for "Workload Index Advisor Recommendations".

Text in the recommendations section states: "The following information shows the index recommendations for this workload. Y recommendations are applied. There is the option to run index analysis again wit recommendations." (Note: 'wit' is likely a typo for 'with').

Two summary metrics are highlighted with red boxes and connected by a red arrow:

- Estimated performance improvement: 38 %
- Disk space required(DASD space): 60.78 MB

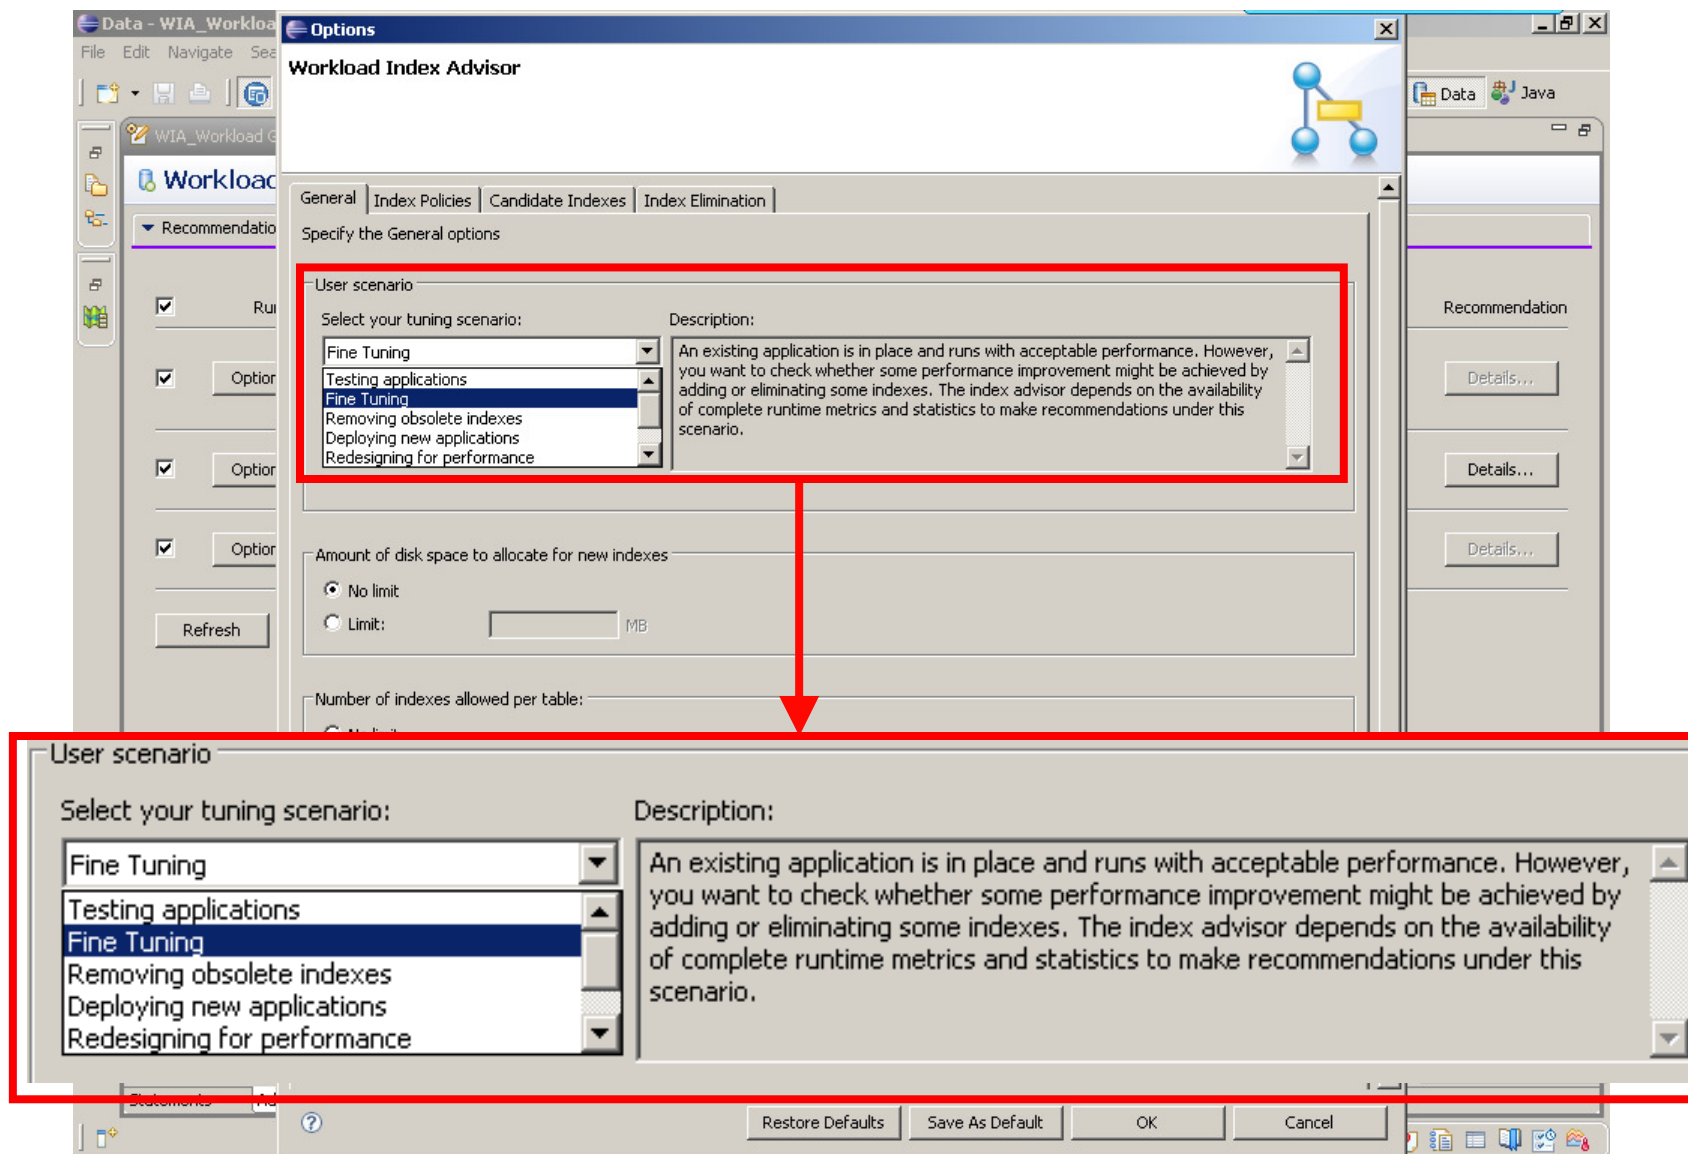
Below these metrics is a table of recommendations:

Feature Details	Action	Object N...	Columns
<input checked="" type="checkbox"/> DSN_WIA_QBLOCK			
<input checked="" type="checkbox"/> Index	Create	DSN_WIA...	STMT_ID(ASC)
<input checked="" type="checkbox"/> Index	Create	DSN_WIA...	SESSION_ID(ASC)
<input checked="" type="checkbox"/> DSN_DET_COST_TABLE			
<input checked="" type="checkbox"/> Index	Create	DSN_DET...	QUERYNO(ASC), EXPLAIN...
<input checked="" type="checkbox"/> DSN_STATEMENT_TABLE			
<input checked="" type="checkbox"/> Index	Create	DSN_STA...	QUERYNO(ASC), EXPLAIN...
<input checked="" type="checkbox"/> PLAN_TABLE			
<input checked="" type="checkbox"/> Index	Create	PLAN_TAB...	QUERYNO(ASC), BIND_TI...
<input checked="" type="checkbox"/> Index	Create	PLAN_TAB...	QUERYNO(ASC), BIND_TI...
<input checked="" type="checkbox"/> Index	Create	PLAN_TAB...	QUERYNO(ASC), BIND_TI...
<input checked="" type="checkbox"/> SYSINDEXPART			
<input checked="" type="checkbox"/> Index	Create	SYSINDEX...	PARTITION(ASC)

To the right of the table are several buttons: "Show DDL...", "Show Related SQL...", "What-If Analysis...", "Run DDL...", "Select All", and "Deselect All".

At the bottom of the interface, a blue callout box contains the text: "Recommend indexes to improve workload performance".

WIA – New User scenarios (New in V2.2)



Workload index advisor recommendations

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Workload Tuning Editor" and shows index recommendations for a workload. The estimated performance improvement is 69% and the disk space required is 673.08 MB. A table of recommendations is shown below.

Feature Details	Action	Object...	Columns	Estimated Disk Space
<ul style="list-style-type: none"> Index 	Create	PART_V...	P_PARTKEY(ASC), P...	17.1953125 M
<ul style="list-style-type: none"> Index 	Create	PARTSU...	PS_SUPPKEY(ASC)	
<ul style="list-style-type: none"> Index 	Create	PARTSU...	PS_SUPPLYCOST(AS...	
<ul style="list-style-type: none"> Index 	Create	LINEITE...	L_SHIPMODE(ASC), ...	
<ul style="list-style-type: none"> Index 	Create	LINEITE...	L_QUANTITY(ASC), L...	
<ul style="list-style-type: none"> Index 	Create	LINEITE...	L_RETURNFLAG(ASC...	
<ul style="list-style-type: none"> Index 	Create	LINEITE...	L_SUPPKEY(ASC), L...	

A "DDL Details" dialog box is open, showing the SQL DDL for the selected recommendation:

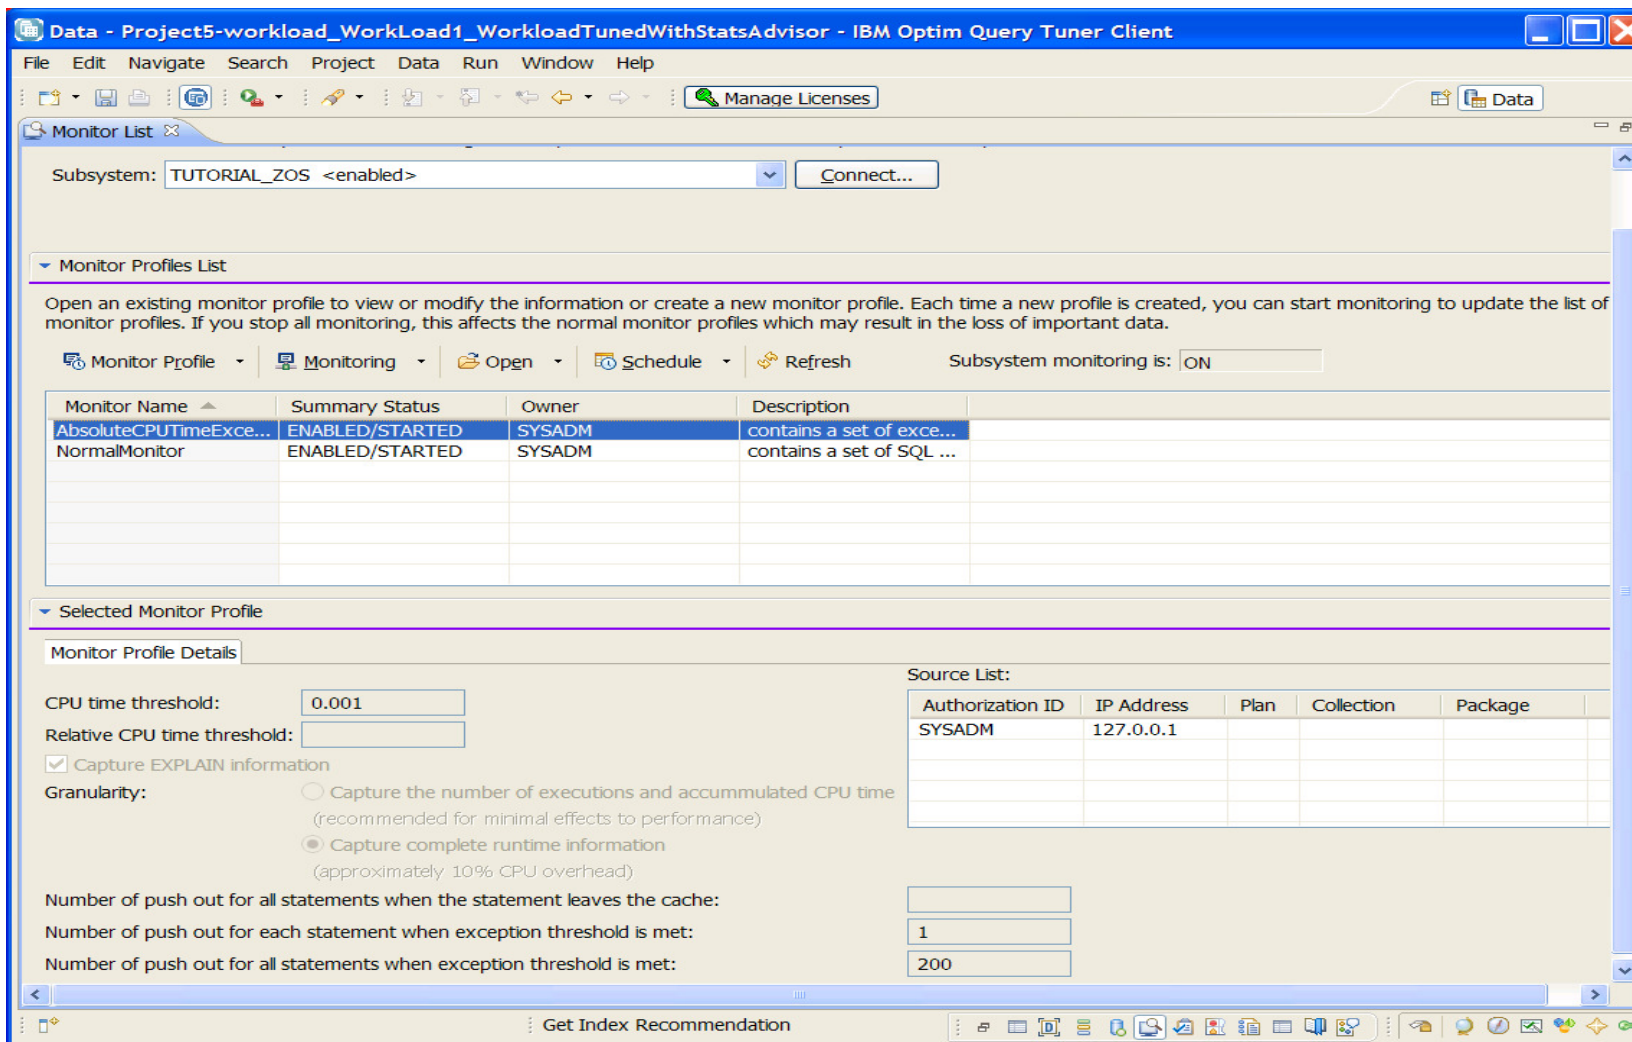
```

CREATE INDEX DB2OE.LINEITEM_VIRT_IDX_115953554645 ON
SYSADM.LINEITEM ( L_SHIPMODE ASC, L_ORDERKEY ASC) FREEPAGE 0
PCTFREE 10;

CREATE INDEX DB2OE.LINEITEM_VIRT_IDX_1159535563807 ON
SYSADM.LINEITEM ( L_QUANTITY ASC, L_EXTENDEDPRICE ASC, L_TAX
ASC) FREEPAGE 0 PCTFREE 10;

CREATE INDEX DB2OE.LINEITEM_VIRT_IDX_1159535564028 ON
SYSADM.LINEITEM ( L_RETURNFLAG ASC, L_LINESTATUS ASC,
L_EXTENDEDPRICE ASC, L_TAX ASC, L_DISCOUNT ASC, L_QUANTITY
ASC, L_SHIPDATE ASC) FREEPAGE 0 PCTFREE 10;
    
```


Monitor workload for CPU Time



The screenshot displays the IBM Optim Query Tuner Client interface. The window title is "Data - Project5-workload_WorkLoad1_WorkloadTunedWithStatsAdvisor - IBM Optim Query Tuner Client". The interface includes a menu bar (File, Edit, Navigate, Search, Project, Data, Run, Window, Help) and a toolbar with icons for file operations and a "Manage Licenses" button. The main area is divided into several sections:

- Monitor List:** Shows a dropdown for "Subsystem: TUTORIAL_ZOS <enabled>" and a "Connect..." button. Below it is a "Monitor Profiles List" section with a descriptive paragraph and a toolbar with icons for "Monitor Profile", "Monitoring", "Open", "Schedule", and "Refresh". A "Subsystem monitoring is: ON" checkbox is also present.
- Monitor Profiles List Table:**

Monitor Name	Summary Status	Owner	Description
AbsoluteCPUTimeExce...	ENABLED/STARTED	SYSADM	contains a set of exce...
NormalMonitor	ENABLED/STARTED	SYSADM	contains a set of SQL ...
- Selected Monitor Profile:** Contains a "Monitor Profile Details" section with various configuration options:
 - CPU time threshold: 0.001
 - Relative CPU time threshold: (empty field)
 - Capture EXPLAIN information
 - Granularity:
 - Capture the number of executions and accumulated CPU time (recommended for minimal effects to performance)
 - Capture complete runtime information (approximately 10% CPU overhead)
 - Number of push out for all statements when the statement leaves the cache: (empty field)
 - Number of push out for each statement when exception threshold is met: 1
 - Number of push out for all statements when exception threshold is met: 200
- Source List Table:**

Authorization ID	IP Address	Plan	Collection	Package
SYSADM	127.0.0.1			

The bottom of the window features a status bar with a "Get Index Recommendation" button and a standard Windows taskbar with various system icons.

Tuning a workload monitored into a project

The screenshot displays the IBM Optim Query Tuner Client interface. The main window is titled "Workload Tuning Editor" and shows details for a workload named "AbsoluteCPUTimeExceptionMonitor". The workload owner is "SYSADM" and the status is "ENABLED/STARTED". The description states: "contains a set of exception SQL statements that satisfy the predefined".

Below the description, there are several tool buttons: Capture, Workload Tools, Query Tools, Schedule, History, and Remove. The filter name is set to "DEFAULT_VIEW".

A table titled "Monitor" displays the following data:

Execution Count	Source	Accumulated Elapsed Time (sec)	Average Elapsed Time (sec)	Accumulated CPU Time (sec)	Average CPU Time
5		47.044178	9.408835	4.14425	0.82885
5		6.865626	1.373125	0.112362	0.022472
16		0.219881	0.013743	0.036937	0.002309
16		0.32213	0.020133	0.01921	0.001201

The interface also includes a left-hand navigation pane with a tree view showing project structures like "Project5-workload" and "Query Tuner Sample WorkLoad1". The bottom status bar includes a "Get Index Recommendation" button.




***Acceleration of SQL Performance and better
application development with OMPE and
Optim Development Studio***

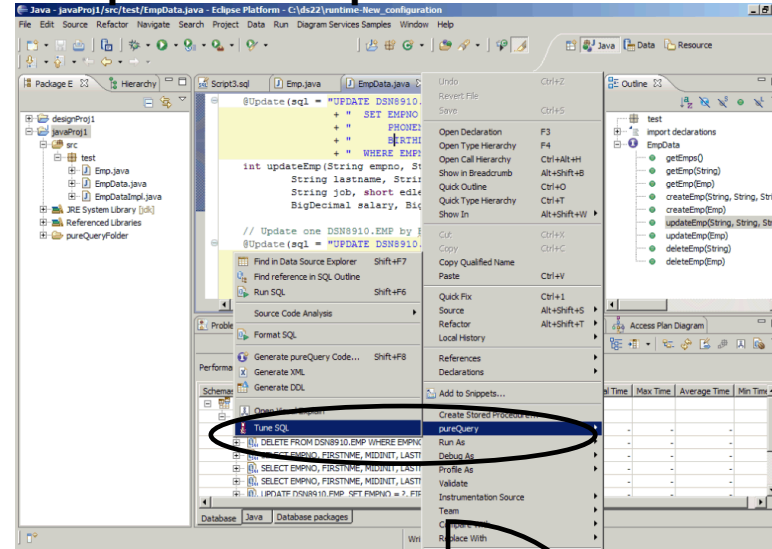
- A peek at product integration support*

Accelerate SQL Performance with OPME and Optim Development Studio

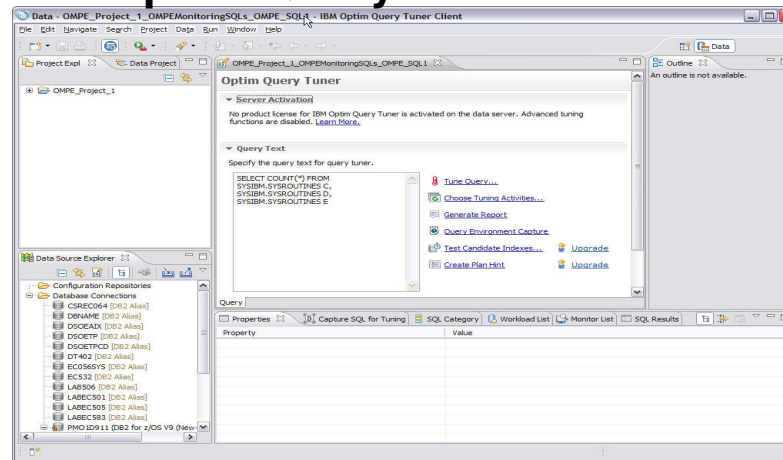
■ Omegamon XE (z)



■ Optim Development Studio



■ Optim Query Tuner



- OPME launch
- OQT to tune
- SQL performance

- ODS Launch OQT
- For application tuning

Product Integration Support:

- OMPE Launch Query Tuner for SQL Performance tuning

The screenshot displays the IBM Optim Query Tuner Client interface. On the left, a window titled 'D911 - Statistics Details' shows the execution statistics for a query. The SQL statement being analyzed is:

```
select count(*) from sysibm.sysroutines c, sysibm.sysroutines d, sysibm.sysroutines e
```

The statistics table below shows various performance metrics:

Property	Value	Property	Value
Users	0	Lock and Latch Wait Time	0.000000
Copies	0	Global Locks Wait Time	0.000000
Status	Active	Reads by other Threads Wait Time	0.000000
Insert	6/3/09 9:24:04 PM	Writes by other Threads Wait Time	0.000000
Executions	2	Program Name	SQLC2E07
Synchronous Read	6	Table Name	SYSROUTINES
Synchronous Write	0	Dynamicrules Bind	RUN
Getpage	168,928	Current Rules	DB2
Rows examined	0	DB2 Limit Exceeded	0
Rows examined / Executions	0	Storage Limit Exceeded	0
Rows processed	2	Statement Line Number	210
Rows processed / Executions	1	Transaction	db2bp.exe
Sort	0	User ID	MTS

The main window, 'Data - OMPE_Project_1_OMPEMonitoringSQLs_OMPE_SQL1 - IBM Optim Query Tuner Client', shows the 'Optim Query Tuner' interface. It includes a 'Server Activation' section with a message: 'No product license for IBM Optim Query Tuner is activated on the data server. Advanced tuning functions are disabled. Learn More.' Below this is the 'Query Text' section, which contains the same SQL statement as shown in the statistics window. The interface also features a 'Data Source Explorer' on the left and a 'Query' section at the bottom with various analysis options like 'Capture SQL for Tuning', 'SQL Category', 'Workload List', 'Monitor List', and 'SQL Results'.

Product Integration ODS (with pureQuery) -> QT

ODS launch Query Tuner for (PQ) Tuning

The screenshot displays the Eclipse IDE interface. The main editor shows a SQL script with a context menu open over a query. The 'Tune SQL' option is selected, revealing a sub-menu where 'pureQuery' is highlighted. The background shows the project structure and the 'Outline' view.

Database

Time	Max Time	Average Time	Min Time
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-
-	-	-	-

Optim Query Tuner V2.2 – What's new?

- **Optim Query Tuner V2.2**
 - *Made available in June/July, 2009*
 - *IBM Optim Query Tuner for DB2 for z/OS*
 - *Single Query based*
 - *IBM Optim Workload Tuner for DB2 for z/OS*
 - *Single query and workload SQL tuning*
 - *IBM Optim Query Tuner for DB2 LUW*
 - *Support Single Query SQL Tuning tools and advisors*

IBM Optim Query Tuner V2.2 New Feature Highlights

- ***Usability and workflow enhancements***
 - *Integrated look and feel when evaluating advisor recommendations*
 - *Expanded project management*
 - *provides the ability to group statements*
 - *maintain versions of statements and analysis results*
 - *Common DB connection and management*
 - *Ability to capture and analyze SQL from*
 - *Code editors*
 - *Routine editor*
 - *Integrated Query Editor (IQE)*
 - *pureQuery Outline*
 - *Data Source Explorer*
 - *Packages, UDFs ,Stored Procedures, Triggers and Views*

IBM Optim Query Tuner V2.2 New Feature Highlights...

- **SQL tuning Tools and Advisors enhancements**
 - *Index Advisors*
 - *Ability to maintain user-defined virtual indexes without invoking the Index Advisor*
 - *New predefined scenarios for Workload Index Advisor Recommendations*
 - *New and enhanced rules for Index Advisor recommendations*
 - *Statistics Advisor*
 - *New and enhanced rules for Statistics advisor recommendations*
 - *New Statistics profile function*
 - *ability to save implemented RUNSTATS recommendations to a user-created profile table for later recall*
 - *Access Plan Graph to present the use of virtual indexes and its statistics*
 - *Access plan graph comparison*
 - *Ability to compare APG between two statements or tuning analysis results before and after*
 - *Integrated advisor recommendations into APG*

IBM Optim Query Tuner V2.2 New Feature Highlights...

- **Provide performance enhancement for**
 - Workload SQL capturing
 - Workload Advisors
- **Support product integration with**
 - *IBM Optim Development Studio (ODS)*
 - *OMEGAMON XE(z) for DB2 PM/PE (APAR PK88290)*
 - With ability to launch Query Tuner during application development (with pureQuery) by ODS and SQL performance monitoring for further performance analysis by OMPE.

Optim Query Tuner – The Step Beyond

- *More usability enhancements*
- *More performance enhancements*
- *More enhancements on advisor recommendations*
 - *provide more best-practice rule for better application development*
 - *for both novice and experienced DBAs and developers*
- *More workload SQL tuning features/functions*
- *More product integration support*

Summary

- *IBM Optim Query Tuner*
 - *Made substantial enhancements in V2.2 for DB2 z*
 - *Provide Single Query Tuning tools and advisors for DB2 LUW (V2.2 new)*
- *Provides Key Business Value to*
 - Increase quality of service by proactively solving problems before they occur
 - Accelerate problem analysis/resolution by providing expert guidance and decreasing reliance on special skills
 - Aggregate performance data in a query warehouse to readily see and track the affect of changes over time
 - **Query Workload Tuner can:**
 - Define and analyze workloads to proactively optimize physical database design on indexes
 - Provide SQL workload and query tuning advice to maximize application performances and reduce the total cost of ownership
 - Eliminating the needs for costly manual analysis of large applications with huge number of SQL in a large workload



Optim Query Tuner – Feature offering

Backup charts

IBM Optim Query Tuner key functions

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

• *DB2 for z/OS V9.1 NFM

• ** functions for DB2 for z/OS

IBM Optim Solutions

- **Optim Solutions Page:**
 - <http://www.ibm.com/software/data/optim/>

- **IBM Integrated Data Management (Optim and Data Studio):**
 - <http://www.ibm.com/developerworks/spaces/optim>
 - Tutorials
 - Downloads
 - [Forums](#) / Blogs
 - Join the community!

Important 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, OR SHALL HAVE THE EFFECT OF:

- CREATING ANY WARRANTY OR REPRESENTATION FROM IBM (OR ITS AFFILIATES OR ITS OR THEIR SUPPLIERS AND/OR LICENSORS); OR
- ALTERING THE TERMS AND CONDITIONS OF THE APPLICABLE LICENSE AGREEMENT GOVERNING THE USE OF IBM SOFTWARE.

THANK
YOU